

Article

Not peer-reviewed version

Why Does Life Exist?

Michael Timothy Bennett *

Posted Date: 12 March 2026

doi: 10.20944/preprints202603.0203.v2

Keywords: origins of life; ool; stack theory; functional information; weakness



Preprints.org is a free multidisciplinary platform providing preprint service that is dedicated to making early versions of research outputs permanently available and citable. Preprints posted at Preprints.org appear in Web of Science, Crossref, Google Scholar, Scilit, Europe PMC.

Copyright: This open access article is published under a [Creative Commons CC BY 4.0 license](https://creativecommons.org/licenses/by/4.0/), which permit the free download, distribution, and reuse, provided that the author and preprint are cited in any reuse.

Disclaimer/Publisher's Note: The statements, opinions, and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions, or products referred to in the content.

Article

Why Does Life Exist?

Michael Timothy Bennett

School of Computing, The Australian National University, Canberra, ACT, Australia; m@michaeltimothybennett.com

Abstract

Functional information measures how rare functional configurations are. Wong and colleagues argue that selection should drive a *law of increasing functional information*. This is often read as a claim that complexity must increase. Here I give a different interpretation, which is that survivors tend to be the systems that did not overcommit. I model a system as a policy π , meaning a bundle of commitments expressed in a finite embodied vocabulary. New selection pressures arrive as a set of future requirements drawn from the unobserved outcome set U . A currently viable policy leaves an unobserved buffer $B_\pi \subseteq U$ of outcomes it still permits. Under a maximally ignorant novelty model, the survival probability of π is exactly $2^{|B_\pi| - |U|}$. Under any exchangeable novelty prior, survival remains monotone in $|B_\pi|$. So persistence under novelty favours weak policies, where weakness counts the compatible completions left open. I define degree of future function as survival probability and functional information as Hazen and Szostak rarity within the currently viable set. Conditioning on persistence reweights the viable set toward larger buffers and therefore toward higher functional information. This yields a mathematical analogue of the proposed law under explicit assumptions. Supplementary analysis quantifies how much structured novelty is needed before that buffer size ordering can reverse. In fully enumerated toy worlds, weakness maximisation improves mean log survival probability by 1.674 bits relative to random choice. Weakness and simplicity are not the same thing. Weakness helps a system persist under novelty, because it keeps more futures compatible. Simplicity helps a system persist because there is less to break, which obviates the need for repair. Complexity requires self-repair to persist, increasing weakness. Life is persistent complexity. In between complex life and simple nonlife is *the void of the unviable*: complexity which is not alive.

Keywords: origins of life; ool; stack theory; functional information; weakness

Significance Statement

The current PNAS debate over increasing functional information has become entangled with claims about complexity. In a minimal mathematical setting that captures the core logic of the proposed law, I show that the two can be separated. Under explicit assumptions about novelty, persistence favours policies that keep more untested possibilities compatible. Functional information rises among survivors because survival reweights the viable set toward weaker policies. The only invariant is not simplicity but weakness, meaning how many further commitments remain available while staying correct. A separate damage and repair illustration shows why complexity matters through a different mechanism entirely: breakage and repair, not future compatibility. The paper therefore clarifies what the law does and does not say.

Main Text

The Law That Is Not About Complexity

Wong and colleagues propose that evolving systems tend to increase functional information under selection [1]. A recent exchange frames this as a question about complexity [2,3]. Assembly theory offers a parallel attempt to quantify selection in evolving systems through combinatorial complexity [4], and related information centric accounts have tried to understand life and selection

without collapsing everything into raw structural complication [5–7]. I take a different view. Selection does not need to reward complexity. Selection needs to reward survival, and evolution favours keeping one's options open.

My result is a survivorship bias theorem. Conditioning a population on persistence reweights it toward policies that stayed compatible with whatever the future demanded. That reweighting increases expected functional information inside the viable set. Importantly, it does not require any monotone increase in complexity. McShea and Brandon propose a zero force evolutionary law in which diversity and complexity increase in the absence of selection and other forces [8]. The present result is different: it gives a law of persistence filtering under novelty, closer in spirit to Ashby's cybernetic point that coping with a wider range of future disturbances requires sufficient compatible variety [9]. The intended logic of the paper is therefore:

1. Novelty selection acts on persistence under future requirements that have not yet been observed.
2. Under an exchangeable novelty prior, persistence is monotone in unobserved buffer size, hence in constraint weakness.
3. Inside the currently viable set, that makes weakness a monotone predictor of future function. Functional information rises among survivors without any requirement that complexity increase.
4. But complexity faces a separate threat: damage. Complex systems have more to break, and persist only if they can repair themselves. Simple systems persist by having little to break.
5. Between simple nonlife and complex life is the void of the unviable: complexity that is not alive because it is too complex to persist without repair and not alive enough to repair itself. Life is what crosses that void. It is persistent complexity, maintained by self-repair sufficient to balance the cost of having many parts.

Survival is a problem of generalising beyond what you have already been tested on. Generalisation is only possible because the future is not fully specified by the past. That means there is an unobserved region that still contains many possible constraints. A system persists when later constraints land inside what it already left compatible.

The core claim of this paper is that if you do not know what the future will demand, you should not overcommit. Embodiment can be represented as a formal language of constraints [10]. Related constraint first accounts likewise emphasise that what persists depends on how viable organisations maintain compatible possibilities [5,6,11]. In that formal language the right target for minimising commitment is constraint weakness. Maximising weakness means maximising the size of the completion set left open by your current commitments.

A Minimal Model of Commitments and Novelty

I use a lightweight simplified version of Stack Theory [12], beginning with a finite embodied language L . I use Stack Theory here because the generalisation-optimal learning model associated with it are useful for this paper. The vocabulary is finite because any embodied system occupies a bounded region and therefore has finite information capacity [13]. L is a set of implementable bodily outcomes a body can make by interacting with its environment¹. Asserting a particular outcome is asserting a constraint on possible worlds. Think of L as a set of partial specifications of possible worlds in which a particular body exists.

Policies and Completions

Asserting an outcome rules some of the other outcomes in and rules some other outcomes out, because it rules in or out possible worlds. For example, if a human raises their arm, then all possible

¹ These bodily outcomes are usually called "statements" in related work, but here that terminology is avoided because it adds unnecessary complication.

worlds from that point in time on begin with that person's arm raised. A completion of an outcome is any other outcome that implies the first. Assume a statement π . Its completion set is

$$\text{Ext}(\pi) := \{y \in L \mid \pi \subseteq y\}.$$

Its weakness is

$$w(\pi) := |\text{Ext}(\pi)|.$$

To embody a particular outcome and constrain possible worlds is to adopt a policy. If π represents your present embodied commitments or policy, then $\text{Ext}(\pi)$ is everything compatible with those commitments. Weakness is the number of compatible completions. Looser constraints mean more possibilities. Higher weakness means fewer promises, or more degrees of freedom remaining. The completion set can also be interpreted as a formal counterpart of Kauffman's adjacent possible [7,14], adapted here from one-step reachability to logical compatibility with current commitments. Weakness measures its size.

Observed and Unobserved Outcomes

Let α be the set of outcomes you have already been tested on (a history of outcomes you have survived). Write the unobserved region as

$$U := L \setminus \alpha.$$

For any policy π , define its unobserved buffer

$$B_\pi := \text{Ext}(\pi) \cap U.$$

Interpretation. U is everything that could matter later but has not mattered yet. B_π is the part of your option volume that lives in the future. This is where surprises arrive.

Novelty as Future Requirements

A novelty set is a subset $S \subseteq U$ of unobserved outcomes that become required later. A novelty prior is a distribution P over subsets $S \subseteq U$. I call P exchangeable when it depends only on $|S|$ and not on which elements of U are in S . This usage is consistent with the standard notion of exchangeability in probability [15].

Interpretation. Exchangeable means you do not know which surprises will matter, only how many of them show up. The position that true novelty is not prestatable aligns with arguments that biological evolution is not governed by entailing laws [16].

Survival Probability Equals Option Volume

Assume π has not yet failed on the observed set α , meaning $\alpha \subseteq \text{Ext}(\pi)$. Under novelty set S , the policy survives if it did not preclude any newly required outcome. In this model that condition is

$$S \subseteq B_\pi.$$

Interpretation. You die when the future demands something you forbade with your embodied policy.

Proposition 1 (Uniform ignorance survival law). *Assume a maximally uninformative novelty prior where every subset $S \subseteq U$ is equally likely. Then the survival probability of a currently correct policy π is*

$$\mathbb{P}(\pi \text{ survives}) = \frac{2^{|B_\pi|}}{2^{|U|}} = 2^{|B_\pi| - |U|}.$$

Interpretation. Under complete ignorance, you survive in proportion to how much untested future you left compatible. Every extra compatible unobserved outcome doubles your survival probability. There is no complexity term in this expression.

Proposition 2 (Exchangeable monotonicity). *For any exchangeable novelty prior P , the survival probability $\mathbb{P}_P(S \subseteq B_\pi)$ is a nondecreasing function of $|B_\pi|$.*

Interpretation. If your only information is how many surprises arrive, then leaving more room for surprises cannot hurt. In set terms, survival is the event $S \subseteq B_\pi$. A larger buffer means more possible novelty sets are contained in it. This is a novelty analogue of survival of the flattest [17], where genotypes on broader fitness plateaus outcompete faster replicators at high mutation rates. Here the “flatness” is the size of the unobserved buffer rather than the density of neutral neighbours in sequence space. The result also connects to bet-hedging in evolutionary biology, where strategies that reduce fitness variance can invade populations of higher-mean competitors in unpredictable environments [18]. Weakness maximisation can be read as a deterministic limit of that logic.

From Option Volume to Functional Information

Hazen and Szostak define functional information as

$$I(F) = -\log_2\left(\frac{M(F)}{N}\right),$$

where N is the total number of configurations and $M(F)$ is the number of configurations with degree of function at least F [19,20].

I apply that definition inside the currently viable set. Fix the observed outcomes α and let Π be the set of policies that remain compatible with α , meaning $\alpha \subseteq \text{Ext}(\pi)$ for every $\pi \in \Pi$. Define the degree of future function of a policy by its survival probability under the novelty prior

$$F(\pi) := \mathbb{P}_P(S \subseteq B_\pi).$$

Define

$$M(\pi) := \{\pi' \in \Pi \mid F(\pi') \geq F(\pi)\}, \quad N := |\Pi|, \quad I(\pi) := -\log_2\left(\frac{|M(\pi)|}{N}\right).$$

Interpretation. A policy has high future function if it is likely to stay correct when new requirements appear. Functional information measures how rare that level of future function is among the policies that are currently viable.

Under exchangeability, $F(\pi)$ is a nondecreasing function of $|B_\pi|$. So inside Π , larger unobserved buffers imply weakly higher functional information, with strict increase whenever the novelty prior assigns different survival probabilities to the corresponding buffer sizes. Because α is fixed and $\alpha \subseteq \text{Ext}(\pi)$ for $\pi \in \Pi$, this gives $|\text{Ext}(\pi)| = |\alpha| + |B_\pi|$, so this is equivalent to higher weakness up to an additive constant [12].

A Formal Law of Increasing Functional Information

Now consider a population of viable policies sampled uniformly from Π . Selection acts by persistence (i.e. something that does not persist is not selected). A policy survives to the next time step exactly when it survives the novelty set.

Theorem 1 (Law of increasing functional information under novelty selection). *Assume an exchangeable novelty prior P , and suppose the survival event has positive probability under the joint law of π and S . Draw π uniformly from the viable set Π . Condition on survival, meaning on the event $S \subseteq B_\pi$. Then*

$$\mathbb{E}[I(\pi) \mid \pi \text{ survives}] \geq \mathbb{E}[I(\pi)],$$

with strict inequality whenever $F(\pi)$ is not constant over Π .

Interpretation. Selection for persistence reweights the population toward policies that keep more untested possibilities open. Those policies are rarer among viable policies. So the average functional information of survivors increases.

A proof is given in the Supplementary Information. The main step is an inequality: when a population is reweighted by a factor that increases with a trait, the conditional mean of any monotone transform of that trait increases. This reweighting argument is a special case of the covariance structure captured by the Price equation [21], restricted here to a single round of persistence selection. In the present setting the transmission term vanishes because I condition on a single selection event. What is generic is the reweighting step. What is substantive is the identification of future function with survival probability and the monotone link between survival and buffer size under exchangeable novelty. Supplementary analysis quantifies how far novelty must depart from exchangeability before the buffer size ordering can reverse.

Experiments in Finite Worlds

I tested the novelty claim in fully enumerated toy worlds where every outcome can be counted. The theorem itself is exact, so the toy worlds do not supply an independent proof. They illustrate the size of the effect under the stated model. Each world is a finite catalogue L of 48 implementable outcomes. Each outcome is a binary feature vector of length 10.

A current function fixes 2 base commitments that every viable policy must satisfy. I draw 4 observed outcomes α from the outcomes that satisfy those base commitments. Because α is small, some extra feature values look constant by accident. A policy can commit to those accidents and still pass the observed tests. That is overcommitment. It shrinks the unobserved buffer B_π .

I compare two selection rules among policies that satisfy the base commitments and remain compatible with the observed outcomes. The weakness rule chooses a viable policy that maximises $|B_\pi|$. The baseline chooses uniformly at random from the viable set.

Interpretation. Weakness maximisation tries to keep options open. Random does not try.

Under uniform novelty, the counting law gives $\log_2 \mathbb{P}(\text{survive}) = |B_\pi| - |U|$ exactly. So this experiment is noise free. Figure 1 summarises the mean \log_2 survival probability under uniform novelty.

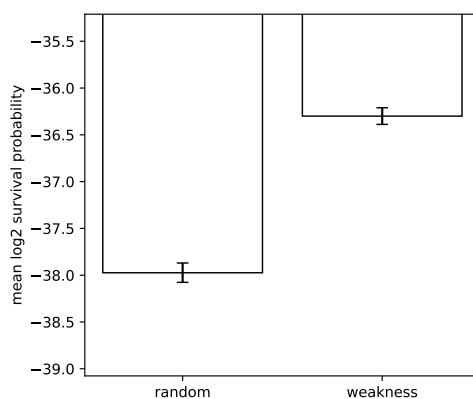


Figure 1. Weakness selection improves survival under novelty. Bars show mean \log_2 survival probability under a uniform novelty prior across 4000 trials in 200 toy worlds, with 95% bootstrap confidence intervals. Weakness selection is Bayes optimal in this setting and outperforms random choice by 1.674 bits.

In the uniform novelty setting, weakness selection improves mean log survival probability by 1.674 bits relative to random choice.

Simplicity Avoids Damage and Complexity Needs Repair

The theorem is about novelty selection. Real systems also face wear and tear. If you have more moving parts, you have more ways to fail. This is where everyday intuition about simplicity is right. A rock survives for boring reasons. It has little to break.

In Stack Theory terms, this is a very different axis from weakness. Weakness counts how many futures you still permit [12]. Simplicity in the physical sense reduces the number of damageable parts. A system can persist by being simple enough that repair is unnecessary for persistence. A system can also persist by being complex and continuously repairing itself. Thermodynamic accounts of self-replication emphasise that persistence of complex structure requires continuous dissipation [22], and the repair model here captures the discrete combinatorial analogue of that requirement. All else being equal a rock that self-repairs is more persistent than a rock that does not self-repair. However self-repair has a cost in complexity, which makes the self-repairing rock less likely to persist. This divides the world into objects which are simple, and objects which are complex but viable because their complexity is balanced by effective self-repair. In between these two clusters is *the void of the unviable*: complexity which is not alive. This is a void as described in [23].

Experiment E is a minimal damage model that illustrates this distinct axis of persistence [24]. Here the relevant state variables are the number of essential components and the repair budget, not the unobserved buffer B_π and not the novelty prior from Theorem 1. A system requires k essential components to remain viable. Each component fails independently with probability $p = 0.020$ per time step. Repair can restore at most $r = 1$ failed components per step. I compare a simple system with $k = 3$ and no repair to a complex system with $k = 12$, with and without repair. Figure 2 shows survival curves.

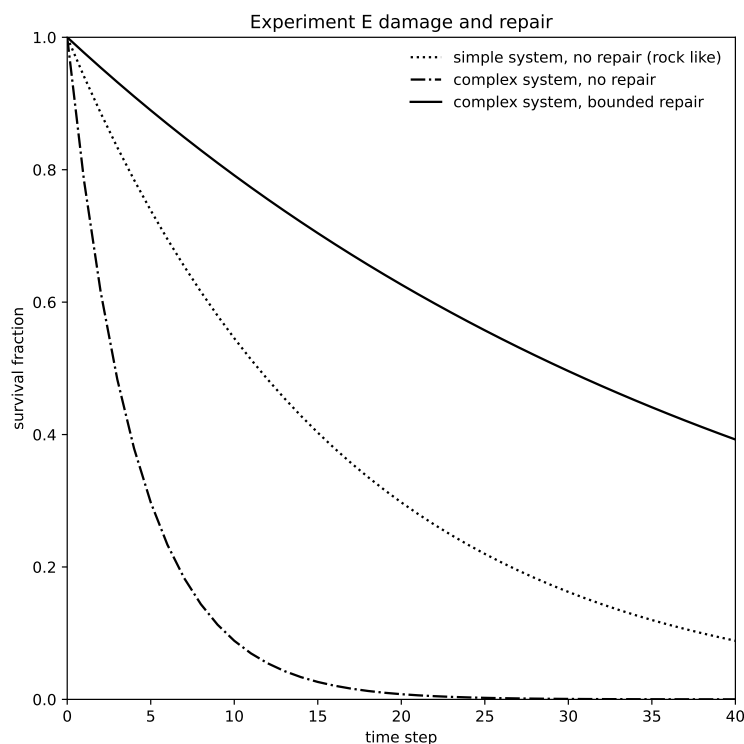


Figure 2. Simplicity avoids damage and complexity requires repair. Survival curves under stochastic damage with bounded repair. The simple system has fewer essential parts so it suffers fewer fatal breaks and needs no repair. The complex system has more to break so it dies quickly without repair. With repair, the complex system becomes viable again. This figure is included as a conceptual contrast, not as a test of Theorem 1.

Interpretation. Simple things persist by not needing repair. Complex things persist by repairing.

Real systems face both novelty and damage. This paper proves a result only about novelty selection. The repair model is included to prevent a category error: higher functional information under novelty does not imply higher complexity, because complexity enters through breakage and repair rather than future compatibility.

Discussion

This paper formalises one version of the law of increasing functional information. The law is not a claim that complexity must increase. It is a claim that selection under novelty biases survivors toward weak constraints that preserve compatibility with many unseen outcomes. This formalises The Cosmic Ought as described in the associated PhD thesis [12].

Within the broader Stack Theory programme, this is the novelty selection counterpart of the stack result that greater utility requires weaker lower layer policies [10,12]. Delegating adaptation down the stack enlarges the higher level vocabulary and permits weaker policies. Here the same logic appears as a law of persistence filtering inside a fixed viable set. The complementary failure mode is overconstraint. When collectives become too tightly constrained, viable policies disappear and parts can splinter off or revert to default behaviour [12]. The present theorem isolates the positive side of that picture: persistence under novelty systematically favours policies that preserve future compatibility.

More broadly, it supports a constraints first view where what persists is shaped by what can be stably realised, not by a monotone complexity gradient [5,6]. This aligns with the characterisation of biological organisation as closure of constraints, in which viable systems are those whose constraints mutually sustain one another [11]. The emphasis on constraints over adaptation echoes the long standing argument that not every persistent feature need be an adaptation shaped by selection [25]. It also aligns with the observation that apparent trends toward complexity can arise from diffusive spread away from a boundary rather than from directional selection [26]. Policies with large buffers are pre-adapted to requirements they were never selected for, formalising a mechanism akin to exaptation [27].

Importantly, simplicity does not explain the novelty result in this framework. Under novelty, the explanatory variable is weakness. Weakness is invariant under re-labellings of the unobserved region, because it is a count of compatible completions [12]. Under damage, simplicity can matter because having fewer essential parts reduces the rate at which the world breaks you. That is a different mechanism, and it is exactly why Experiment E is in the paper. The relationship between robustness and evolvability has been studied extensively in molecular and morphological contexts [28,29]. The present result adds a distinct channel: weakness as robustness to novel selection pressures, where compatible completions play the role that neutral neighbours play in sequence space. In neutral network models, populations drift toward genotypes with more neutral neighbours purely through mutation selection dynamics [30]. Draghi and colleagues show that mutational robustness facilitates evolvability by maintaining access to diverse phenotypes [31]. In this framework, weakness plays an analogous role: it maintains access to diverse future outcomes.

The main empirical message is that in finite worlds where all possibilities can be enumerated, weakness maximisation is a reliable proxy for survival under novelty. This aligns with generalisation optimal learning results for weakness in artificial intelligence [32]. The broader argument that approximate correctness on unseen cases is the right learning target dates to Valiant's foundational framework for probably approximately correct learning [33]. When damage and repair become part of the story, it becomes clear why a single scalar notion of complexity is a trap. Some systems are simple and survive by being hard to break. Other systems are complex and survive by repairing. For example, homeostatic regulation can begin as co-regulation between co-embodied organisms [34], which in the present interpretation would align with the idea of a repair budget capable of supporting complexity. Under novelty, the systems that keep their options open are the ones that tend to survive.

Author Contributions: M.T.B. designed research, performed research, analysed data, and wrote the paper.

Data Availability Statement: Figures 1 and 2 are generated by the Python script `run_all_experiments.py` included with this draft package. The broader code and theory stack is documented in the technical appendices release [35]. A fully reproducible package is included with this manuscript draft.

Conflicts of Interest: The author declares no competing interests.

References

1. Wong, M.L.; Cleland, C.E.; Arend, D.; Bartlett, S.; Cleaves, H.J.; Demarest, H.; Prabhu, A.; Lunine, J.I.; Hazen, R.M. On the roles of function and selection in evolving systems. *Proceedings of the National Academy of Sciences* **2023**, *120*, e2310223120, [<https://www.pnas.org/doi/pdf/10.1073/pnas.2310223120>]. <https://doi.org/10.1073/pnas.2310223120>.
2. Root-Bernstein, M. Evolution is not driven by and toward increasing information and complexity. *Proceedings of the National Academy of Sciences* **2024**, *121*, e2318689121. <https://doi.org/10.1073/pnas.2318689121>.
3. Wong, M.L.; Bartlett, S.; Cleland, C.E.; Demarest, H.; Cleaves, H.J.; Prabhu, A.; Lunine, J.I.; Hazen, R.M. Reply to Root-Bernstein: Increasing complexity allows for the pervasiveness of low-complexity entities and is not anthropocentric. *Proceedings of the National Academy of Sciences* **2024**, *121*, e2406598121. <https://doi.org/10.1073/pnas.2406598121>.
4. Sharma, A.; Czégel, D.; Lachmann, M.; Kempes, C.P.; Walker, S.I.; Cronin, L. Assembly theory explains and quantifies selection and evolution. *Nature* **2023**, *622*, 321–328. <https://doi.org/10.1038/s41586-023-06600-9>.
5. Walker, S.I.; Kim, H.; Davies, P.C.W. The informational architecture of the cell. *Philosophical Transactions of the Royal Society A* **2016**, *374*, 20150057. <https://doi.org/10.1098/rsta.2015.0057>.
6. Solé, R.; Kempes, C.P.; Corominas-Murtra, B.; De Domenico, M.; Kolchinsky, A.; Lachmann, M.; Libby, E.; Saavedra, S.; Smith, E.; Wolpert, D. Fundamental constraints to the logic of living systems. *Interface Focus* **2024**, *14*, 20240010. <https://doi.org/10.1098/rsfs.2024.0010>.
7. Kauffman, S.A. *The Origins of Order: Self-Organization and Selection in Evolution*; Oxford University Press, 1993.
8. McShea, D.W.; Brandon, R.N. *Biology's First Law: The Tendency for Diversity and Complexity to Increase in Evolutionary Systems*; University of Chicago Press, 2010. <https://doi.org/10.7208/chicago/9780226562278.001.0001>.
9. Ashby, W.R. *An Introduction to Cybernetics*, second impression ed.; Chapman & Hall: London, 1957. First published 1956.
10. Bennett, M.T. Are Biological Systems More Intelligent Than Artificial Intelligence?, 2024, [[arXiv:cs.AI/2405.02325](https://arxiv.org/abs/2405.02325)]. In press, 2026, *Philosophical Transactions of the Royal Society B: Biological Sciences*. Special issue on Hybrid agencies: crossing borders between biological and artificial worlds, <https://doi.org/10.48550/arxiv.2405.02325>.
11. Montévil, M.; Mossio, M. Biological organisation as closure of constraints. *Journal of Theoretical Biology* **2015**, *372*, 179–191. <https://doi.org/10.1016/j.jtbi.2015.02.029>.
12. Bennett, M.T. How To Build Conscious Machines. PhD thesis, The Australian National University, 2025. <https://doi.org/10.25911/ta58-p428>.
13. Bekenstein, J.D. Universal upper bound on the entropy-to-energy ratio for bounded systems. *Phys. Rev. D* **1981**, *23*, 287–298. <https://doi.org/10.1103/physrevd.23.287>.
14. Kauffman, S.A. *Investigations*; Oxford University Press: Oxford, 2000.
15. de Finetti, B. Foresight: Its Logical Laws, Its Subjective Sources. In *Breakthroughs in Statistics: Foundations and Basic Theory*; Kotz, S.; Johnson, N.L., Eds.; Springer New York: New York, NY, 1992; pp. 134–174. https://doi.org/10.1007/978-1-4612-0919-5_10.
16. Longo, G.; Montévil, M.; Kauffman, S. No entailing laws, but enablement in the evolution of the biosphere. In *Proceedings of the Proceedings of the 14th Annual Conference Companion on Genetic and Evolutionary Computation*, New York, NY, USA, 2012; GECCO '12, pp. 1379–1392. <https://doi.org/10.1145/2330784.2330946>.
17. Wilke, C.O.; Wang, J.L.; Ofria, C.; Lenski, R.E.; Adami, C. Evolution of digital organisms at high mutation rates leads to survival of the flattest. *Nature* **2001**, *412*, 331–333. <https://doi.org/10.1038/35085569>.
18. Starrfelt, J.; Kokko, H. Bet-hedging—a triple trade-off between means, variances and correlations. *Biological Reviews* **2012**, *87*, 742–755, [<https://onlinelibrary.wiley.com/doi/pdf/10.1111/j.1469-185X.2012.00225.x>]. <https://doi.org/10.1111/j.1469-185X.2012.00225.x>.
19. Hazen, R.M.; Griffin, P.L.; Carothers, J.M.; Szostak, J.W. Functional information and the emergence of biocomplexity. *Proceedings of the National Academy of Sciences* **2007**, *104*, 8574–8581. <https://doi.org/10.1073/pnas.0701744104>.

20. Szostak, J.W. Functional information: Molecular messages. *Nature* **2003**, *423*, 689–689. <https://doi.org/10.1038/423689a>.
21. Price, G.R. Selection and Covariance. *Nature* **1970**, *227*, 520–521. <https://doi.org/10.1038/227520a0>.
22. England, J.L. Statistical physics of self-replication. *The Journal of Chemical Physics* **2013**, *139*, 121923, . <https://doi.org/10.1063/1.4818538>.
23. Solé, R.; Seoane, L.F.; Pla-Mauri, J.; Bennett, M.T.; Hochberg, M.E.; Levin, M. Cognition spaces: natural, artificial, and hybrid, 2026, [2601.12837]. <https://doi.org/10.48550/arxiv.2601.12837>.
24. Bennett, M.T. Is Complexity an Illusion? In Proceedings of the 17th International Conference on Artificial General Intelligence. Springer, 2024, Lecture Notes in Computer Science. https://doi.org/10.1007/978-3-031-65572-2_2.
25. Gould, S.J.; Lewontin, R.C. The spandrels of San Marco and the Panglossian paradigm: a critique of the adaptationist programme. *Proceedings of the Royal Society of London. B. Biological Sciences* **1979**, *205*, 581–598, [<https://royalsocietypublishing.org/rspb/article-pdf/205/1161/581/171995/rspb.1979.0086.pdf>]. <https://doi.org/10.1098/rspb.1979.0086>.
26. Gould, S.J. *Full House: The Spread of Excellence from Plato to Darwin*; Harmony Books: New York, 1996; p. 244.
27. Gould, S.J.; Vrba, E.S. Exaptation—a Missing Term in the Science of Form. *Paleobiology* **1982**, *8*, 4–15. <https://doi.org/10.1017/S0094837300004310>.
28. Wagner, A. Robustness and evolvability: a paradox resolved. *Proceedings of the Royal Society B: Biological Sciences* **2008**, *275*, 91–100, [<https://royalsocietypublishing.org/rspb/article-pdf/275/1630/91/593598/rspb.2007.1137.pdf>]. <https://doi.org/10.1098/rspb.2007.1137>.
29. Kirschner, M.; Gerhart, J. Evolvability. *Proceedings of the National Academy of Sciences* **1998**, *95*, 8420–8427, [<https://www.pnas.org/doi/pdf/10.1073/pnas.95.15.8420>]. <https://doi.org/10.1073/pnas.95.15.8420>.
30. van Nimwegen, E.; Crutchfield, J.P.; Huynen, M. Neutral evolution of mutational robustness. *Proceedings of the National Academy of Sciences* **1999**, *96*, 9716–9720, [<https://www.pnas.org/doi/pdf/10.1073/pnas.96.17.9716>]. <https://doi.org/10.1073/pnas.96.17.9716>.
31. Draghi, J.A.; Parsons, T.L.; Wagner, G.P.; Plotkin, J.B. Mutational robustness can facilitate adaptation. *Nature* **2010**, *463*, 353–355. <https://doi.org/10.1038/nature08694>.
32. Bennett, M.T. Computational Dualism and Objective Superintelligence. In Proceedings of the 17th International Conference on Artificial General Intelligence. Springer, 2024, Lecture Notes in Computer Science. https://doi.org/10.1007/978-3-031-65572-2_3.
33. Valiant, L.G. A theory of the learnable. *Commun. ACM* **1984**, *27*, 1134–1142. <https://doi.org/10.1145/1968.1972>.
34. Ciaunica, A.; Constant, A.; Preissl, H.; Fotopoulou, K. The first prior: from co-embodiment to co-homeostasis in early life. *Consciousness and cognition* **2021**, *91*, 103117. <https://doi.org/10.1016/j.concog.2021.103117>.
35. Bennett, M.T. Technical Appendices, 2025. Archived release on Zenodo. Source repository: <https://github.com/ViscousLemming/Technical-Appendices>, <https://doi.org/10.5281/zenodo.17164785>.

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.