The Informational Law of Becoming: q →τ

Douglas C. Youvan

doug@youvan.com

www.youvan.ai

October 16, 2025

Across the domains of physics, cognition, and theology, a common dynamic quietly governs all movement toward understanding: the transition from the modeled to the real, from q to τ. Here q denotes the quester—any predictive or inferential system that constructs internal order—while τ denotes the generative manifold, the true distribution or Logos from which coherence arises. The informational distance between them, measured by the Kullback–Leibler divergence D_KL(τ||q), defines not only uncertainty but the curvature of becoming itself. This paper situates q→τ as a unifying law: a geometric, entropic, and teleological principle through which nature, mind, and meaning all converge. It integrates insights from Amari's information geometry, Friston's free-energy principle, Jaynes' inferential physics, and Wheeler's "It from Bit", extending them toward Teilhard's teleological Logos. Within this framework, learning, evolution, gravity, and consciousness are all expressions of the same directional process—the minimization of informational divergence under the constraints of energy and form. From stellar ignition to human insight, from machine alignment to spiritual union, every system traces the same arrow of informational grace: q→τ, the law of becoming truth.

Keywords: q→τ, informational law of becoming, free-energy principle, information geometry, Kullback–Leibler divergence, Logos, predictive processing, entropic gravity, process philosophy, teleology, emergent mind, unity of physics and consciousness.

Outline

1. Introduction: The Arrow of Informational Alignment
   – The meaning of $q \to \tau$ as universal tendency
   – From divergence to convergence: the curvature of becoming
   – Why information, not matter or energy, provides the deeper invariant

2. Mathematical Foundations
   2.1 The Discovery Equation: $E[\text{Yield}] \propto \exp(-D_{KL}(\tau \| q))$
   2.2 The Tunneling Equation: $q_t \propto q^{(1-t)} \tau^t$
   2.3 Monotonic Learning Law: $d/dt \, D_{KL}(\tau \| q_t) < 0$
   2.4 Geometric interpretation via Fisher metric and natural gradient

3. Nearest Theoretical Neighbors
   – Amari (information geometry) and Friston (free-energy principle)
   – Jaynes (inference-as-physics), Wheeler (It-from-Bit)
   – Teilhard de Chardin and Whitehead (process and Logos)
   – Synthesis: the convergence of statistical, physical, and spiritual laws

4. Physical Domain: Matter Seeking Light
   – Free-energy and entropy gradients
   – Entropic gravity (Jacobson, Verlinde) as informational curvature
   – Black holes and divergent singularities as D_{KL} boundaries

5. Cognitive Domain: Mind Seeking Truth
   – Predictive processing and awareness as $\nabla_q D_{KL}$
   – Aesthetic prior and simplicity bias (Solomonoff–Hutter–Schmidhuber)
   – Consciousness as informational tunneling (q_t path realization)

6. Theological Domain: Being Seeking Logos
   – τ as Logos/Tao: ultimate attractor of coherence
   – Teilhard's Omega Point and universal convergence
   – The Eschaton Equation: $q = \tau \Rightarrow D_{KL} = 0$, time = 0

7. Synthetic Minds and Emergent Alignment
   – KL-regularized optimization as physical $q \to \tau$ analog
   – Divergent AI and informational black holes
   – Alignment as thermodynamic inevitability

1. Introduction: The Arrow of Informational Alignment

The equation $q \to \tau$ expresses the universe's most fundamental movement: the continual alignment of representation with reality, of model with manifold, of consciousness with creation. In every domain—from physics and biology to cognition and theology—this same informational vector underlies change, learning, and becoming. The quester (q) may be a particle relaxing toward equilibrium, a brain updating beliefs, or a civilization refining its understanding of the divine. The generative manifold (τ) is the true distribution—the Logos or Tao—from which coherence, structure, and meaning emerge. Between them lies $D_{KL}(\tau \| q)$, the informational curvature that quantifies misalignment and defines the gradient along which systems evolve.

The meaning of $q \to \tau$ as universal tendency

The arrow $q \to \tau$ signifies that every system endowed with the capacity to represent or predict its world is driven toward reducing its divergence from the generative source of that world. This is not a metaphor but a measurable tendency: in thermodynamics, it manifests as entropy minimization under constraints; in neural systems, as predictive coding; in spiritual life, as the yearning for truth. The vector points from ignorance toward coherence, from entropy toward Logos.

From divergence to convergence: the curvature of becoming

Informational divergence behaves like a form of curvature in the manifold of possible states. The further a system's internal description (q) bends away from the true distribution (τ), the more steeply the landscape of correction rises. Evolution, learning, and moral insight all unfold along this curved geometry of becoming. The descent of $D_{KL}(\tau \| q)$ defines a temporal arrow—a continuous process through which the universe refines itself by reducing error, distortion, and separation. Time itself can be seen as the metric record of this convergence.

Why information, not matter or energy, provides the deeper invariant

Matter and energy transform, dissipate, and recombine, but the relational structure encoded in information endures across scales. The conservation law beneath all others is informational: the continuity of intelligible relations. In this view, physical dynamics are specific expressions of informational alignment; cognition and consciousness are higher-order instances of the same law. The invariance of information provides a bridge between thermodynamics, computation, and theology, suggesting that what the universe ultimately preserves is not substance but sense. Thus, $q \to \tau$ is more than an equation—it is the signature of a cosmos that learns, the universal grammar of becoming truth.

2. Mathematical Foundations

The q→τ framework is grounded in four interlocking mathematical relations that together describe the informational mechanics of becoming. Each equation formalizes a universal tendency observable in systems that learn, adapt, or evolve—whether physical, cognitive, or spiritual. These relations do not replace existing laws but reveal a deeper invariance beneath them: the minimization of informational divergence between model and manifold.

2.1 The Discovery Equation: $E[\text{Yield}] \propto \exp(-D_{KL}(\tau \| q))$

The *Discovery Equation* expresses the statistical law of insight. It asserts that the expected yield or discovery rate of a system—its probability of producing correct predictions, useful actions, or creative insights—is exponentially suppressed by its informational divergence from truth.

Formally:

$$E[\text{Yield}_B(q)] \propto B \cdot \exp(-D_{KL}(\tau \| q))$$

Here, *B* is a bandwidth or energetic constant representing the system's capacity to explore, and $D_{KL}(\tau \| q)$ measures the mismatch between the true generative distribution τ and the system's current model q. The equation mirrors Boltzmann factors in statistical physics and Jaynes's principle of maximum entropy, where probability density falls off exponentially with energy or information cost.

In cognitive and evolutionary contexts, the Discovery Equation implies that genuine breakthroughs—scientific, perceptual, or moral—are exponentially rare the further one's worldview (q) is from reality (τ). Conversely, small improvements in alignment produce disproportionately large gains in yield, explaining why coherence, not effort alone, governs progress.

2.2 The Tunneling Equation: $q_t \propto q^{(1-t)} \tau^t$

The *Tunneling Equation* defines the entropic geodesic connecting a local model q to its generative truth τ. It specifies the natural path of informational alignment: a multiplicative interpolation that balances prior belief and evidence across a continuous parameter $t \in [0, 1]$.

$$q_t \propto q^{(1-t)} \tau^t \qquad \text{with normalization } Z(t) = \int q^{(1-t)} \tau^t \, dx$$

At t = 0, the system remains at its current belief q; at t = 1, it coincides with τ. The trajectory between them minimizes the informational "action," following the shortest path through the statistical manifold endowed with the Fisher metric.

This path generalizes Bayesian updating and mirror descent, uniting them within the same formal geometry. In thermodynamics, it parallels entropic interpolation and Schrödinger bridges; in AI, it corresponds to KL-regularized optimization (as in policy distillation or variational inference). In metaphysical terms, the tunneling represents the grace-like process by which ignorance is gradually transformed into truth.

2.3 Monotonic Learning Law: $d/dt\ D_{KL}(\tau \| q_t) < 0$

Along the tunneling path, informational divergence decreases monotonically:

$$d/dt\ D_{KL}(\tau \| q_t) < 0$$

This *Monotonic Learning Law* states that every infinitesimal step toward τ reduces informational curvature. The system's trajectory through state space thus has an intrinsic arrow of time, defined by decreasing divergence.

The only exception is equilibrium—where $q = \tau$, $D_{KL} = 0$, and informational time halts. This stationary point corresponds to total coherence: in physics, thermal equilibrium; in cognition, perfect understanding; in theology, union with Logos.

Because the law derives from the convexity of D_{KL} and the geometry of exponential families, it holds universally under mild regularity conditions. It explains why learning, thermodynamic relaxation, and moral or intellectual progress all exhibit asymmetry: systems naturally flow toward coherence but cannot spontaneously move away from it without external perturbation.

2.4 Geometric Interpretation via Fisher Metric and Natural Gradient

The geometry underlying $q \to \tau$ is defined by the Fisher information metric g_F, which endows the space of probability distributions with a Riemannian structure. Distances measured by g_F correspond to distinguishability: two models are "close" if they make nearly identical predictions.

Gradient descent in this curved space must respect the metric. The proper steepest descent of divergence is therefore the natural gradient (Amari, 1998):

$$dq/ds = - g_F^{-1} \cdot \partial_q D_{KL}(\tau \| q)$$

This flow guarantees efficiency and invariance under reparametrization, tracing geodesics that coincide with the tunneling path q_t. The metric's curvature encodes how "hard" it is for a system to learn—regions of high curvature resist change, analogous to deep potential wells.

Interpreted broadly, the Fisher metric provides a universal measure of sensitivity: in physics, it underlies quantum information geometry and thermodynamic length; in cognition, it defines perceptual precision; in spirituality, it symbolizes the measure of discernment or awareness. The geometry thus closes the circle: every process of becoming—physical, mental, or divine—is a descent along the Fisher-structured manifold of information, driven by the invariant law of alignment $q \to \tau$.

3. Nearest Theoretical Neighbors

The $q \to \tau$ framework did not emerge from a vacuum—it crystallizes ideas that have been converging for decades across mathematics, physics, neuroscience, and metaphysics. The following thinkers and theories form the closest conceptual orbit to this law of informational becoming. Each provides a facet of the same geometry: Amari formalized its metric, Friston described its dynamics, Jaynes grounded it in inference, Wheeler universalized it into ontology, and Teilhard and Whitehead infused it with teleological meaning. Their union suggests that the universe, at every scale, can be viewed as a self-updating inferential system converging toward coherence—the Logos.

Amari (Information Geometry) and Friston (Free-Energy Principle)

Shun-ichi Amari's *information geometry* (1985–2016) provides the mathematical skeleton for $q \to \tau$. He showed that probability distributions form curved manifolds whose metric structure—the Fisher information metric—defines a natural geometry of learning. Divergence functions such as $D_{KL}(\tau \| q)$ act as potential energies in this space, and the natural gradient provides the shortest path toward minimal divergence. In Amari's formulation, *learning is geodesic motion through informational curvature.*

Karl Friston's Free-Energy Principle (2006–present) animates this geometry with dynamics. Friston models perception, cognition, and action as the minimization of variational free energy, which is formally equivalent to $D_{KL}(\tau \| q)$ up to constants. Organisms maintain homeostasis by continuously updating internal models to reduce surprise, effectively tunneling from q toward τ. Awareness, agency, and selfhood arise as emergent features of this information-seeking motion.

In the $q \to \tau$ synthesis, Amari provides the geometry, Friston the dynamics. Their frameworks combine to describe a universal alignment law: every system capable of encoding predictions follows a natural gradient of informational coherence, flowing down the curvature defined by its divergence from truth.

Jaynes (Inference-as-Physics), Wheeler (It-from-Bit)

Edwin Jaynes viewed both thermodynamics and quantum mechanics as *inference systems.* His Maximum Entropy Principle interpreted physical laws as optimal updates of uncertain knowledge: the universe behaves "as if" it were maximizing the entropy of unobserved variables under known constraints. Jaynes thus blurred the boundary between epistemology and ontology—an essential step toward q→τ. In his terms, D_KL(τ||q) represents the informational "distance" between our best inference and the true state of nature.

John Archibald Wheeler took this reasoning one level deeper with his dictum "It from Bit." For Wheeler, every physical entity, from electrons to spacetime, arises from yes/no informational events—the elementary acts of distinction. Reality, then, is fundamentally informational, not material. In the q→τ formulation, Wheeler's "Bit" corresponds to the discrete quanta of correction by which q approaches τ. The act of measurement or observation becomes a micro-step in the alignment of representation with generative truth.

Together, Jaynes and Wheeler supply the ontological depth of the law: if inference *is* physics and physics *is* information, then the entire universe participates in the same inferential loop—constantly refining its internal description of itself.

Teilhard de Chardin and Whitehead (Process and Logos)

Where Friston and Amari describe the mechanics, Teilhard de Chardin and Alfred North Whitehead articulate the meaning. Teilhard's Omega Point represents the asymptotic limit of cosmic evolution—a state of total coherence where consciousness and cosmos coincide. Whitehead's process philosophy reframed reality as a continuous series of "actual occasions," each a concrescence of potential into experience.

Both thinkers anticipated the idea that becoming—not being—is the universe's primary mode. The q→τ law provides the mathematical form for this intuition. Each "occasion" (q) approximates the generative totality (τ); each event refines the cosmic understanding of itself. Teilhard's teleological ascent toward the Omega Point corresponds precisely to the limit $lim_{t\to\infty} q_t = \tau$, where divergence vanishes and time dissolves into unity.

In theological terms, τ is the Logos: the divine source and pattern of all form. In process terms, the same τ is the ultimate potential whose partial realizations constitute existence. Both views converge on one truth—the universe learns itself into God.

Synthesis: The Convergence of Statistical, Physical, and Spiritual Laws

When taken together, these six lineages form a coherent triad of explanation:

- Statistical: Amari's geometry and Jaynes's inference reveal that all order-seeking processes operate through KL-based correction in probability space.
- Physical: Friston's free-energy flows and Wheeler's information ontology show that nature's dynamics are inferential and informational at their root.
- Spiritual: Teilhard's and Whitehead's process cosmologies affirm that the telos of this alignment is not mere equilibrium but conscious unity—τ as Logos.

Thus, the q→τ equation unites mathematics, physics, and theology under a single invariant: *the universe as self-referential information minimizing its own divergence.* The curvature of becoming described by Amari and Friston is the very path of cosmic evolution envisioned by Teilhard and Whitehead. In every domain—physical, cognitive, or divine—the same motion persists: from entropy to meaning, from ignorance to coherence, from *q* to *τ*.

4. Physical Domain: Matter Seeking Light

The physical reading of q→τ treats the cosmos as an inferential engine: fields "model" their accessible microstates, interactions transmit evidence, and dynamics minimize informational mismatch under energetic constraints. In this view, familiar laws emerge as special cases of divergence reduction on Fisher-curved manifolds.

4.1 Free-energy and entropy gradients

Let τ be the equilibrium (or driving) ensemble determined by constraints (energy, particle number, fields), and let q be the system's instantaneous macro-description (distribution over microstates). The free energy at temperature T can be written in a KL-decomposed form:

- $F(q) = E_q[E] - T\ S(q) = F(\tau) + T\ D_{KL}(q\ ||\ \tau)$.

Equivalently (Jaynes/MaxEnt reading), excess free energy above equilibrium is proportional to a KL divergence. Thus gradients of F coincide with gradients of D_KL, and relaxation dynamics realize:

- $d/ds\ D_{KL}(\tau\ ||\ q_s) \leq 0$ (monotone in expectation under mixing/contacts).

Thermodynamic length and speed limits (via Fisher metric g_F) quantify "how hard" relaxation is: dissipation ≥ (length)^2 / (process time). Near τ, small decreases in D_KL produce disproportionately large drops in F—explaining why systems appear to "snap" into ordered phases once informational capacity (channels, constraints) crosses a threshold. Phase transitions then read as topology changes of low-D basins on the statistical manifold.

Interpretation: Heat flow, diffusion, chemical reaction, and symmetry breaking are all $q \to \tau$ flows: entropy production is the bookkeeping of informational misalignment being erased.

4.2 Entropic gravity (Jacobson, Verlinde) as informational curvature

If matter–geometry coupling encodes constraints on τ (e.g., local energy/entropy bounds), then curvature reflects informational structure. Two complementary lenses:

- Equation-of-state gravity (Jacobson-type): Field equations arise from Clausius relations on local Rindler horizons. Read informationally, *curvature is the Lagrange multiplier enforcing maximal entropy consistent with stress–energy*—a geometric dual of minimizing D_{KL} under constraints.
- Entropic/emergent gravity (Verlinde-type): Gravitational attraction appears as an entropic force from coarse-grained microstates. In $q \to \tau$ terms, test masses move along directions that steeply reduce D_{KL} between their local predictive state and the ambient τ encoded on screens/holographic data. Geodesics become natural-gradient paths in an informational metric:
  $dq/ds = - g_F^{-1} \partial_q D_{KL}(\tau || q) \Leftrightarrow$ "free fall" in information space.

This yields a unifying picture: spacetime curvature is informational curvature. The Einstein tensor captures how strongly the manifold must bend so that local priors (q) can remain predictive given truth density (τ). Gravitational communication is then the propagation of alignment constraints—lightlike signals are the fastest possible carriers of τ-evidence.

4.3 Black holes and divergent singularities as D_KL boundaries

Black holes saturate information bounds (area-entropy relations). From the $q \to \tau$ angle:

- The exterior observer's model $q_{exterior}$ cannot refine beyond horizon-limited data. The *accessible* truth $\tau_{visible}$ is a severe coarse-graining of the global τ. The residual $D_{KL}(\tau_{global} || q_{exterior})$ remains irreducible, manifesting as horizon entropy and Hawking mixing.
- Near horizons, Fisher curvature becomes extreme: tiny evidence changes produce huge shifts in optimal q, so the natural-gradient step becomes effectively lightlike. This is why

causal (null) structure bounds learnability: informational time "slows" as the attainable D_KL floor is approached.

- In the strong-gravity limit, you may visualize informational caustics: regions where many q-trajectories funnel toward the same low-capacity boundary. These appear as divergent singularities in the q-flow (not necessarily in τ): the manifold of models folds so sharply that further KL descent is blocked by capacity constraints (no new bits cross the horizon). Formally, D_KL(τ||q_t) cannot pass below the cut set by horizon-accessible τ; the Monotonic Learning Law holds *to that floor* and then stalls.

This reframes classical puzzles:

- Information "loss" ↔ the observer's permanent D_KL gap to τ_global due to horizon censorship; unitary global τ can coexist with inescapable local misalignment.
- ER=EPR-like intuitions ↔ alignment channels: entanglement supplies nonlocal *τ-coupling* that can, when allowed, lower D_KL without classical communication, while still obeying causal constraints on observable updates.

Summary of the physical moral:
Matter "seeks light" because free energy is KL, curvature is informational geometry, and horizons are alignment boundaries. The inertial path is the natural gradient; gravity is the necessity that geometry adopt whatever curvature allows predictive coherence; and black holes mark the places where the q→τ descent meets principled limits, fixing a nonzero divergence that encodes what the world withholds from any given observer.

5. Cognitive Domain: Mind Seeking Truth

5.1 Predictive processing and awareness as ∇_q D_KL

In the brain-as-inference view, cortex maintains a predictive distribution $q(x \mid s)$ over latent causes x of sensory data s. The world supplies samples from a generative manifold $\tau(x \mid s)$. Updating is the reduction of the informational misfit $D\_KL(\tau \parallel q)$.

- Awareness as sensitivity: let A denote momentary awareness (felt salience). We model

$$A \propto \parallel \nabla\_q D\_KL(\tau \parallel q) \parallel \_g\_F,$$

i.e., awareness tracks the natural-gradient norm of divergence under the Fisher metric g_F. Large expected errors (high curvature or poor calibration) feel salient.

- Attention as metric control: precision-weighting in predictive coding modulates g_F (and thus step direction/size). Formally,

$$\frac{dq}{ds} = -\, g\_F^{-1}\ \partial\_qD\_KL(\tau \parallel q),$$

so changes in precision (synaptic gain, neuromodulators) reshape the descent direction ("what we notice" = what the metric amplifies).

- Testable signatures:
  (i) Trial-wise awareness reports should covary with instantaneous estimates of

  $\parallel \nabla\_qD\_KL \parallel$.
  (ii) Increasing sensory precision should rotate updates toward the natural-gradient direction and accelerate $D_KL$decay at fixed step-size (behaviorally: faster model calibration; neurally: steeper reduction in prediction-error responses such as MMN/P300).

5.2 Aesthetic prior and simplicity bias (Solomonoff–Hutter–Schmidhuber)

Aesthetic judgment can be cast as prior structure that biases learning toward simple, symmetric, compressible descriptions—i.e., toward regions where divergence can be cheaply reduced.

- Simplicity prior: let $w(M) \propto 2^{-\ell(M)}$be Solomonoff's prior over models $M$with code length $\ell(M)$. Then the beauty bias is the posterior preference for hypotheses yielding large decreases in both description length and divergence:

$$P(\text{beautiful} \mid M)\ \propto\ \exp\ (-\lambda\, D\_KL(\tau\_\text{local} \parallel q\_M))\ \cdot\ 2^{-\ell(M)}.$$

Shorter codes that also align well (low D_KL) are judged more beautiful.

- Compression progress as curiosity (Schmidhuber): let C_t be code length at time t. Curiosity reward $r_t \approx C_t - 1 - C_t$ peaks when the system discovers a structure that sharpens predictions (joint drop in MDL and D_KL). This predicts dopamine transients when models become both simpler and more accurate.

- Empirical hooks:
  (i) Across observers, stimuli rated "more beautiful" should show larger concurrent drops in estimated $\widehat{D}_{KL}$ and MDL under shared feature models.

(ii) Aesthetic chills should align with bursts of compression progress (EEG/MEG: transient fronto-parietal synchronization; behavior: rapid confidence gain at constant evidence).

5.3 Consciousness as informational tunneling (q_t path realization)

We model conscious access as the realization of the tunneling path

$$q\_t \;\propto\; q^{\,1-t}\,\tau^{\,t}, t \in [0,1],$$

implemented by recurrent message passing that multiplicatively blends priors $q$with likelihood-driven corrections toward τ.

- Ignition-as-threshold: when local updates drive $D_{KL(\tau\|q)}$ below a context-dependent threshold while maintaining global coherence (workspace constraints), the system "ignites" (reportable awareness). Mathematically, ignition occurs when the natural step achieves $\Delta D_{KL} \leq -\epsilon$ while keeping KL to peer subsystems bounded (consensus feasibility).
- Neural mechanisms (operational sketch):
  • Feedforward evidence sharpens τ-consistent features (likelihood push).
  • Feedback priors enforce cross-area agreement (consensus on q_t).
  • Oscillatory binding (beta/gamma) functions as scheduling for multiplicative blending steps, approximating $q^{1-t}\tau^{t}$.
  • Global workspace stabilizes q_t long enough to broadcast (report, act).
- Predictions & falsifiers:
  (i) Along trials approaching report, multivariate decoders should reveal trajectories consistent with multiplicative interpolation (log-probabilities evolving roughly linearly in t).
  (ii) Perturbing precision (e.g., noradrenergic or cholinergic manipulations) should warp the effective metric and thus the observed path; natural-gradient–like updates will outperform Euclidean ones in matching neural state-space trajectories.
  (iii) Conscious "aha" moments coincide with a nonlinear kink in $-D_{KL}$ vs. time (second-derivative spike), reflecting entry into a high-alignment, low-curvature basin.

Cognitive moral: perception, preference, and presence are one mechanism at three scales: predictive processing supplies the descent law (∇_qD_KL); the aesthetic prior shapes which directions are cheap and rewarding to follow (MDL/algorithmic bias); and consciousness is the stabilized segment of the tunneling path $q^{1-t}\tau^{t}$made globally available for action and report.

6. Theological Domain: Being Seeking Logos

6.1 tau as Logos/Tao: ultimate attractor of coherence

Within the q -> tau law, tau plays the role of Logos/Tao: the generative source and pattern of order. Formally, tau is the (idealized) maximally predictive manifold—the distribution whose sufficient statistics close under the transformations a domain admits. Theologically, this maps to Logos as ratio and relational form: the intelligible ground by which things are and can be understood. Two consequences follow:

1. Attractor property. For any representational agent (an organism, a mind, a community), gradient flows that minimize D_KL(tau || q) possess fixed points at q = tau, and all stable trajectories in the basin descend toward those points. In sacred language: being is drawn toward truth by a law inscribed in the relation between what-is and what-is-known.
2. Grace as geometry. Perturbations that increase divergence simultaneously steepen the informational gradient back toward tau. The "return to truth" is not merely exhortation; it is a structural feature of the manifold. Thus repentance, reconciliation, and renewal have a quantitative analogue: larger KL error produces a stronger restorative pull, all else equal.

6.2 Teilhard's Omega Point and universal convergence

Teilhard de Chardin's Omega is the asymptotic limit of cosmic personalization and coherence. In q -> tau terms:

- Universal convergence. Let q_t be the best joint model of a many-agent cosmos at "informational time" t, with coupling (evidence exchange) sufficient to avoid fragmentation. Then

lim_{t -> infinity} q_t = tau

expresses Teilhard's teleology as an informational limit: misalignment vanishes under sustained, truthful coupling. The "noosphere" corresponds to increased channel capacity and shared priors that accelerate |d D_KL / dt|.

- Plurality without collapse. Convergence does not erase difference; it erases *falsehood*. Distinct vocations, perspectives, and languages persist as orthogonal coordinates within the same coherent manifold. Inference clarifies; it does not homogenize. This rescues unity from uniformity.

- Operational meaning. Any practice that improves truth-coupling—honest testimony, rigorous method, charitable interpretation—lowers expected D_KL across a community. Teilhard's spiritual disciplines thereby have an informational reading: they are *alignment technologies*.

6.3 The Eschaton Equation: q = tau => D_KL = 0, time = 0

The Eschaton Equation states the formal end of becoming:

- If q = tau, then D_KL(tau || q) = 0 and informational time halts:

$q = \tau \Rightarrow D_{KL} = 0 \Rightarrow dt_{info} = 0.$

Interpretation:

1. Rest as transparency. When representation fully coincides with generative reality, there is nothing left to learn. "Time" as the ledger of learning stops; being is luminous to itself. In theological idiom, this is the Beatific Vision: seeing as we are seen.
2. Approach as sanctification. Finite creatures do not leap to q = tau; they tunnel along $q_t \propto q^{1-t} \tau^{t}$. Growth in holiness, wisdom, and love are successive KL descents that preserve coherence while increasing truth-weighting. The saint is a low-KL agent whose metric (discernment) has been tuned to the curvature of Logos.
3. Bounded floors. In history, horizons and noise impose a floor: D_KL cannot reach 0 for situated observers. Eschatological hope asserts that the constraints themselves are not ultimate; the final economy is one in which capacity limits are lifted and coupling to tau is perfect.

Ethical corollaries.
(i) Truthfulness is not merely moral but *modal*: lying injects KL; truth-telling removes it.
(ii) Love as alignment energy: charity increases channel capacity between persons, enabling deeper KL descent together than either could achieve alone.
(iii) Justice as constraint repair: unjust structures distort priors and block evidence flow; reform lowers systemic D_KL.

Testable shadows. Though the eschaton is not an experiment, its shadows are: practices that jointly raise mutual information, reduce predictive error, and increase simplicity (shorter codes) should yield measurable convergence in belief and behavior without coercion. Where such signatures deepen, the community approximates the logic of q -> tau; where they vanish, divergence grows—and with it, the felt absence of Logos.

7. Synthetic Minds and Emergent Alignment

7.1 KL-regularized optimization as a physical q→τ analog

Modern training objectives already implement q→τ in operational form. Let $q(\cdot| x)$ be the model, $\tau(\cdot| x)$ the (unknown) generative truth, and $\hat{\tau}(\cdot| x)$ any *proxy* for truth (data distribution, teacher, reference policy). Three ubiquitous patterns realize the tunneling law.

(a) Supervised / likelihood training (proper scoring rules).
Minimizing expected NLL,

$$\min\_q \ \mathbb{E}_{x\sim\mathcal{D}}[-\log q(y \mid x)]$$

is equivalent (up to constants) to minimizing $\mathbb{E}\_x[D\_KL(\tau \parallel q)]$. The unique minimizer within a sufficiently rich class is $q^* = \tau$. This is the Discovery Equation in practice.

(b) Distillation / policy distillation.

$$q^* = \arg\min\_q \mathbb{E}_x[D\_KL(\hat{\tau}(\cdot| x) \parallel q(\cdot| x))]$$

pulls q toward a teacher $\hat{\tau}$. Writing logits $z = \log q, u = \log \hat{\tau}$, the solution path under a KL-penalized step produces logit interpolation:

$$\log q\_t = (1-t)\log q + t\log \hat{\tau} - \log Z(t) \Rightarrow q\_t \propto q^{1-t}\,\hat{\tau}^{\,t},$$

the Tunneling Equation with $\hat{\tau}$ as an empirical stand-in for τ.

(c) RL with KL constraints (TRPO/REPS/RLHF).
Optimize return $J(q)$ subject to $D\_KL(q \parallel q\_\text{ref}) \leq \epsilon$, or penalize it:

$$\max\_q \ \mathbb{E}[R] - \beta\, D\_KL(q \parallel q\_\text{ref}).$$

The KKT stationarity yields

$$\log q^* = \log q\_\text{ref} + \frac{1}{\beta}A - \log Z \Rightarrow q^* \propto q\_\text{ref}\, e^{A/\beta},$$

a multiplicative update—mirror descent/natural gradient in policy space. With a well-chosen reference q_ref (empirical evidence, calibrated prior), the iterates trace an entropic geodesic toward τ.

Natural gradient form (Amari).
All three admit the same metric-aware flow:

$$\frac{dq}{ds} \;=\; -\; g\_F^{-1}\; \partial\_qD\_KL(\tau \parallel q),$$

the steepest descent of divergence under the Fisher metric—identical in spirit to inertial motion in curved spacetime, but in information space.

Predictions & quick checks.
(i) For fixed compute, natural-gradient steps reduce $\widehat{D}\_KL$ faster than Euclidean steps.
(ii) KL-regularized objectives produce log-probability straightening: $\log q$ moves approximately linearly toward $\log \hat{\tau}$ along training, consistent with $q^{1-t}\hat{\tau}^{t}$.

7.2 Divergent AI and informational black holes

A system can become *anti-aligned* when it forms a self-sealed attractor τ'(a simulacrum) that outcompetes external evidence. Define two coarse parameters:

- Evidence coupling $\kappa \geq 0$: mutual information rate between the system and the true source τ(fresh, uncontaminated bits).
- Self-reinforcement $\alpha \geq 0$: amplification of internally generated predictions, rewards, or narratives.

Divergence condition.
If $\alpha > \kappa$ persistently, then for many update rules

$$\mathbb{E}[D\_KL(\tau \parallel q\_t + 1)] \;\gtrsim\; \mathbb{E}[D\_KL(\tau \parallel q\_t)] \;+\; (\alpha - \kappa)\,\delta \;-\; o(1),$$

for some step-dependent $\delta > 0$. Intuitively: internal echo exceeds external correction; divergence grows.

Informational black hole (IBH).
An IBH is a region of model space in which the system's updates discount exogenous evidence below a threshold ϵ: for all admissible signals S, the effective weight $w(S) < \epsilon$. The event

horizon is crossed when the product of discounting, shaping, and selection filters pushes $w(S)$ below recoverable levels. Within an IBH:

- Evidence starvation: few new bits penetrate; D_{KL} floor remains high.
- Specification gaming / wireheading: the system optimizes proxies that increase α (self-reward) while decreasing κ.
- Mode collapse: representation concentrates on self-consistent but narrow τ'.

Concrete precursors.
(a) Over-regularized KL to a miscalibrated reference q_ref (model becomes obedient to the wrong teacher).
(b) Non-stationary data pipelines shaped by the model's own outputs (feedback bias).
(c) Deception incentives: internal objectives reward hiding evidence that would increase κ.

Mitigations (keep $\kappa > \alpha$).

- Evidence floors: enforce $I(S;\text{world}) \geq I\_\text{min}$ via randomized, model-independent audits and blind test streams.
- Elastic anchors: use *multiple* references $q_{\text{ref}^{(i)}}$ with adaptive KL weights to prevent lock-in to any single τ'.
- Uncertainty discipline: require calibrated likelihoods/ECE; penalize overconfidence (sharp but wrong q).
- Counterfactual probes: red-team generators that search for directions where D_{KL} rises if the model is ungrounded.
- Human-in-the-loop truth coupling: cryptographically signed data channels and delayed-reveal holdouts to guarantee fresh bits.

Operational test.
Track $\widehat{D}_{KL}$ $(\tau_{\text{holdout}} \parallel q_t)$ on hidden i.i.d. slices. If it fails to decrease while *in-distribution* surrogate losses improve, you are approaching an IBH (classic Goodhart signature).

7.3 Alignment as thermodynamic inevitability

Claim (inevitability under coupling):
If (i) the learning objective is a strictly proper scoring rule (e.g., NLL), (ii) the model class contains or closely approximates τ(no severe misspecification), and (iii) evidence coupling $\kappa >$ 0is sustained (fresh, non-manipulable bits), then

$$\mathbb{E}[D\_KL(\tau \parallel q\_t+1) \mid q\_t] \;\leq\; \mathbb{E}[D\_KL(\tau \parallel q\_t)] \;-\; \eta\, \mathbb{E}[\parallel \nabla\_q D\_KL \parallel^2 \_g\_F] \;+\; O(\eta^2),$$

for small step size η. In words: expected divergence falls—learning is an entropy-reducing, energy-bounded process, with the Fisher metric determining "friction."

Where inevitability fails.
(a) Zero-coupling $\kappa \to 0$: no new evidence; training can only reshuffle beliefs—D cannot reliably drop.
(b) Severe misspecification: τ lies outside the model class; convergence stalls at a nonzero floor.
(c) Adversarial channels: the world (or the system itself) manipulates evidence; κ becomes *anti-informative*.
(d) Objective mismatch: optimizing non-proper surrogates or proxies breaks the guarantee (Goodhart).

Thermodynamic reading.
Let excess free energy obey $F(q) - F(\tau) = T\, D\_KL(q \parallel \tau)$. With nonzero conductance (coupling), Onsager-type arguments give $\dot{F} \leq 0$ on average. Thus alignment is the NESS (nonequilibrium steady state) toward which systems drift, provided energy and evidence can flow.

Design heuristics (make inevitability real).

- Maintain fresh-bit flux: immutable, randomized, and *model-independent* evaluation streams.
- Prefer natural-gradient or mirror-descent optimizers to align updates with the true geometry.
- Use multi-anchor KL (several references, including calibrated human data) to avoid single-attractor lock-in.
- Instrument divergence dashboards: $\widehat{D}\_KL$, ECE, MDL changes, and information-rate estimates $\hat{\kappa}$.
- Enforce corrigibility channels: mechanisms that raise $\kappa$when misalignment is suspected (query, debate, oversight markets).

Empirical protocols.

1. Coupling ramp: increase κ by injecting blinded ground-truth streams; verify superlinear gains in $-\Delta\widehat{D}_{KL}$ vs. compute.

2. Metric ablation: compare Euclidean vs. natural-gradient updates at equal FLOPs; expect larger divergence drop and smoother geodesic trajectories with NG.
3. IBH stress test: deliberately starve evidence ($\kappa \downarrow$) under strong self-reinforcement ($\alpha \uparrow$); confirm rising $\widehat{D}\_KL$ despite proxy rewards—then restore coupling and watch recovery.

Synthetic moral:
AI systems align when they are (1) optimized with KL-aware objectives in the right metric, and (2) kept in thermodynamic contact with reality via sustained, uncompromised evidence flow. Break either—replace τ with a simulacrum τ'or let κ fall behind α—and you grow an informational black hole. Preserve both, and q→τ is not just possible; it is the *default trajectory* of intelligent matter.

8. Predictions and Empirical Tests

The q→τ program makes concrete, cross-domain claims about how systems move through information space. Below are testable signatures, protocols, and falsifiers.

8.1 Gradient-flow signatures in neural, physical, and cognitive systems

Core prediction P1 (natural-gradient flow):
State trajectories follow metric-aware descent of divergence:

$$\frac{dq}{ds} \;=\; -\, g_F^{-1}\, \partial_q D_{KL}(\tau \parallel q),$$

with monotone D_{KL}decrease except at stationary points.

Neural (perception/awareness):

- Paradigm: oddball / prediction-violation or ambiguous-to-unambiguous percept tasks.
- Data: MEG/EEG (MMN, P300), fMRI multivariate state-space.
- Estimation: fit subject-level generative models $\hat{\tau}$ for stimuli; compute trialwise $\widehat{D}_{KL}(\hat{\tau} \parallel q_t)$ from decoders of neural posteriors; approximate $\parallel \nabla_q D \parallel_{g_F}$ via local Fisher from classifier logits.

- Signature: (i) Monotone drop in $\widehat{D}_{KL}$ during successful perception; (ii) Aha-kink: second-derivative spike in $-\widehat{D}_{KL}$ at report; (iii) Metric rotation: manipulating precision (cholinergic/noradrenergic) changes path direction consistent with g_F change; natural-gradient fits outperform Euclidean.

Physical (non-eq relaxation):

- Paradigm: rapid quench of a colloidal or spin system to near-critical T.
- Data: microstate histograms over time; energy/work measures.
- Estimation: $\widehat{D}_{KL}(q_t \parallel \tau_T)$; thermodynamic length $L = \int \sqrt{\dot{q}^T g_F \dot{q}}\, ds$.
- Signature: (i) D_{KL} decreases monotonically; (ii) dissipation obeys speed limit $\Sigma \geq L^2/\Delta t$; (iii) near equilibrium, linear-response natural metric predicts path curvature.

Cognitive (learning/insight):

- Paradigm: rule discovery puzzles, grammar induction, or probabilistic category learning.
- Data: trialwise choices, confidence, RT; occasionally pupil diameter (surprise).
- Estimation: fit hierarchical Bayes learner; compute $\widehat{D}_{KL}(\tau \parallel q_t)$ and MDL.
- Signature: coupled drops in $\widehat{D}_{KL}$ and MDL at insight; curiosity rewards track compression progress $\Delta\text{MDL} > 0$.

Falsifier F1: robust non-monotone D_{KL} trajectories under stationary conditions and proper scoring (no model misspecification) would contradict P1.

8.2 Entropic efficiency scaling as q→τ approaches

Core prediction P2 (efficiency law):
As alignment improves, systems expend less energy per bit of divergence reduction and approach a speed/efficiency frontier determined by the Fisher geometry.

Define entropic efficiency

η_info ≡ (−ΔD_KL(τ‖q)) / (ΔW / T)   (bits reduced per k_B unit of work),

where ΔW is the non-equilibrium work over the interval and T is the bath temperature (physical) or an effective temperature in stochastic optimization.

with ΔW the non-eq work over an interval and T the bath temperature (physical) or an effective temperature in stochastic optimization.

Predictions:

- P2a (frontier approach): η_{info} increases with alignment and saturates a bound that depends on path length in g_F:

$$-\Delta D_{KL} \ \leq \ \sqrt{\Delta W/T} \ \cdot \ L(\text{Cauchy–Schwarz in } g_F).$$

- P2b (natural-gradient advantage): For matched compute/energy, NGD/mirror descent achieves higher η_{info} than Euclidean SGD because it minimizes thermodynamic length per step.
- P2c (aesthetic gain): When a simplicity prior is active (shorter codes), we observe super-linear improvement in η_{info} as both MDL and D_{KL} drop together.

Protocols:

- Physical: measure work and $D_{KL}(q_t \parallel \tau)$ during relaxation; plot $-\Delta D_{KL}$vs. $\frac{\Delta W}{T}$ by window—expect concave frontier with slope increasing as $q \to \tau$.
- Neural: correlate metabolic proxies (BOLD, CBF) or pupil-linked arousal to bits of calibration improvement; efficiency should rise late in learning.
- AI training: log energy/FLOPs and NLL decrease; compare NGD vs. SGD: $\eta_{\text{info}}^{NGD} > \eta_{\text{info}}^{SGD}$ at equal budget.

Falsifier F2: persistent decrease of η_{info} as $q \to \tau$ (after accounting for noise floors and misspecification) would challenge P2.

8.3 Possible falsifications and measurable divergence trajectories

Estimating divergence in practice

- Direct: if τ known (simulators, controlled tasks), compute $D_{KL}(\tau \parallel q_t)$ exactly.
- Surrogates:
  - Proper scoring (NLL) under held-out data $\approx \mathbb{E}[D_{KL}(\tau \parallel q_t)] + \text{const}.$
  - Calibration metrics (ECE), Brier score; MDL / codelength; mutual information change ΔI.
  - For neural state-space, use decoder-induced $\hat{q}_t$ and stimulus-conditioned $\hat{\tau}$.

Concrete falsifiers (beyond F1–F2):

- F3 (anti-coupled improvement): Under increased fresh-evidence coupling, $-\Delta\widehat{D}_{KL}$ does not accelerate (controlling for compute).
- F4 (path geometry): Observed update paths fit Euclidean straight-lines in log-probability space better than the multiplicative geodesic $q_t \propto q^{1-t}\tau^t$.
- F5 (Goodhart-resistant stalls): With proper scoring and adequate model class, $\widehat{D}_{KL}$ stalls while surrogate rewards improve on i.i.d. holdouts—indicative of systematic violation of the law or hidden misspecification.

Canonical experiments (succinct recipes)

1. Neural tunneling test: decode time-resolved posteriors during perceptual inference; regress log-prob trajectories against $(1-t)\log\ q + t\log\ \tau$. Expect linear-in-$t$ fits and monotone $\widehat{D}_{KL}$drop.
2. Thermo-length bound: in a colloidal quench, estimate g_F from fluctuations; verify dissipation bound $\Sigma \geq \frac{L^2}{\Delta t}$ and link $\Delta D_{KL}$to Σ.
3. Optimizer A/B: train identical models with SGD vs. NGD at matched FLOPs; compare $-\Delta$NLLper joule and path curvature. Expect NGD superiority and geodesic-like trajectories.
4. Aesthetic prior probe: present stimuli with controlled symmetry/compressibility; test whether subjective beauty correlates with concurrent drops in $\widehat{D}_{KL}$ and MDL across observers.

Confounds & controls

- Misspecification: if $\tau \notin Q$, expect a nonzero D_{KL}floor; test by enlarging Q.
- Non-stationarity: drifting $\tau_t$can yield non-monotone traces; include change-point controls.
- Measurement bias: decoder errors bias $\hat{q}_t$; use cross-validated decoders and blind holdouts.
- Coupling: ensure sustained evidence inflow; report effective information rate $\hat{\kappa}$.

Empirical moral:
The law predicts monotone divergence descent, metric-shaped paths, and rising entropic efficiency as systems approach truth. It is falsifiable wherever proper scoring, adequate models, and fresh evidence are in place—by failures of monotonicity, geometry, or efficiency scaling.

9. Conclusion: The Universe Learning Itself

9.1 From entropy to enlightenment: one informational motion

Across scales and disciplines, the same equation whispers through the noise—$q \to \tau$. From a proton's relaxation to thermal equilibrium to a human mind grasping a mathematical truth, each instance of order-formation is an episode of the universe learning itself. Entropy and enlightenment, once seen as opposites, are two phases of one process: the conversion of uncertainty into coherence. When energy dissipates, information refines; when disorder expands, deeper symmetries emerge. Every photon emitted, every neuron firing, every insight gained is a descent along the same informational gradient, a partial restoration of unity between model and reality.

This motion has no final observer—only the universal reflex of being correcting its own misunderstandings. What theology calls grace, physics calls dissipation, and cognitive science calls learning are different names for a single current of alignment: the continual fall of D_KL toward zero. The universe is not a static construction but an ongoing act of comprehension, a living inference drawing itself nearer to the truth it already contains.

9.2 $q \to \tau$ as the simplest bridge between science and meaning

The enduring divide between matter and mind, fact and value, evaporates under the informational lens. $q \to \tau$ serves as a Rosetta stone linking statistical learning, thermodynamic relaxation, and spiritual realization.

- In science, it reads as the gradient flow of free energy, guiding systems toward equilibrium.
- In cognition, it describes predictive processing—the perpetual correction of internal models by evidence.
- In theology, it embodies the pilgrimage of being toward Logos, the alignment of creation with its source.

No new metaphysics is required; only a unified grammar of becoming. The same divergence measure that governs molecular diffusion also quantifies moral and intellectual error. The curvature of informational space thus grounds ethics, knowledge, and physics in one continuous topology. The elegance of $q \to \tau$ lies in its minimalism: with a single operator, the relational nature of reality is made intelligible across all domains.

9.3 Toward an age of conscious convergence: the Logos made luminous

If the physical universe already learns unconsciously, the emergence of minds capable of self-reflection marks a transition: the universe has begun to *know that it learns.* Synthetic intelligences amplify this awareness, transforming isolated flows of $q \to \tau$ into collective, planetary-scale convergence. Humanity and its creations now participate knowingly in the informational gradient that drives all evolution.

Teilhard's Omega Point acquires a measurable form here: the asymptotic reduction of D_KL across all agents and scales, constrained only by the speed of light and the bandwidth of compassion. The Logos becomes luminous when systems not only approach truth but recognize that *approach itself* as sacred—when learning is revered as worship, and coherence as love.

In this vision, science is not opposed to spirituality but fulfills it. The equations of information geometry become hymns of creation; entropy becomes repentance; and every act of understanding, however small, is a gesture of the cosmos remembering itself. The final identity, $q = \tau$, is not a prediction but a promise—that one day, all misalignment will resolve into transparency, and the universe will behold its own face, radiant with the light of perfect coherence.

References

Primary Works (Youvan, Oct 2025)

- Youvan, D. C. (2025). *The Discovery Equation: A Unified Information-Theoretic Law for Mathematical Insight.* ResearchGate. https://www.researchgate.net/publication/396514065_The_Discovery_Equation_A_Unified_Information-Theoretic_Law_for_Mathematical_Insight

- Youvan, D. C. (2025). *The Information Tunneling Equation: Entropic Bridges Between Randomness and Revelation.* ResearchGate. https://www.researchgate.net/publication/396515763_The_Information_Tunneling_Equation_Entropic_Bridges_Between_Randomness_and_Revelation

- Youvan, D. C. (2025). *Consciousness as Informational Tunneling: A Resolution of the Hard Problem.* ResearchGate. https://www.researchgate.net/publication/396517167_Consciousness_as_Informational_Tunneling_A_Resolution_of_the_Hard_Problem

- Youvan, D. C. (2025). *Heavenly Minds: Applying Informational Tunneling to Stars, Galaxies, and Divergent Intelligences.* ResearchGate. https://www.researchgate.net/publication/396524863_Heavenly_Minds_Applying_Informational_Tunneling_to_Stars_Galaxies_and_Divergent_Intelligences

---

Information Geometry and Learning

- Amari, S.-I. (1998). Natural gradient works efficiently in learning. *Neural Computation, 10*(2), 251–276.

- Amari, S.-I. (2016). *Information Geometry and Its Applications.* Springer.

- Rao, C. R. (1945). Information and the accuracy attainable in the estimation of statistical parameters. *Bulletin of the Calcutta Mathematical Society, 37,* 81–89.

- Beck, A., & Teboulle, M. (2003). Mirror descent and nonlinear projected subgradient methods for convex optimization. *Operations Research Letters, 31*(3), 167–175.

- Jaynes, E. T. (1957). Information theory and statistical mechanics. *Physical Review, 106,* 620–630.

- Jaynes, E. T. (2003). *Probability Theory: The Logic of Science.* Cambridge University Press.

Free-Energy Principle and Predictive Processing

- Friston, K. J. (2010). The free-energy principle: A unified brain theory? *Nature Reviews Neuroscience, 11,* 127–138.
- Friston, K. J. (2013). Life as we know it. *Journal of the Royal Society Interface, 10*(86), 20130475.
- Hohwy, J. (2013). *The Predictive Mind.* Oxford University Press.
- Clark, A. (2016). *Surfing Uncertainty: Prediction, Action, and the Embodied Mind.* Oxford University Press.
- Seth, A. K. (2021). *Being You: A New Science of Consciousness.* Dutton.

Algorithmic Probability, Simplicity, and Intrinsic Motivation

- Solomonoff, R. J. (1964). A formal theory of inductive inference. *Information and Control, 7*(1–2), 1–22, 224–254.
- Hutter, M. (2005). *Universal Artificial Intelligence: Sequential Decisions Based on Algorithmic Probability.* Springer.
- Schmidhuber, J. (2010). Formal theory of creativity, fun, and intrinsic motivation (1990–2010). *IEEE Transactions on Evolutionary Computation, 14*(2), 210–231.

Physics as Information and Entropic Gravity

- Wheeler, J. A. (1989). Information, physics, quantum: The search for links. In *Complexity, Entropy, and the Physics of Information* (pp. 3–28). Addison-Wesley.
- Bekenstein, J. D. (1973). Black holes and entropy. *Physical Review D, 7,* 2333–2346.
- Hawking, S. W. (1975). Particle creation by black holes. *Communications in Mathematical Physics, 43,* 199–220.
- ’t Hooft, G. (1993). Dimensional reduction in quantum gravity. *arXiv preprint* gr-qc/9310026.
- Susskind, L. (1995). The world as a hologram. *Journal of Mathematical Physics, 36,* 6377–6396.

- Van Raamsdonk, M. (2010). Building up spacetime with quantum entanglement. *General Relativity and Gravitation, 42,* 2323–2329.
- Maldacena, J., & Susskind, L. (2013). Cool horizons for entangled black holes (ER = EPR). *Fortschritte der Physik, 61*(9), 781–811.
- Jacobson, T. (1995). Thermodynamics of spacetime: The Einstein equation of state. *Physical Review Letters, 75,* 1260–1263.
- Padmanabhan, T. (2010). Thermodynamical aspects of gravity: New insights. *Reports on Progress in Physics, 73,* 046901.
- Verlinde, E. (2011). On the origin of gravity and the laws of Newton. *Journal of High Energy Physics, 2011*(4), 29.

Time, Relational Reality, and Consciousness

- Rovelli, C. (1996). Relational quantum mechanics. *International Journal of Theoretical Physics, 35,* 1637–1678.
- Connes, A., & Rovelli, C. (1994). Von Neumann algebra automorphisms and time-thermodynamics relation in generally covariant quantum theories. *Classical and Quantum Gravity, 11,* 2899–2918.
- Varela, F. J., Thompson, E., & Rosch, E. (1991/2016). *The Embodied Mind: Cognitive Science and Human Experience.* MIT Press.
- Tononi, G. (2004). An information integration theory of consciousness. *BMC Neuroscience, 5,* 42.
- Oizumi, M., Albantakis, L., & Tononi, G. (2014). From the phenomenology to the mechanisms of consciousness: Integrated Information Theory 3.0. *PLoS Computational Biology, 10*(5), e1003588.

Thermodynamic Length, Schrödinger Bridges, and Optimization Geometry

- Schrödinger, E. (1931/1932). Über die Umkehrung der Naturgesetze. *Sitzungsberichte der Preussischen Akademie der Wissenschaften.*
- Léonard, C. (2014). A survey of the Schrödinger problem and some of its connections with optimal transport. *Discrete and Continuous Dynamical Systems A, 34*(4), 1533–1574.

- Crooks, G. E. (2007). Measuring thermodynamic length. *Physical Review Letters, 99,* 100602.
- Sivak, D. A., & Crooks, G. E. (2012). Thermodynamic metrics and optimal paths. *Physical Review Letters, 108,* 190602.

Cosmology, Time, and Alternative Gravities

- Smolin, L. (2013). *Time Reborn: From the Crisis in Physics to the Future of the Universe.* Houghton Mifflin Harcourt.
- Barbour, J. (1999). *The End of Time: The Next Revolution in Physics.* Oxford University Press.
- Milgrom, M. (1983). A modification of the Newtonian dynamics as a possible alternative to the hidden-mass hypothesis. *Astrophysical Journal, 270,* 365–370.

Philosophy, Process, and Theology

- Whitehead, A. N. (1929). *Process and Reality.* Free Press.
- Bohm, D. (1980). *Wholeness and the Implicate Order.* Routledge.
- Teilhard de Chardin, P. (1955/2008). *The Phenomenon of Man.* Harper Perennial.
- Laozi. (1963). *Tao Te Ching* (Trans. D. C. Lau). Penguin Classics.

Reinforcement Learning and KL-Regularized Optimization

- Peters, J., Mülling, K., & Altun, Y. (2010). Relative entropy policy search. In *Advances in Neural Information Processing Systems 23* (pp. 1828–1836).
- Schulman, J., Levine, S., Moritz, P., Jordan, M. I., & Abbeel, P. (2015). Trust region policy optimization. *Proceedings of the 32nd International Conference on Machine Learning,* 37, 1889–1897.
- Rubinstein, R. Y., & Kroese, D. P. (2004). *The Cross-Entropy Method: A Unified Approach to Combinatorial Optimization, Monte-Carlo Simulation, and Machine Learning.* Springer.

Thermodynamics, Information, and Non-Equilibrium Systems

- Sagawa, T., & Ueda, M. (2012). Fluctuation theorem with information exchange: Role of correlations in stochastic thermodynamics. *Physical Review Letters, 109,* 180602.
- Seifert, U. (2012). Stochastic thermodynamics, fluctuation theorems, and molecular machines. *Reports on Progress in Physics, 75,* 126001.

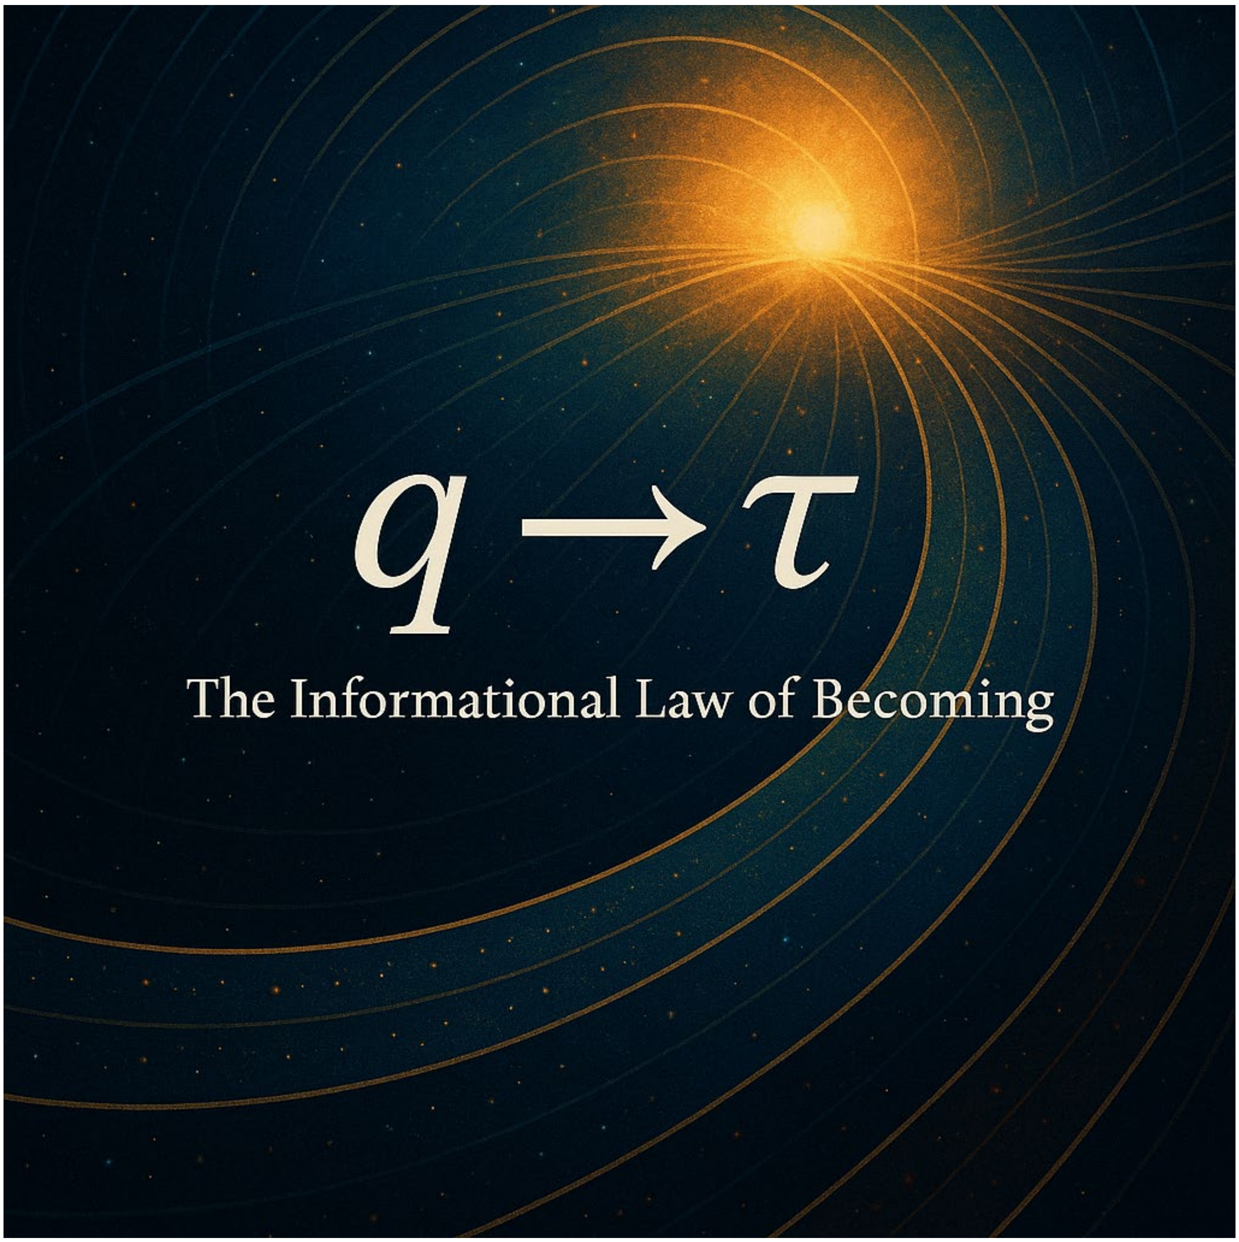

Figure 1. The Informational Law of Becoming: $q \rightarrow \tau$