Admissibility Boundaries and Edge-Selected Structure

Keith Taylor
VERSF Theoretical Physics Program

Table of Contents

General Reader Summary

Why do so many things in physics seem to "just happen" to match up? The Sun and Moon appear almost exactly the same size in our sky. Certain physical constants seem suspiciously well-balanced for stars and atoms to exist. Scientists often invoke either extraordinary coincidence or the idea that observers like us could only exist in universes where things line up this way.

This paper proposes a simpler explanation: **things that last tend to exist right at the edge of what's physically possible.**

Here's the intuition. Imagine you're a surfer. You want to be as far from the dangerous rocks as possible, but you also need to stay where the waves are. There's a sweet spot—the farthest point from danger where surfing is still possible. You don't choose this spot because it's special; you end up there because everywhere else either crashes you into rocks or leaves you without waves. The edge is not chosen because it is special; it is selected because everything else is either unstable or unattainable under the constraints.

Physical systems work similarly. A moon orbits as far from its planet as it can while still being gravitationally bound. A star burns at the edge between gravitational collapse and explosive dispersal. These systems don't "choose" their configurations—they evolve toward boundaries and stay there because that's where persistence is possible.

The key insight is this: **when multiple boundaries intersect, the numbers describing different aspects of the system end up being similar to each other.** Not because of coincidence, not because of observers, but because the geometry of "what's allowed" forces them together.

This paper works through the math, tests the idea on real astrophysical systems, and makes predictions that future observations could prove wrong.

Technical Abstract

Long-lived structures in nature do not occupy arbitrary regions of parameter space but instead settle at the maximal distance from destabilizing interfaces consistent with persistence. When multiple such admissibility boundaries intersect, near-equalities between otherwise unrelated scales arise generically, without fine-tuning or anthropic assumptions. This note formalizes the principle, demonstrates it through classical and astrophysical examples, and grounds it in a time-free formulation of bit conservation capacity.

---

1. Persistence as an Inequality Constraint

General reader version: For something to last, the forces holding it together must be stronger than the forces trying to tear it apart. This section puts that intuition into mathematical form.

Consider a system characterized by a control parameter x, interpreted as distance from a generative or destabilizing interface. Here x is a generalized separation coordinate (spatial distance, energetic decoupling, or causal depth) along which stabilizing coupling and disruptive influence trade off monotonically. Let g(x) denote a stabilizing or coherence-supplying mechanism required for persistence, and Φ(x) denote disruptive influences. Persistence requires the inequality:

$$g(x) / \Phi(x) \geq C_{min}$$

where C_{min} is a minimum coherence threshold. This inequality defines an *admissible region* $\mathcal{A}$ of parameter space, with *admissibility boundary*:

$$\partial\mathcal{A} = \{ x : g(x)/\Phi(x) = C_{min} \}$$

In plain terms: g(x) is whatever holds the system together (self-gravity, chemical bonds, phase coherence), Φ(x) is whatever threatens to destroy it (tidal forces, thermal fluctuations, external shear), and C_{min} is the minimum safety margin needed. The system can only exist where the ratio exceeds this threshold.

---

2. Edge Selection

General reader version: Systems don't just exist anywhere in the "allowed zone"—they migrate to its outer edge. This section explains why.

If g(x) decreases monotonically with x (loss of coupling) and Φ(x) also decreases with x (reduced disturbance), but Φ(x) approaches a nonzero background floor, then there exists a finite maximal x* such that:

$$g(x^*) / \Phi(x^*) = C_{min}$$

Systems evolve toward increasing x to minimize disturbance, but cannot exceed x* without losing persistence. The equilibrium configuration therefore lies at the *admissibility boundary*, not in the interior.

Why the edge? Moving away from a disruptive source (like a planet's tidal forces) is generally beneficial—less stress, less energy expenditure, calmer conditions. So systems migrate outward. But they can't go infinitely far, because at some point the stabilizing influence (like gravitational binding) becomes too weak. The edge is where these two effects balance: as far as possible from disruption while still being held together.

This is not fine-tuning. It's the natural outcome of any system that (a) benefits from distance from disruption and (b) requires some minimum connection to persist.

[Figure 1 should appear here] *Schematic of edge selection: x-axis shows the control parameter x (e.g., orbital distance). The shaded region marks the admissible zone where g/Φ ≥ C_min. The boundary x* is marked. A second constraint (dashed) intersects, narrowing the admissible region. A slow dynamical trajectory (arrow) shows a system migrating toward and through the boundary intersection—illustrating how "windows" of near-equality arise during passage.*

2.1 Edge Consistency Principle: Facts Require Saturation

General reader version: Why do systems end up at the edge rather than somewhere in the middle of the allowed zone? This subsection argues it's not just about optimization—it's about whether the system can maintain a consistent identity at all.

In a time-free framework where temporal ordering is defined by irreversible commitments, a "persistent structure" is precisely a system for which the set of required commitments closes consistently from one stable cycle to the next. If a system lies deep inside an admissible region, there is generically slack: multiple distinct micro-updates can satisfy the inequalities while leading to macroscopically different outcomes, so the mapping from commitments to state is not uniquely well-defined. At an admissibility boundary—where one or more constraints are saturated—the slack is removed: the commitments required to maintain identity are pinned to the maximal capacity available, yielding a self-consistent, reproducible bookkeeping of facts.

In this sense, edge selection is not merely an optimization tendency; it is the condition under which fact formation becomes consistent and repeatable.

Lemma (Slack implies non-uniqueness). Let $\mathcal{A} = \{x : g(x)/\Phi(x) \geq C\_min\}$ be an admissible region and suppose x lies in its interior. Then there exists $\varepsilon > 0$ such that $x \pm \varepsilon \in \mathcal{A}$. If observables depend nontrivially on x, the interior permits multiple distinct realizations compatible with persistence, hence fact-encoding is underdetermined by the constraints. Saturation at $\partial\mathcal{A}$ removes this underdetermination.

In particular, if the macroscopic identity map $x \mapsto \mathcal{S}(x)$ is non-constant over $\mathcal{A}$, then interior slack implies multiple admissible macrostates compatible with persistence.

In plain terms: Inside the allowed zone, there's wiggle room—multiple configurations could work, so the system's "identity" isn't pinned down. At the boundary, the wiggle room vanishes. The system must commit to exactly one configuration to survive. That's what makes facts definite rather than fuzzy.

2.2 Why Admissibility Boundaries Generically Yield Logarithmic Dependence

General reader version: The mathematical assumptions in later sections aren't arbitrary—they follow from the physics of how stabilizing forces weaken with distance.

The logarithmic dependence of the boundary location x* on physical parameters is not an assumption but a common structural outcome when a decaying coupling is balanced against a non-vanishing disturbance floor.

Exponential coupling vs constant floor. Suppose the stabilizing mechanism decays as:

$$g(x; p) = g_0(p)\, e^{-x/\ell},\quad \Phi(x; p) \to \Phi_\infty(p) > 0$$

Saturation $g/\Phi = C_{min}$ yields:

$$x^* = \ell \ln[\, g_0(p) / (C_{min}\, \Phi_\infty(p)) \,]$$

i.e., x* is logarithmic in the underlying physical parameters entering g_0 and Φ_∞.

Power-law coupling vs constant floor. If $g(x) = g_0(p)\, x^{-k}$ and $\Phi \to \Phi_\infty(p)$, then:

$$x^* = [\, g_0(p) / (C_{min}\, \Phi_\infty(p)) \,]^{1/k}$$

which is sub-linear in parameter variations and likewise suppresses extreme sensitivity.

In both cases, boundary selection converts multiplicative parameter dependence into additive dependence in x*, explaining why ratios evaluated at x* are frequently compressed rather than hierarchical. This provides the physical basis for the assumptions in Proposition 1 (Section 5).

---

3. Example: Orbital Mechanics (Roche Limit)

General reader version: This section applies the abstract principle to something concrete—why moons and rings exist where they do around planets.

In classical orbital mechanics, tidal shear provides the disruptive term:

$$\Phi(x) \sim G M_p r / x^3$$

while self-gravity provides cohesion:

$$g \sim G \rho_m r$$

Persistence requires $g \geq \Phi(x)$. This yields the Roche scaling [1, 2]:

$$x \propto R_p (\rho_p / \rho_m)^{1/3}$$

The full Roche limit includes an order-one coefficient (≈ 2.44 for fluid bodies). Material migrates inward until it reaches this boundary, inside which persistence fails. Moons and rings therefore occupy the edge of admissibility.

In plain English: Get too close to a planet, and its gravity will pull harder on your near side than your far side—this differential pull (tidal force) will rip you apart. The Roche limit is the closest distance at which an object's own gravity can hold it together against this tidal stretching.

Saturn's rings exist inside Saturn's Roche limit—that's why they're rings (rubble that can't coalesce) rather than moons. Saturn's moons exist outside the Roche limit—where self-gravity wins and objects can hold together.

This is edge selection in action: material settles at the boundary between "can exist" and "gets destroyed."

3.1 BCB Derivation of the Roche Boundary

General reader version: The standard derivation compares forces. Here we show the same result emerges from comparing information capacities—how much "distinctness" a system needs to maintain versus how much it can support. This demonstrates that BCB isn't just relabeling physics; it's a different route to the same destination.

The Roche limit is usually derived by equating tidal differential acceleration across a satellite to its self-gravity. Here we show the same boundary arises from a time-free bit-capacity criterion.

(i) Time-free BCB form

Let $\mathcal{P}$ be the process "maintain a bound aggregate of radius r as a coherent object under an external tidal field." In a time-free formulation, persistence requires:

$$\mathscr{B}(\mathcal{P}; x) \le \mathcal{K}(E, R; x)$$

where $\mathscr{B}$ is the irreversible commitment complexity required to preserve the object's identity through one stable structural cycle, and $\mathcal{K}$ is the maximum admissible commitment capacity supportable by the physical substrate. To avoid introducing a clock, we compare commitment counts per structural cycle, not rates.

(ii) Capacity scaling from known physics (Bekenstein-type)

A universal, time-independent capacity proxy is the Bekenstein-style information bound [3, 4], which scales as:

$$I \lesssim 2\pi ER / (\hbar c \ln 2)$$

For any structural maintenance task whose "cost" is set by an energy budget E over a characteristic scale R, a natural time-free capacity functional is:

$$\mathcal{K}(E, R) \propto E \cdot R \text{ (universal prefactor cancels in ratios)}$$

We use the Bekenstein-style scaling $\mathcal{K} \propto ER$ as a universal capacity proxy. For consistent comparison, we express both the disruptive requirement and intrinsic capacity in the same ER units, taking $R \sim r$ for the satellite's coherence scale. This choice affects only order-one prefactors, not the density scaling or the existence of the boundary.

(iii) Apply to a satellite in a tidal field

Consider a satellite (moon) of radius r and density ρ_m orbiting a planet of mass M_p at distance x.

Disruptive "commitment requirement." The external tidal field produces a differential acceleration across the satellite of order:

$$\Delta a_{tidal} \sim 2GM_p r / x^3$$

To preserve the satellite as a coherent object, internal degrees of freedom must continually "correct" (resist) the induced shear. A natural time-free proxy for the required commitment complexity is the tidal work scale across the object:

$$E_{tidal} \sim m \cdot \Delta a_{tidal} \cdot r \sim (8\pi/3)\, GM_p \rho_m r^5 / x^3$$

Here E_{tidal} is used as a proxy for the minimum irreversible "repair" work required to prevent tidal shear from destroying the object's identity over one structural cycle.

Available "commitment capacity." The satellite's ability to maintain distinctions is bounded by its internal binding resource, naturally proxied by its self-gravitational binding energy:

$$E_{bind} \sim \alpha\, Gm^2 / r = \alpha\, (16\pi^2/9)\, G \rho_m^2 r^5$$

where α is an order-one structural factor (geometry, rigidity versus fluid response).

Because the capacity proxy is of Bekenstein form ($\propto ER$), we compare disruptive and binding resources in the same ER units; this changes only order-one coefficients. Since $\mathscr{B} \propto E_{tidal} \cdot r$ and $\mathcal{K} \propto E_{bind} \cdot r$, the universal prefactor $(2\pi/\hbar c \ln 2)$ cancels and the BCB inequality reduces to the purely mechanical condition:

$$E_{tidal} \lesssim E_{bind}$$

Substituting the expressions above:

$$(8\pi/3)\, GM_p \rho_m r^5 / x^3 \lesssim \alpha\, (16\pi^2/9)\, G \rho_m^2 r^5$$

$$\Longrightarrow x^3 \gtrsim (3/2\pi\alpha)\, M_p / \rho_m$$

Writing $M_p = (4\pi/3)\, \rho_p R_p^3$:

$x \gtrsim \kappa R_p (\rho_p / \rho_m)^{1/3}$, $\kappa \equiv (2/\alpha)^{1/3} \times$ (geometry factors)

For fluid bodies the exact Roche coefficient is $\kappa \simeq 2.44$ [1]. In this BCB derivation, κ collects the order-one constants arising from the precise tidal potential, shape response, and binding profile; the density scaling and the existence of a sharp boundary are the robust outputs.

(iv) Why this is not a relabeling

The point is not that "tidal stress > binding stress," which is standard. The point is that the same boundary is obtainable by comparing universal information-capacity functionals:

- a disruption-driven commitment requirement $\mathscr{B}$ controlled by the external tidal environment, and
- an intrinsic capacity $\mathcal{K}$ controlled by the satellite's binding resource.

This is the same BCB logic later instantiated exactly by black-hole saturation of the Bekenstein bound, but here it reproduces a classical orbital admissibility boundary.

BCB therefore reproduces a standard "stability" boundary from a time-free capacity principle, showing the informational formulation is compatible with and reproduces known admissibility boundaries, while enabling generalizations beyond force-based stability analysis.

---

4. Intersection of Multiple Boundaries

General reader version: Real systems face multiple constraints simultaneously. Where these constraints overlap, interesting things happen.

In many systems, multiple independent persistence constraints apply simultaneously. Each defines an admissible region. Their intersection is typically narrow, forcing the system to occupy a constrained edge where several inequalities saturate at once. Observable quantities that depend differently on x then acquire near-equal magnitudes.

Analogy: Imagine you need to find an apartment that's (1) affordable, (2) close to work, and (3) in a safe neighborhood. Each requirement rules out part of the city. The overlap of all three constraints might be just a few blocks. If you measure different things about apartments in that zone—say, commute time and monthly cost in hundreds of dollars—they might end up being similar numbers, not because rent and commuting are related, but because the tight constraint zone forces everything into a narrow range.

Physical systems work the same way. When multiple survival constraints intersect, the different quantities describing the system get squeezed toward similar values.

---

5. Emergence of Near-Equal Ratios

General reader version: This section explains mathematically why "coincidences" pop out when systems live at constraint boundaries.

Proposition 1 (Ratio suppression at admissibility boundaries)

Assume the admissibility boundary x^* is defined implicitly by saturation of an inequality:

$$g(x; p) / \Phi(x; p) = C_{min}$$

where p denotes underlying physical parameters. Suppose x^* depends at most logarithmically on p:

$$x^* = x_0 + \Sigma_i \alpha_i \ln p_i$$

with α_i of order ℓ (a characteristic scale in x-space). As shown in Section 2.2, this logarithmic dependence arises generically when decaying couplings are balanced against non-vanishing disturbance floors. Let observables A(x), B(x) be smooth and satisfy:

$$\ln A(x) = a_0 + x/\ell_A + o(1), \ln B(x) = b_0 + x/\ell_B + o(1)$$

over the admissible window. (Here o(1) denotes subleading variation over the admissible interval containing x^*.)

Then the boundary-evaluated ratio obeys:

$$\ln[A(x^*)/B(x^*)] = (a_0 - b_0) + x^*(1/\ell_A - 1/\ell_B) + o(1)$$

and therefore:

$$A(x^*)/B(x^*) = e^{(a_0-b_0)} \prod_i p_i^{[\alpha_i(1/\ell_A - 1/\ell_B)]} \times (1 + o(1))$$

Bounded Consequence (the real "order-one" claim)

If parameters satisfy $p_i \in [10^{-m}, 10^{m}]$ for some modest m (i.e., no inserted hierarchy), and $|\alpha_i(1/\ell_A - 1/\ell_B)| \lesssim 1$, then:

$$10^{-mn} \lesssim A(x^*)/B(x^*) \lesssim 10^{mn}$$

so the ratio is bounded away from parametrically extreme values. In typical applications, the effective number of controlling parameters n_{eff} is small (order unity), so this bound represents "same-ballpark" rather than extreme separation.

Conversely: If any p_i is hierarchical (e.g., $p_i \sim 10^6$ or larger), then parametrically large or small ratios are expected and are not exceptions to the framework.

Proof sketch. Substitute the assumed forms for x*, ln A, and ln B, then exponentiate. The boundedness follows from bounding each p_i and noting the exponents are order-one under the stated conditions. ■

What This Means

In plain terms: If the underlying physics doesn't already contain huge disparities (like a factor of a million between two quantities), then quantities measured at a constraint boundary will typically be within a factor of 10 or 100 of each other—"same ballpark" in physics terms. The proposition shows this isn't coincidence; it's a geometric consequence of where persistent systems can exist.

The framework's power: It explains both when near-equalities *should* occur (moderate parameter ranges, logarithmic boundary dependence) and when they should *not* (hierarchical parameters, extreme scale separations).

Assumptions for the order-one claim:

1. The admissibility boundary x* depends logarithmically or sub-linearly on underlying physical parameters.
2. Observables A(x) and B(x) vary smoothly (exponentially or polynomially) with x.
3. The control parameters span no more than a few orders of magnitude—i.e., no pre-existing hierarchy of 10^6 or greater is inserted a priori.

---

6. Implications

This framework explains why near-equalities recur across physics—orbital systems, phase transitions, replication thresholds, and cosmology—without invoking fine-tuning or observers. They arise because persistent structure is only possible near constraint boundaries, and multiple such boundaries often coincide.

The shift in perspective: Instead of asking "why do these numbers happen to match?" we ask "what constraints determine where this system can exist?" Often, the near-equality falls out automatically.

6.1 Fine-Tuning Reinterpretation

Many apparent near-equalities traditionally cited as evidence of fine-tuning can be reinterpreted as signatures that persistent structure resides on narrow intersections of admissible regions, rather than evidence of parameter tuning. The numbers match not because they were set to match, but because matching is where persistence is possible. This perspective complements discussions of naturalness and fine-tuning in fundamental physics [17, 18].

6.2 Limits of Anthropic Reasoning

The framework does not deny observational selection effects [19] but reduces their explanatory necessity. Near-equalities may arise from constraint geometry even in the absence of observers. Where the admissibility-boundary explanation suffices, anthropic reasoning becomes unnecessary—though both explanations may apply in some cases.

6.3 Predicting Absence as Well as Presence

In systems with strong pre-existing hierarchies (control parameters spanning more than ~6 orders of magnitude), the framework predicts that near-equalities should *not* generically appear at boundary intersections. This provides a way to audit where "coincidence explanations" are appropriate and where they are not. A systematic survey of physical near-equalities, sorted by the hierarchy of their underlying parameters, would test this prediction.

Scope

This note is a framework statement: it explains why persistent structures concentrate at admissibility boundaries and why intersections generically yield near-equalities. It does not claim that all near-equalities must occur, nor that every boundary intersection produces matching. When strong hierarchies are present in the control parameters, the framework predicts that near-equality is generally *not* expected.

The BCB formulation is currently at the stage of (i) proof of principle via extremal systems (black holes), (ii) compatibility demonstration with classical boundaries (Roche limit), and (iii) conceptual distinction from stability analysis. Demonstrating quantitative BCB constraints in ordinary natural systems—where bit conservation is the dominant rather than merely compatible explanation—remains an open research direction.

---

7. Summary Principle

Stable structure maximizes distance from destabilizing interfaces subject to persistence constraints; when multiple constraints intersect, near-equalities between physical scales are generic outcomes.

Appendix A: Solar Eclipse Angular Size Near-Equality

General reader version: The Sun and Moon appear almost exactly the same size in Earth's sky. This appendix applies the framework to explain why this might not be a coincidence—and why it won't last forever.

A1. Angular size and the near-equality

For small angles, the angular diameter θ of an object is approximately:

$\theta \approx 2R / d$

where R is physical radius and d is distance to the observer. A total solar eclipse is possible when the Moon's angular diameter is comparable to (and sometimes slightly larger than) the Sun's.

In everyday terms: How big something looks depends on both its actual size and how far away it is. A basketball held at arm's length might look the same size as the Moon. The Sun is about 400 times larger than the Moon, but it's also about 400 times farther away—so they look almost identical in size.

A2. Observed solar-system values

Parameter	Value
Moon radius	1,737 km
Earth–Moon distance (mean)	384,400 km
Sun radius	696,340 km
Earth–Sun distance (1 AU)	149,600,000 km

Computing the small-angle angular diameters (in radians):

$\theta\_Moon \approx 2R\_Moon / d\_EM \approx 0.009037$ $\theta\_Sun \approx 2R\_Sun / d\_ES \approx 0.009309$ $\theta\_Moon / \theta\_Sun \approx 0.971$

They match to within 3%.

A3. The "~400×" coincidence

The same result can be expressed by comparing the ratios of distances and radii:

$d\_ES / d\_EM \approx 389$, $R\_Sun / R\_Moon \approx 401$

Because these two ratios are close, the apparent angular sizes are close.

A4. Temporal qualification

General reader version: This "coincidence" is temporary. The Moon is slowly drifting away from Earth, and in the distant future, total solar eclipses will no longer be possible.

The eclipse near-equality is not a static equilibrium configuration. Tidal dissipation causes the Moon's orbit to expand at approximately 3.8 cm/year [7], so the system passes through a finite temporal window during which the Sun's and Moon's apparent angular sizes are comparable.

In the admissibility-boundary framework, this is not a flaw but an expected outcome: near-equalities arise generically when a slowly evolving system crosses the intersection of multiple constraints. The framework therefore predicts *windows* of near-equality rather than permanent attractors.

For the Earth–Moon system, the ±10% angular-match window spans on the order of 10^8 years—a significant fraction of the system's multi-billion-year evolution, finite but not implausibly narrow [8]. The fact that complex observers exist during such a window is a separate biological question and is not required for the physical argument.

Key point: The framework doesn't claim the Moon is "stuck" at this special distance. It claims that when a slowly-changing system passes through a constraint intersection, near-equalities naturally occur during that passage. We happen to exist during such a window—but the window itself is a significant fraction of the system's lifetime, not a miraculous instant.

A5. How this supports the boundary-intersection claim

This numerical example demonstrates the *observable signature* of the framework: near-equality in an apparent ratio (angular size) arising from the interaction of multiple constrained scales (orbital distance constraints for moons and stellar/planetary distance constraints for habitable planets). The framework predicts that when each of those scales is pinned by admissibility boundaries, their intersection naturally produces order-one apparent ratios rather than extreme separations.

Appendix B: Multi-Boundary Test in M-Dwarf Habitable Zones

General reader version: If the framework is correct, it should make predictions about other star systems. This appendix asks: around what types of stars could moons produce Sun–Moon-like *eclipse near-equality? The answer is surprising—for the smallest stars, it's geometrically impossible.*

B1. Two boundaries for a moon orbit

A moon's orbit is constrained from both sides:

(i) Inner bound (tidal disruption / Roche-type): Too close and the planet tears the moon apart.

$a \geq a_{min} \approx 2.44 R_p (\rho_p / \rho_m)^{(1/3)}$

(ii) Outer bound (stellar unbinding / Hill stability): Too far and the star's gravity steals the moon away [9, 10].

$a \leq a_{max} \approx f R_H$, $f \approx 0.49$ for long-term prograde stability

$R_H \approx a_p (M_p / 3M_*)^{(1/3)}$

where a_p is the planet–star distance, M_p is planet mass, and M_* is stellar mass.

In plain terms: There's a "Goldilocks zone" for moon orbits—not too close (or tidal forces destroy the moon), not too far (or the star's gravity captures it). The size of this zone depends on the star's mass and the planet's distance from it.

B2. Eclipse near-equality distance

Using the small-angle approximation, the angular diameter ratio is:

$\theta_{Moon} / \theta_* = (R_{Moon} / a) / (R_* / a_p) = (R_{Moon} \cdot a_p) / (R_* \cdot a)$

Define a_{eq} as the orbital radius at which $\theta_{Moon} = \theta_*$:

$a_{eq} = (R_{Moon} \cdot a_p) / R_*$

A near-equality band (e.g., ±10%) corresponds to a within $[a_{eq}/1.1, a_{eq}/0.9]$.

Translation: For any given star and moon size, there's a specific orbital distance where the moon and star would appear the same size. The question is: does this distance fall within the range where a moon can stably exist?

B3. A key simplification

Requiring $a_{eq} \leq a_{max}$ gives:

$R_{Moon} / R_* \leq f (M_p / 3M_*)^{(1/3)}$

The planet-star distance cancels out. Whether eclipse near-equality is possible depends primarily on stellar mass and radius, not on the exact habitable-zone distance.

B3.1 Sensitivity Analysis

The feasibility condition scales as $R_m \lesssim R_* f (M_p / 3M_*)^{(1/3)}$. Thus increasing planet mass widens the feasible region only weakly ($\propto M_p^{(1/3)}$), while increasing moon size tightens it linearly.

For M dwarfs with $R_* \sim M_*^{\beta}$ *(with $\beta \approx 1$ over much of the low-mass main sequence) [11], the stellar threshold scales approximately as* $M_*^{(\beta-1/3)}$, indicating that the ~0.3 $M_\odot$ cutoff is moderately robust to order-unity changes in M_p and R_m, but should shift predictably for super-Earth planets or sub-lunar satellites.

Scaling summary:

- Doubling planet mass shifts the threshold by only $2^{(1/3)} \approx 1.26$
- Halving moon radius doubles the feasible stellar mass range
- The prediction is robust for Earth-like systems but should be recalculated for significantly different planet/moon configurations

Corollary (Near-equality feasibility). Eclipse near-equality is feasible if and only if:

$R_{Moon} / R_* \leq f (M_p / 3M_*)^{(1/3)}$

B4. Worked examples (Earth-mass planet, Moon-sized moon)

Key finding: For $M_* < 0.3\ M_\odot$, the admissible orbital window does not intersect the ±10% eclipse near-equality band.

Star type	M_*	HZ distance	a_max	a_eq	Reachable?
Deep M-dwarf	0.2 $M_\odot$	0.060 AU	75,000 km	89,300 km	**No**
Late M / early K	0.4 $M_\odot$	0.201 AU	200,200 km	187,700 km	Yes
K-dwarf	0.5 $M_\odot$	0.297 AU	274,700 km	221,900 km	Yes

What this means: For $M_* < 0.3\ M_\odot$, the system cannot reach the ±10% near-equality band. Since $\theta_{Moon}/\theta_* \propto 1/a$, and a cannot exceed a_{max}, the moon remains *larger* in angular size than the star across the entire stable window. Total eclipses could still occur (the moon would more than cover the star), but Sun–Moon-like near-equality cannot. The claim here concerns near-equality ($0.9 \leq \theta_{Moon}/\theta_* \leq 1.1$), not eclipse occurrence per se.

By contrast, around Sun-like and larger stars, the allowed zone is big enough to include the near-equality distance.

B5. What this demonstrates

This is a direct test of the multi-boundary framework:

1. Physical admissibility (Roche + Hill) produces a finite allowed window for moon orbits.
2. The eclipse near-equality condition defines a target band in that same coordinate.
3. Whether coincidence-like near-equality is possible is determined by whether the target band intersects the admissible window.

The framework makes a real prediction: Sun–Moon-like angular near-equality should be impossible around the smallest stars, because the Hill-stability outer boundary compresses the orbital window below the near-equality distance. Future exomoon surveys [12, 13] can test this prediction.

Appendix C: BCB Recasting — Admissibility as Bit Conservation Capacity

General reader version: This section connects the admissibility idea to a deeper principle from information theory. The claim is that "what can persist" is ultimately determined by how much information a system can maintain against disruption.

C1. Bit commitment as a physical requirement

Any long-lived physical structure must continuously maintain a set of irreversible distinctions—facts—that define its identity through time. These distinctions correspond to committed bits: once established, they cannot be undone without physical cost. Persistence therefore requires ongoing bit commitment against environmental disturbance.

Analogy: Think of a sandcastle. To persist, it must maintain distinctions: "this grain is here, not there." Waves constantly try to erase these distinctions. The castle persists only if it can maintain its defining structure faster than the waves destroy it. The "bits" are the facts about the castle's shape; "commitment" is the physical process that maintains them.

C2. Time-free formulation of the BCB

In a framework where time emerges from irreversible commitments, expressions involving bit *rates* are circular. The BCB must be expressed without reference to time.

Let:

- $\mathscr{B}(\mathcal{P})$ = the number of irreversible bit commitments required for a process $\mathcal{P}$ to complete one stable cycle
- $\mathcal{K}(E, R)$ = the maximum admissible number of irreversible commitments supportable within a region of total energy E and characteristic size R

Persistence requires the inequality:

$\mathscr{B}(\mathcal{P}) \leq \mathcal{K}(E, R)$

This formulation is time-free and compares required commitment complexity to admissible commitment capacity.

In simpler terms: Every stable process needs to "write" a certain amount of information to complete one cycle. The universe imposes a limit on how much information can be written in a given region with a given amount of energy. A process can only persist if its information requirement fits within this budget.

C2.1 Structural Cycles as Closed Macroscopic Updates

A "structural cycle" is defined operationally as one closed update of the system's coarse-grained state under its internal dynamics: a mapping $\mathcal{S} \mapsto \mathcal{S}'$ that returns the system to the same macroscopic equivalence class $[\mathcal{S}]$ (up to tolerated fluctuations).

Examples include one orbital period for a moon (closure of orbital phase), one oscillation cycle for a mode-locked system, one reaction loop for a catalytic cycle, or one relaxation ringdown mode for a compact object. Importantly, this definition is time-free: it counts closure events of macroscopic state, not seconds.

For a tidally stressed satellite, the structural cycle is completion of one orbit—the period over which tidal deformation completes one full cycle and the moon returns to the same orbital phase. The commitment count $\mathscr{B}$ measures how many irreversible distinctions must be maintained to preserve the satellite's identity through this closure.

C3. Edge selection in BCB language

Moving away from a generative or destabilizing interface typically reduces environmental disturbance and therefore lowers the required commitment count $\mathscr{B}$. However, retreat also reduces coupling to the interface, decreasing the available capacity $\mathcal{K}$. Systems evolve toward the maximal distance x* at which the BCB inequality is just satisfied:

$x^* = \max\{x : \mathscr{B}(\mathcal{P}; x) \leq \mathcal{K}(E, R; x)\}$

C4. Why BCB is not a relabeling of stability

General reader version: "Stable" usually means "doesn't fall apart when poked." BCB is different—it asks whether a system can maintain its own identity at all, regardless of external pokes.

Traditional stability criteria describe whether perturbations grow or decay under specific dynamics. The Bit Conservation Boundary instead limits whether a system can maintain the distinctions required to define its state at all.

Stability analysis asks whether perturbations grow; BCB asks whether the system can sustain the irreversible bookkeeping required to remain the same system over the relevant cycle.

Two configurations may both be dynamically stable, yet only one lies below the irreversible information capacity required for persistence. A satellite orbit may be linearly stable against perturbations yet still exceed the commitment capacity required to maintain phase coherence over secular timescales; BCB excludes it where stability analysis does not.

Example: A spinning top can be stable against small pushes (dynamical stability), but if there's too much friction, it eventually can't maintain "which way is up"—it loses the information that defines its spinning state. BCB captures this second kind of limit; classical stability analysis often doesn't.

Unifying principle. Roche limits, Hill stability, and black-hole entropy bounds share no common force-based description, yet all impose hard ceilings on sustainable information structure. BCB unifies them at the level of admissibility rather than dynamics.

C4.1 A Configuration That Is Dynamically Stable but BCB-Excluded

General reader version: Here's where BCB does something standard stability analysis cannot: it can rule out configurations that are "stable" in the usual sense but cannot persist as repeatable facts.

Dynamical stability is a statement about the local growth of perturbations under idealized evolution. Persistence as a "fact" is stronger: it requires that the system's macroscopic state remain distinguishable and repeatable across many structural cycles under unavoidable dissipation channels.

Consider a bound orbit that is linearly stable under conservative dynamics. If the orbit radiates energy (e.g., via gravitational waves in strong-field regimes [14] or electromagnetic radiation for accelerated charges), then over successive cycles the orbit's parameters drift. The configuration may remain locally stable at every instant, yet cease to define a persistent macroscopic fact unless the system can maintain sufficient irreversible bookkeeping resolution to re-identify "the same orbit" across cycles.

In BCB terms, persistence requires that the commitment complexity needed to maintain orbital identity across N cycles, $\mathscr{B}\_N(\mathcal{P})$, remain below the admissible capacity $\mathcal{K}(E, R)$ of the orbital region:

$$\mathscr{B}\_N(\mathcal{P}) \leq \mathcal{K}(E, R)$$

This yields an exclusion criterion on long-lived orbits that is not captured by linear stability: there exist parameter ranges in which an orbit is dynamically stable but cannot persist as a repeatable, distinguishable structure over many cycles because dissipation-driven drift overwhelms the available distinguishability capacity.

The logical wedge: Stability ≠ persistent fact. A metastable state in statistical mechanics provides another example: a configuration can be dynamically stable against small perturbations but cannot persist because thermal fluctuations erase memory faster than the system can maintain distinguishable macrostates. Classical stability doesn't bound the maximum number of distinguishable macrostates; BCB does.

Important qualification: In purely classical gravitational systems, dissipation timescales already track orbital drift with effectively infinite resolution, so BCB-based persistence bounds are not expected to be parametrically tighter than standard inspiral calculations. The conceptual distinction becomes quantitative only when distinguishability is finite (e.g., due to decoherence, noise floors, or discrete state resolution). Identifying regimes where finite distinguishability imposes a stricter persistence limit than classical dissipation is an open problem, and one where BCB may provide genuinely new constraints.

C5. Multiple BCB constraints and near-equalities

In realistic systems, multiple independent commitment budgets apply simultaneously (structural binding, phase coherence, thermal regulation). Each budget defines a BCB inequality and an admissible region. The intersection of these regions is typically narrow. Observables evaluated at this intersection acquire comparable magnitudes, producing near-equal ratios without fine-tuning.

C6. Talking about propagation without rates

Where language such as "faster" or "slower" propagation is used, it should be understood as referring to the number of irreversible commitments required per unit spatial advance or per structural update, not per unit time. Different physical interactions correspond to different commitment depths per advance, without introducing an external temporal parameter.

Appendix D: A Quantitative BCB from Known Physics

General reader version: Is there a real, physical limit on information that we can actually calculate? Yes—it's called the Bekenstein bound, and black holes exactly reach this limit.

D1. The Bekenstein bound as a universal capacity constraint

For any system of total energy E confined to a sphere of radius R, the Bekenstein bound states [3]:

$$S \leq 2\pi kER / (\hbar c)$$

Converting entropy to bits using $S = k \ln 2 \cdot I$ gives a BCB-style capacity bound:

$$I \leq 2\pi ER / (\hbar c \ln 2) = \mathcal{K}(E, R)$$

This is a quantitative BCB: it upper-bounds the number of irreversibly distinguishable bits that can be stored in a region of size R given total energy E.

What this means: Physics itself limits how much information can exist in a region. Pack too much information into too small a space, and you violate fundamental principles. This isn't a technological limit—it's a law of nature.

D2. Black holes saturate the bound

For a non-rotating black hole of mass M:

$$E = Mc^2,\ R = R_s = 2GM/c^2$$

Inserting into the Bekenstein bound:

$$I \leq 4\pi GM^2 / (\hbar c \ln 2)$$

The Bekenstein–Hawking black-hole entropy is [4, 5]:

$$S_{BH} = kA / (4\ell_P^2) = 4\pi kGM^2 / (\hbar c)$$

Converting to bits:

$$I_{BH} = S_{BH} / (k \ln 2) = 4\pi GM^2 / (\hbar c \ln 2)$$

A Schwarzschild black hole *exactly saturates* the Bekenstein information bound. In BCB language: a black hole is the maximally information-dense object allowed by known physics.

Translation: Black holes aren't just dense in the usual sense—they're maximally dense in *information*. They contain exactly as many bits as physics permits for their size and energy. They sit precisely at the edge of what's allowed.

D3. Concrete numbers

System	Mass	Information capacity
Solar-mass black hole	1.99×10^{30} kg	$\sim 10^{77}$ bits
Earth-mass black hole	5.97×10^{24} kg	$\sim 10^{66}$ bits

For perspective: The entire observable universe contains roughly 10^{80} atoms. A single solar-mass black hole has information capacity within a few orders of magnitude of this number. These figures illustrate both the enormity of the bound and the fact that black holes achieve it.

D4. Significance

This provides a concrete instantiation of the BCB idea: there exist universal, quantitative bounds on how many bits can be irreversibly distinguished within a region. In the body of the note, admissibility boundaries and their intersections explain why persistent structures generically live at constraint edges. Appendix D shows that at least one such edge is already known physics and is saturated in nature: black-hole horizons maximize information density.

Why this matters for the main argument: This isn't speculation. The Bekenstein bound is established physics. Black holes saturating it is established physics. The BCB framework is proposing that similar capacity limits—perhaps less extreme but structurally identical—govern ordinary persistent structures too.

D5. Ordinary Structures Need Not Saturate Bekenstein to Be BCB-Limited

General reader version: Moons and molecules don't need to be as information-dense as black holes for capacity limits to matter. They encounter tighter, local bottlenecks first.

The Bekenstein bound provides a universal ceiling on distinguishability capacity, saturated by black holes. Ordinary structures (moons, molecules, stars) are not expected to approach this ceiling. The BCB framework does not require saturation; it requires that persistence be constrained by *some* capacity $\mathcal{K}$, which in non-extremal systems is typically set by tighter, local bottlenecks (thermal noise, dissipation channels, decoherence, finite resolution of phase space).

Thus, black holes provide a proof of principle that (i) physically meaningful capacity bounds exist and (ii) nature can saturate them. The framework's claim is that many "everyday" admissibility boundaries are governed by non-extremal BCBs derived from the relevant local bottleneck, not by the global Bekenstein ceiling.

Empirical precedent for capacity-limited structure:

- **Landauer limit in computing [15]:** The minimum energy dissipation per bit erasure (kT ln 2) is a capacity bound approached in modern low-power experiments.
- **Shannon capacity in communications [16]:** Channel capacity bounds are approached with near-optimal coding schemes.

These are real, non-extremal capacity bounds that engineered systems actually approach. They demonstrate that capacity limits govern structure well below the Bekenstein extremum.

D6. Regimes of BCB Relevance

Physical systems fall into three broad categories with respect to bit-conservation limits:

(i) Extremal systems saturate a known universal capacity bound (e.g., black holes and the Bekenstein–Hawking entropy).

(ii) Engineered near-capacity systems are deliberately driven toward known bounds (e.g., Landauer-limited computation, Shannon-capacity communication).

(iii) Ordinary natural systems are constrained by local and typically sub-extremal capacity limits (e.g., decoherence thresholds, thermal noise floors, finite resolution of phase space).

The present work establishes (i) as proof of principle and shows that classical admissibility boundaries are compatible with a BCB formulation. Identifying unambiguous cases in category (iii) where bit conservation provides the dominant explanatory constraint—rather than merely being consistent with energetics—is an important direction for future work.

Promising candidates for future investigation:

- Ultracold quantum systems near decoherence thresholds
- Biological replicators near minimal genome sizes
- Phase transitions with finite-resolution order parameters

These represent systems where finite distinguishability may impose constraints tighter than classical energetics alone.

Appendix E: Falsifiable Predictions

General reader version: A good scientific theory must make predictions that could turn out to be wrong. Here are three specific predictions from this framework.

The admissibility-boundary framework makes quantitative predictions that can be tested against observation.

E1. Exomoon eclipse statistics (M-dwarf systems)

Prediction: For Earth-mass planets with Moon-sized satellites orbiting stars with $M_* < 0.3\ M_\odot$, no dynamically stable satellite orbit permits $0.9 \leq \theta_{Moon}/\theta_* \leq 1.1$.

Test: Future exomoon surveys (e.g., via transit timing variations or direct imaging [12, 13]) that identify such systems would falsify the boundary-intersection explanation for eclipse near-equalities.

In plain terms: If we find a Moon-like satellite around an Earth-like planet around a small red dwarf star, and that moon produces Sun–Moon-like eclipse near-equality (appearing the same size as the star to within ±10%), the framework is wrong. The framework predicts this is geometrically impossible.

E2. Research Direction: BCB Contribution to Baryon Retention Thresholds

Programmatic prediction: If early-universe structure formation is BCB-constrained, the minimum halo mass capable of retaining baryons should correlate with the Bekenstein capacity at the virial radius.

Research program: The capacity at the virial radius scales as $\mathcal{K} \propto ER \sim (M_{vir} c^2) R_{vir}$. A testable BCB contribution would require modeling the minimum distinguishability budget $\mathcal{B}$ for multiphase baryons (scaling with baryon mass, temperature resolution, and phase-space granularity). We outline this as a research direction rather than a fixed quantitative prediction, as the functional form of $\mathcal{B}(M)$ requires further development.

Test: Observations of ultra-faint dwarf galaxies near the cosmic baryon retention threshold provide a test of whether the cutoff reflects a capacity limit rather than purely cooling-driven physics. A BCB contribution would predict correlations between retention threshold and information-theoretic quantities (entropy, phase-space density) beyond what cooling models alone specify.

In plain terms: The smallest galaxies that can hold onto their gas should have a characteristic mass that reflects information-capacity limits, not just temperature and gravity. This is a promising research direction rather than a sharp prediction at present.

E3. Absence of coincidences in hierarchical systems

Prediction: In systems where control parameters span more than ~6 orders of magnitude, the framework predicts that order-one ratios should *not* generically appear at boundary intersections.

Test: Surveys of physical coincidences across domains should find that apparent fine-tuning correlates with moderate parameter ranges, not extreme hierarchies.

In plain terms: The framework predicts where coincidences *shouldn't* happen. If a system already has huge built-in disparities (like a trillion-to-one ratio), the framework doesn't predict that constraint boundaries will magically produce matching numbers. If we find order-one coincidences in highly hierarchical systems, the framework is wrong—or at least incomplete.

Summary

For the general reader: This paper argues that many apparent "coincidences" in physics aren't lucky accidents or evidence of cosmic design—they're the natural consequence of where stable things can exist. Like water pooling in the lowest available spot, physical systems settle at the edges of what's permitted, and when multiple constraints intersect, the numbers naturally line up.

This note has formalized a general physical principle: persistent structures occupy admissibility boundaries, not parameter-space interiors. When multiple such boundaries intersect, near-equalities between otherwise unrelated scales emerge without fine-tuning.

The framework:

1. **Explains** why near-equalities recur across physics (Section 6)
2. **Demonstrates** the mechanism through classical examples (Section 3, Appendix A)
3. **Derives** a classical boundary (Roche limit) from time-free BCB principles (Section 3.1)
4. **Shows** why logarithmic boundary dependence is generic, not assumed (Section 2.2)
5. **Establishes** that BCB excludes configurations that are dynamically stable but cannot persist as facts, while identifying where this distinction becomes quantitative (Section C4.1)
6. **Makes testable predictions** about where near-equalities should and should not occur (Appendix B, E)
7. **Grounds** the principle in time-free bit conservation capacity (Appendix C, D)
8. **Unifies** disparate constraint types (Roche, Hill, Bekenstein) under a common informational criterion

The framework predicts windows of near-equality rather than permanent attractors, and explains both the presence and absence of apparent coincidences depending on the structure of the underlying parameter space.

References

[1] Roche, É. (1849). "Mémoire sur la figure d'une masse fluide soumise à l'attraction d'un point éloigné." *Académie des sciences de Montpellier*, 1, 243–262.

[2] Chandrasekhar, S. (1969). *Ellipsoidal Figures of Equilibrium*. Yale University Press.

[3] Bekenstein, J. D. (1981). "Universal upper bound on the entropy-to-energy ratio for bounded systems." *Physical Review D*, 23(2), 287–298.

[4] Bekenstein, J. D. (1973). "Black holes and entropy." *Physical Review D*, 7(8), 2333–2346.

[5] Hawking, S. W. (1975). "Particle creation by black holes." *Communications in Mathematical Physics*, 43(3), 199–220.

[6] Williams, J. G., & Boggs, D. H. (2016). "Secular tidal changes in lunar orbit and Earth rotation." *Celestial Mechanics and Dynamical Astronomy*, 126(1), 89–129.

[7] Dickey, J. O., et al. (1994). "Lunar laser ranging: A continuing legacy of the Apollo program." *Science*, 265(5171), 482–490.

[8] Walker, J. C., & Zahnle, K. J. (1986). "Lunar nodal tide and distance to the Moon during the Precambrian." *Nature*, 320(6063), 600–602.

[9] Hill, G. W. (1878). "Researches in the lunar theory." *American Journal of Mathematics*, 1(1), 5–26.

[10] Domingos, R. C., Winter, O. C., & Yokoyama, T. (2006). "Stable satellites around extrasolar giant planets." *Monthly Notices of the Royal Astronomical Society*, 373(3), 1227–1234.

[11] Boyajian, T. S., et al. (2012). "Stellar diameters and temperatures. II. Main-sequence K- and M-stars." *The Astrophysical Journal*, 757(2), 112.

[12] Kipping, D. M. (2009). "Transit timing effects due to an exomoon." *Monthly Notices of the Royal Astronomical Society*, 392(1), 181–189.

[13] Teachey, A., & Kipping, D. M. (2018). "Evidence for a large exomoon orbiting Kepler-1625b." *Science Advances*, 4(10), eaav1784.

[14] Peters, P. C., & Mathews, J. (1963). "Gravitational radiation from point masses in a Keplerian orbit." *Physical Review*, 131(1), 435–440.

[15] Landauer, R. (1961). "Irreversibility and heat generation in the computing process." *IBM Journal of Research and Development*, 5(3), 183–191.

[16] Shannon, C. E. (1948). "A mathematical theory of communication." *Bell System Technical Journal*, 27(3), 379–423.

[17] Giudice, G. F. (2008). "Naturally speaking: The naturalness criterion and physics at the LHC." In *Bentivegna, E., Ferrara, S., & Zumino, B. (Eds.), Perspectives on LHC Physics* (pp. 155–178). World Scientific.

[18] Williams, P. (2015). "Naturalness, the autonomy of scales, and the 125 GeV Higgs." *Studies in History and Philosophy of Science Part B*, 51, 82–96.

[19] Barrow, J. D., & Tipler, F. J. (1986). *The Anthropic Cosmological Principle*. Oxford University Press.