

The Attractor Framework in Astrophysics: Persistence, Entropy, and Gravitational Systems; Robert Galida (July 2026) [A]

Abstract

The attractor framework provides a domain-general vocabulary for describing persistence and change across physical, biological, cognitive, and social systems. This paper extends the framework to astrophysical dissipative systems. We distinguish between conservative gravitational dynamics – which define families of stable invariant solutions – and dissipative processes – which select and can stabilize particular configurations within those families.

The central thesis is:

Gravity defines the landscape. Dissipation selects the configuration.

We provide an operational definition of the excess entropy production functional σ_{excess} for gravitational systems, grounding the persistence functional $D_{\infty} = \int \sigma_{\text{excess}} dt$ in physical dissipation rates above steady-state baselines. We show that:

- Orbital circularization is a dissipative process driven by gravitational radiation and tidal friction
- Tidal locking is an asymptotically stable state reached through dissipative evolution

- Planetary systems settle into metastable low-dissipation configurations through dissipative processes in protoplanetary disks
- Binary inspirals provide a natural setting for the framework's persistence functional

The framework's contribution is not a new mechanism of orbital evolution, but a unifying description of persistence across disparate dissipative systems using a common mathematical quantity: the persistence functional.

Keywords: attractor framework, astrophysics, gravitational radiation, tidal locking, orbital circularization, dissipative structures, Hamiltonian dynamics, planetary systems, binary inspirals, excess entropy production

1. Introduction

The attractor framework has been developed to describe persistence and change across physical, biological, cognitive, and social systems. The core claim is that every dissipative system maintains its attractor through continuous reconfiguration, and that reconfiguration generates excess entropy.

This paper extends the framework to astrophysical dissipative systems. The key insight is a distinction that is often blurred in the literature:

Concept	Role
Conservative gravitational dynamics	Defines the landscape of possible configurations (orbits, resonances, stable solutions)

Concept	Role
Dissipative processes	Select and can stabilize particular configurations within that landscape

Gravity does not provide attractors in the dynamical systems sense – Hamiltonian systems conserve phase-space volume and do not have attractors. However, when dissipative processes are added, the system evolves toward particular asymptotically stable configurations within the family of invariant solutions. The circular orbit is not a dynamical attractor of pure Newtonian gravity; it is the endpoint of dissipative evolution (tidal friction, gravitational radiation, gas drag).

This distinction is central to the paper. Gravity defines the landscape; dissipation determines which configuration is reached.

What is new: Existing astrophysical theory explains *how* dissipative mechanisms drive orbital evolution. The attractor framework proposes a common mathematical quantity – the persistence functional – that measures the cumulative irreversible cost of approaching an asymptotically stable configuration. The novelty is therefore not a new mechanism of orbital evolution, but a unifying description of persistence across disparate dissipative systems.

2. Conservative vs. Dissipative Systems

2.1 Hamiltonian Dynamics

A conservative Hamiltonian system preserves phase-space volume (Liouville's theorem). It does not have attractors in the dynamical systems sense. Orbits are determined by initial conditions and remain on their invariant tori (Arnold, 1989).

Property	Implication
No phase-space contraction	No attractors
Time-reversible	No arrow of time
Energy conserved	No dissipation

2.2 Dissipative Dynamics

When dissipative processes are added, the system loses energy and angular momentum. Phase-space volume contracts, and asymptotically stable states can emerge. For foundational treatments of irreversible thermodynamics, see Onsager (1931) and Prigogine (1947).

Property	Implication
Phase-space contraction	Asymptotically stable states appear
Time-irreversible	Arrow of time
Energy lost	Entropy generated

2.3 The Framework’s Position

The framework treats gravity as defining the landscape of possible configurations. Dissipation determines which of those configurations are actually reached.

Gravity defines the landscape. Dissipation selects the configuration.

This is the core insight of the paper.

3. The Gravitational Persistence

Functional

3.1 Excess Entropy Production

Following Galida (2026c), the excess entropy production rate is defined as:

$$\sigma_{\text{excess}}(x) = \sigma(x) - \sigma_{\text{ss}}(x)$$
$$\sigma_{\text{excess}}[\cdot](x) = \sigma[\cdot](x) - \sigma_{\text{ss}}[\cdot](x)$$

where $\sigma(x)$ is the total entropy production rate and $\sigma_{\text{ss}}(x)$ is the steady-state baseline rate at the attractor.

For gravitational systems, we propose:

$$\sigma_{\text{excess}} = E_{\text{irrev}} - E_{\text{ss}} / T_{\text{eff}}$$
$$\sigma_{\text{excess}}[\cdot] = T_{\text{eff}}[E_{\text{irrev}}[\cdot] - E_{\text{ss}}[\cdot]]$$

where E_{irrev} is the total irreversible energy loss rate, E_{ss} is the steady-state baseline loss rate at the attractor, and T_{eff} is an effective temperature.

This decomposition ensures $\sigma_{\text{excess}} \rightarrow 0$ at the attractor, avoiding the divergence problem that would arise from integrating raw dissipation rates over infinite time. Systems that continue to dissipate at a steady baseline (e.g., a circular binary emitting GWs, a tidally locked moon with residual eccentricity-driven heating) contribute only their excess above baseline to the persistence cost.

3.2 Domain-Specific Definitions

Process	Total E	Baseline E_{ss}	σ_{excess}
Orbital circularization	$L_{\text{GW}}(e)$	$L_{\text{GW}}(e=0)$	$[L_{\text{GW}}(e) - L_{\text{GW}}(0)] / T_{\text{eff}}$
Tidal locking	$P_{\text{tide}}(\Omega, e)$	$P_{\text{tide}}(\Omega=n, e)$	$[P_{\text{tide}}(\Omega, e) - P_{\text{tide}}(n, e)] / T_{\text{eff}}$
Disk dissipation	L_{disk}	$L_{\text{disk, steady}}$	$[L_{\text{disk}} - L_{\text{disk, ss}}] / T_{\text{eff}}$

3.3 The Persistence Functional

Definition 1 (Gravitational Persistence Functional): For a finite horizon $T > 0$:

$$D_T(x) = \int_0^T \sigma_{\text{excess}}(\phi_t(x)) dt$$
$$D_T[\cdot](x) = \int_0^T$$

$$\sigma_{\text{excess}}(\phi t(x)) dt$$

For trajectories that converge to the attractor: $D_{\infty}(x) = \int_0^{\infty} \sigma_{\text{excess}}(\phi t(x)) dt$ $D_{\infty}^{\square}(x) = \int_0^{\infty} \sigma_{\text{excess}}^{\square}(\phi t(x)) dt$

Interpretation: $D_{\infty}(x) D_{\infty}^{\square}(x)$ measures the total excess entropy generated during the approach to an asymptotically stable configuration – the cumulative cost of reconfiguration above the steady-state baseline.

Note on gravitational wave entropy: Classical gravitational waves are coherent radiation and do not automatically carry large thermodynamic entropy. The entropy associated with gravitational wave emission arises from coarse-graining the wave's phase space or from the generalized entropy increase of the sources (e.g., black hole horizons). The proposed definition $\sigma_{\text{excess}} = [L_{\text{GW}}(e) - L_{\text{GW}}(0)] / T_{\text{eff}}$ $\sigma_{\text{excess}}^{\square} = [L_{\text{GW}}^{\square}(e) - L_{\text{GW}}^{\square}(0)] / T_{\text{eff}}^{\square}$ isolates the eccentricity-specific excess above the circular-orbit baseline. Constructing an explicit entropy functional for gravitational radiation remains an open problem.

4. Orbital Circularization

4.1 The Phenomenon

Binary systems (stars, black holes, planets) often have elliptical orbits. Over time, these orbits tend to **circularize** – the eccentricity decreases and the orbit becomes more circular.

This is a dissipative process. The system loses energy and angular momentum through:

- **Gravitational radiation** (for compact objects)
- **Tidal friction** (for fluid bodies)
- **Gas drag** (for protoplanetary disks)

4.2 Framework Interpretation

Component	Role
The landscape	Family of Keplerian orbits (all ellipses)
Asymptotically stable state	Circular orbit (endpoint of dissipative evolution)
The dissipation	Gravitational radiation, tidal friction, gas drag
The cost	$\sigma_{\text{excess}} = [L_{\text{GW}}(e) - L_{\text{GW}}(0)] / T_{\text{eff}}$ $\sigma_{\text{excess}} = [L_{\text{GW}}(e) - L_{\text{GW}}(0)] / T_{\text{eff}}$

The framework proposes: $\kappa \propto 1/D \propto D^{-1}$

where κ is the circularization rate and $D = \int \sigma_{\text{excess}} dt$
 $D = \int \sigma_{\text{excess}} dt$ is the cumulative excess entropy production during circularization.

4.3 The Peters & Mathews Formula

The foundational computation of the gravitational-wave power from a Keplerian orbit was given by Peters & Mathews (1963). The secular decay of semi-major axis and eccentricity was derived by Peters (1964):

$$\frac{da}{dt} = -\frac{64}{5} \frac{G^3 m_1 m_2 (m_1 + m_2)}{c^5 a^3 (1 - e^2)^{7/2}} (1 + \frac{7324}{3} e^2 + \frac{3796}{5} e^4) \frac{dtd}{a} = -\frac{564}{5} \frac{G^3 m_1 m_2 (m_1 + m_2)}{c^5 a^3 (1 - e^2)^{7/2}} (1 + \frac{2473}{5} e^2 + \frac{9637}{5} e^4) \frac{dedt}{de} = -\frac{30415}{2} \frac{G^3 m_1 m_2 (m_1 + m_2)}{c^5 a^4 (1 - e^2)^{5/2}} e (1 + \frac{121304}{5} e^2) \frac{dtd}{de} = -\frac{15304}{5} \frac{G^3 m_1 m_2 (m_1 + m_2)}{c^5 a^4 (1 - e^2)^{5/2}} e (1 + \frac{304121}{5} e^2)$$

Framework Interpretation: The decay of eccentricity $e \rightarrow 0$ is the approach to the asymptotically stable state. The excess entropy production is the eccentricity-dependent component of the gravitational wave luminosity: $\sigma_{\text{excess}} = L_{\text{GW}}(e) - L_{\text{GW}}(0) / T_{\text{eff}}$

(0)

This quantity vanishes as $e \rightarrow 0$, consistent with the e -proportionality of the de/dt equation. Orbital eccentricity may serve as an experimentally accessible proxy for the cumulative excess entropy production.

5. Binary Inspirals

5.1 The Phenomenon

Binary systems of compact objects (neutron stars, black holes) lose energy through gravitational radiation. The orbit shrinks and the binary inspirals.

This is one of the most direct applications of the framework. The inspiral is a dissipative process driven by gravitational wave emission. For general relativistic treatments of binary dynamics and the geometry of spacetime, see Carroll (2004), Schutz (2009), Wald (1984), and Misner, Thorne & Wheeler (1973).

5.2 Framework Interpretation

Component	Role
The landscape	Family of binary orbits
Asymptotically stable state	Quasi-circular orbit (endpoint of circularization)
The dissipation	Gravitational radiation
The cost	$\sigma_{\text{excess}} = [L_{\text{GW}}(e) - L_{\text{GW}}(0)] / T_{\text{eff}}$ $\sigma_{\text{excess}} = [L_{\text{GW}}(e) - L_{\text{GW}}(0)] / T_{\text{eff}}$

5.3 The Persistence Functional

The persistence functional for a binary inspiral is:
$$D_{\infty} = \int_0^{\infty} \sigma_{\text{excess}}(t) \, dt = \int_0^{\infty} \text{LGW}(e(t)) - \text{LGW}(0) T_{\text{eff}} \, dt$$
$$\sigma_{\text{excess}}(t) dt = \int_0^{\infty} T_{\text{eff}} [\text{LGW}(e(t)) - \text{LGW}(0)] \, dt$$

Note on circularization: For compact-object binaries, eccentricity damps on a much shorter timescale than the inspiral itself. Gravitational radiation circularizes the orbit well before merger, so the system reaches a quasi-circular state as a near-asymptotic limit before the final coalescence.

Hypothesis: The inspiral time τ is inversely proportional to D_{∞} : $\kappa = 1/\tau \propto 1/D_{\infty} \Rightarrow \kappa = \tau^{-1} \propto D_{\infty}^{-1}$

6. Tidal Locking

6.1 The Phenomenon

Tidal locking occurs when a body’s rotational period equals its orbital period. The Moon is tidally locked to Earth. Many exoplanets in the habitable zone are expected to be tidally locked.

Tidal locking is a dissipative process. Tidal friction converts rotational energy into heat, gradually slowing the body’s rotation until it matches its orbital period.

6.2 Framework Interpretation

Component	Role
The landscape	Family of rotational states
Asymptotically stable state	Tidal lock (rotational period = orbital period)

Component	Role
The dissipation	Tidal friction (heat generation)
The cost	$\sigma_{\text{excess}} = [P_{\text{tide}}(\Omega, e) - P_{\text{tide}}(\Omega = n, e)] / T_{\text{eff}}$ $\sigma_{\text{excess}} = [P_{\text{tide}}(\Omega, e) - P_{\text{tide}}(\Omega = n, e)] / T_{\text{eff}}$

Hypothesis: The tidally locked state is a low-dissipation configuration for the system. Once locked, tidal dissipation approaches a minimum. The excess entropy production is the despinning-specific component above whatever baseline eccentricity-driven heating persists after lock.

6.3 The Tidal Locking Timescale

The timescale for tidal locking is commonly given as (see, e.g., Murray & Dermott, 1999):
 $\tau_{\text{lock}} \approx 221 Q k_2^2 m M (a/R)^6 1/\Omega$
 $\tau_{\text{lock}} \approx 212 k_2^2 Q^2 M m (R/a)^6 1/\Omega$

where:

- Q is the tidal dissipation factor
- k_2^2 is the Love number
- m is the mass of the body
- M is the mass of the primary
- a is the semi-major axis
- R is the radius of the body
- Ω is the rotation rate

(Different derivations use different prefactors depending on the assumed dissipation model; the $(a/R)^6$ scaling is robust.)

Hypothesis: $\kappa = 1/\tau_{\text{lock}}$. The recovery rate is the inverse of the locking timescale. The cumulative excess entropy production is the total tidal heat dissipated during despinning above the post-lock baseline.

7. Planetary Systems

7.1 Formation and Evolution

Planetary systems form from protoplanetary disks. The disk is a dissipative structure: it loses energy through radiation, viscosity, and accretion.

Over time, the system approaches a stable configuration:

- Planets on nearly circular orbits
- Resonances between orbits
- Stable spin-orbit states

For a comprehensive treatment of solar system dynamics and tidal evolution, see Murray & Dermott (1999).

7.2 Framework Interpretation

Component	Role
The landscape	Family of possible planetary configurations
Metastable configuration	Low-dissipation planetary system
The dissipation	Disk viscosity, radiation, accretion
The cost	$\sigma_{\text{excess}} = [L_{\text{disk}} - L_{\text{disk, ss}}] / T_{\text{disk}}$ $\sigma_{\text{excess}} = [L_{\text{disk}} - L_{\text{disk, ss}}] / T_{\text{disk}}$

Hypothesis: Mature planetary systems approach metastable low-dissipation configurations. The cumulative excess entropy production is the total disk dissipation above the steady-state baseline integrated over the formation epoch.

8. Entropy Generation in Gravitational Systems

8.1 The Subtlety of Gravitational Entropy

Gravitational waves carry energy. Whether they carry entropy is a more subtle question. Classical gravitational waves are coherent radiation; coherent radiation is not obviously high-entropy. Binary mergers ultimately increase the generalized entropy of spacetime, but the bookkeeping is subtle.

Note: Throughout this paper, entropy generation refers to the irreversible processes associated with tidal heating, viscous dissipation, and the generalized entropy increase accompanying gravitational-wave emission. The precise entropy carried by gravitational radiation remains an active topic.

8.2 Operational Definition of σ_{excess}

For the purposes of this framework, we propose the following operational definition: $\sigma_{\text{excess}} = \frac{\dot{E}_{\text{irrev}} - \dot{E}_{\text{ss}}}{T_{\text{eff}}}$

where:

- $\dot{E}_{\text{irrev}}$ is the total irreversible energy loss rate
- $\dot{E}_{\text{ss}}$ is the steady-state baseline loss rate at the attractor
- T_{eff} is an effective temperature for the dissipative process

This definition ensures $\sigma_{\text{excess}} \geq 0$ and vanishes when the system reaches its attractor. For specific astrophysical contexts:

Context	E^{irrev}	E^{ss}	T_{eff}
Orbital circularization	$L_{\text{GW}}(e)$	$L_{\text{GW}}(0)$	Effective GW temperature
Tidal locking	$P_{\text{tide}}(\Omega, e)$	$P_{\text{tide}}(\Omega=n, e)$	Effective body temperature
Disk dissipation	L_{disk}	$L_{\text{disk, ss}}$	Disk temperature
Black hole mergers	L_{GW}	0	Hawking temperature of final black hole

Note: This is a working hypothesis. Constructing an explicit entropy functional for relativistic gravitational systems remains an open problem. The effective temperature T_{eff} is the primary underdetermined quantity in the framework; its derivation from first principles is a priority for future work.

9. The Boundary

The framework’s boundary is not absolute zero. It is the **absence of irreversible processes**. At the boundary, the system becomes conservative and no entropy is generated. Hamiltonian systems exist at nonzero temperature; the boundary is dynamical, not thermal.

10. Testable Predictions

10.1 Core Prediction

Prediction: The circularization rate κ is inversely proportional to the cumulative excess entropy production during circularization. $\kappa \propto 1/D \propto \kappa D \propto 1$

10.2 Specific Predictions

Prediction	Falsification
Tidal locking timescale correlates with total tidal heat dissipated above baseline	If no correlation, the prediction is falsified
Circularization rate correlates with total eccentricity-dependent GW energy emitted	If no correlation, the prediction is falsified
Planetary system stability correlates with total disk dissipation above steady state	If no correlation, the prediction is falsified

11. Open Questions

Question	Status
Q1: Gravitational entropy	What is the entropy of a gravitational system? (Penrose, 1965; Hawking & Ellis, 1973)
Q2: Black hole entropy	How does black hole entropy fit into the framework?

Question	Status
Q3: Entropy of gravitational radiation	Does gravitational radiation carry entropy, and if so, how is it defined? (Zeldovich, 1972)
Q4: Cosmological stability	Do cosmological models admit asymptotically stable late-time solutions?
Q5: Effective temperature for GWs	What is the correct T_{eff} for gravitational wave entropy production? (Galida, 2026d)
Q6: Coarse-graining	What coarse-graining scheme defines the entropy of classical gravitational waves? (Galida, 2026d)

12. Conclusion

The attractor framework extends naturally to astrophysical dissipative systems. The key insight is a distinction that is often blurred:

Gravity defines the landscape. Dissipation selects the configuration.

Conservative gravitational dynamics define families of stable invariant solutions. Dissipative processes – gravitational radiation, tidal friction, gas drag – select and can stabilize particular configurations within those families.

The framework does not claim that gravity provides attractors. It claims that the combination of conservative dynamics and dissipative processes produces asymptotically stable states. This is a more accurate and defensible position.

The contribution is not a new mechanism of orbital evolution,

but a unifying description of persistence across disparate dissipative systems using a common mathematical quantity: the persistence functional $D_\infty = \int \sigma_{\text{excess}} dt$, with σ_{excess} operationally defined as the rate of irreversible energy loss above steady-state baseline divided by an effective temperature.

References

- Arnold, V. I. (1989). *Mathematical Methods of Classical Mechanics*. Springer.
- Carroll, S. M. (2004). *Spacetime and Geometry: An Introduction to General Relativity*. Addison-Wesley.
- Galida, R. (2026a). "The Persistence Functional: A Candidate Formal Foundation for the Attractor Framework." *Fantasy Attractor*.
- Galida, R. (2026b). "Deriving Corrective Permeability from the Cumulative Deviation Functional." *Fantasy Attractor*.
- Galida, R. (2026c). "Excess Entropy Production as a Candidate Universal Cost of Persistence: A Thermodynamic Foundation for the Attractor Framework." *Fantasy Attractor*.
- Galida, R. (2026d). "Deep Research Questions on the Attractor Framework." *Fantasy Attractor*.
- Goldreich, P., & Soter, S. (1966). "Q in the Solar System." *Icarus*, 5(1-6), 375-389.
- Hawking, S. W., & Ellis, G. F. R. (1973). *The Large Scale Structure of Space-Time*. Cambridge University Press.
- Misner, C. W., Thorne, K. S., & Wheeler, J. A. (1973). *Gravitation*. W. H. Freeman.

Murray, C. D., & Dermott, S. F. (1999). *Solar System Dynamics*. Cambridge University Press.

Onsager, L. (1931). "Reciprocal Relations in Irreversible Processes." *Physical Review*, 37(4), 405-426.

Penrose, R. (1965). "Gravitational Collapse and Space-Time Singularities." *Physical Review Letters*, 14(3), 57-59.

Peters, P. C. (1964). "Gravitational Radiation and the Motion of Two Point Masses." *Physical Review*, 136(4B), B1224-B1232.

Peters, P. C., & Mathews, J. (1963). "Gravitational Radiation from Point Masses in a Keplerian Orbit." *Physical Review*, 131(1), 435-440.

Prigogine, I. (1947). *Étude Thermodynamique des Phénomènes Irréversibles*. Dunod.

Schutz, B. F. (2009). *A First Course in General Relativity* (2nd ed.). Cambridge University Press.

Wald, R. M. (1984). *General Relativity*. University of Chicago Press.

Zeldovich, Y. B. (1972). "A Hypothesis Unifying the Structure and the Entropy of the Universe." *Monthly Notices of the Royal Astronomical Society*, 160(1), 1P-4P.

Suggested citation: Galida, R. S. (2026). The Attractor Framework in Astrophysics: Persistence, Entropy, and Gravitational Systems (Final Edition). *Fantasy Attractor*.

Excess Entropy Production as a Candidate Universal Cost of Persistence: A Thermodynamic Foundation for the Attractor Framework; Robert Galida (July 2026) [F]

Abstract

Every dissipative system maintains its attractor through continuous reconfiguration. Reconfiguration requires work; work generates entropy. The recovery rate κ – corrective permeability – is the rate at which a system reconfigures to return to its attractor after perturbation. This paper proposes that κ is a measure of excess entropy generation rate.

We develop an abstract persistence cost framework and prove its equivalence to Lyapunov theory. We then identify entropy production as a physical realization of this cost, deriving: $\kappa = \inf_{\delta} \int_0^\infty \sigma_{\text{excess}}(\phi_t(x)) dt$ $\kappa = \inf_{\delta} \int_0^\infty \sigma_{\text{excess}}(\phi_{t-\delta}(x)) dt$

where $\sigma_{\text{excess}} = \sigma - \sigma_{ss}$ $\sigma_{\text{excess}} = \sigma - \sigma_{ss}$ is the excess entropy production rate above the system's steady-state baseline. For physical systems, the baseline is zero (equilibrium); for biological, cognitive, and social systems, the baseline is the steady-state dissipation rate of the healthy, well-coordinated attractor.

This unifies physical, biological, cognitive, and social systems. The framework is grounded in the second law of

thermodynamics and non-equilibrium steady-state thermodynamics, not analogy. Empirical predictions are provided for each domain.

Keywords: entropy generation, excess entropy production, corrective permeability, attractor framework, dissipative structures, reconfiguration, Lyapunov theory, free energy principle, allostatic load

1. Introduction

The attractor framework defines persistence as the ability of a system to maintain its attractor under perturbation. Historically, persistence has been measured kinematically – as distance traveled or time spent away from equilibrium. This paper proposes that the true cost of persistence is thermodynamic: it is the excess entropy generated during reconfiguration and recovery.

Every dissipative system maintains its attractor through continuous reconfiguration. A bacterium reconfigures its metabolism to maintain homeostasis. A brain reconfigures its synaptic connections to maintain predictive models. A society reconfigures its institutions to maintain order. Reconfiguration requires work; work generates entropy. The second law of thermodynamics applies at every level of organization.

We develop an abstract persistence cost framework first, establishing its equivalence to Lyapunov theory. We then identify entropy production as a physical realization of this cost, deriving the relationship between corrective permeability and excess entropy generation.

The framework unifies physical, biological, cognitive, and

social systems. It is grounded in the second law of thermodynamics and non-equilibrium steady-state thermodynamics, not analogy.

2. The Persistence Cost Functional

Let X be a state space, $\phi_t(x)$ the flow of a dynamical system, and $A \subseteq X$ an attractor set. Let $\delta(x) = d(x, A)$ be the distance from x to the attractor. For a treatment of state-space constraints in viability theory, see Aubin (1991).

Definition 1 (Persistence Cost Functional): A persistence cost functional $C(x)$ is a scalar function on X satisfying:

1. $C(x) \geq 0$ for all x
2. $C(x) = 0$ if and only if $x \in A$
3. $C(\phi_t(x)) \in L^1([0, \infty))$ for all x in the basin

Definition 2 (Cumulative Persistence Cost): For a finite horizon $T > 0$: $DT(x) = \int_0^T C(\phi_t(x)) dt$

For trajectories that converge to the attractor: $D^\infty(x) = \int_0^\infty C(\phi_t(x)) dt$

3. Existence and Lyapunov Equivalence

Theorem 1 (Existence of the Persistence Functional): Assume $C(x) \geq 0$, $C=0$ only on A ,

and $C(\phi_t(x)) \in L^1([0, \infty))$ $C(\phi_{t^*}(x)) \in L^1([0, \infty))$ for all x in the basin. Assume f is locally Lipschitz, the flow is continuously differentiable in the initial condition, and C is continuous and locally bounded. Then:

1. $D^\infty(x) = \int_0^\infty C(\phi_t(x)) dt$ $D^\infty_*(x) = \int_0^\infty C(\phi_{t^*}(x)) dt$ exists and is finite.
2. $D^\infty D^\infty_*$ is continuous.
3. $D^\infty D^\infty_*$ satisfies the transport equation:

$$\nabla D^\infty(x) \cdot f(x) = -C(x) \quad \nabla D^\infty_*(x) \cdot f(x) = -C(x)$$

Proof: The integral exists and is finite by the $L^1 L^1$ assumption. Continuity follows from the dominated convergence theorem under the stated regularity assumptions. To derive the transport equation, compute: $D(\phi_h(x)) = \int_0^\infty C(\phi_t(x)) dt = D(x) - \int_0^h C(\phi_t(x)) dt$ $D(\phi_{h^*}(x)) = \int_0^\infty C(\phi_{t^*}(x)) dt = D(x) - \int_0^h C(\phi_{t^*}(x)) dt$

$$\text{Then: } D(\phi_h(x)) - D(x) = -\int_0^h C(\phi_t(x)) dt \rightarrow -C(x)h \quad D(\phi_{h^*}(x)) - D(x) = -\int_0^h C(\phi_{t^*}(x)) dt \rightarrow -C(x)$$

as $h \rightarrow 0$ $h \rightarrow 0$. By the chain rule: $\nabla D(x) \cdot f(x) = -C(x)$ $\nabla D_*(x) \cdot f(x) = -C(x)$ $\square$

Corollary (Equivalence to Lyapunov Theory): Any Lyapunov function $V(x)$ (with $V \geq 0$, $V=0$ on the attractor, and $V' \leq 0$) yields a persistence cost $C(x) = -V'(x)$. Conversely, any persistence cost $C(x)$ satisfying $\nabla D \cdot f = -C$ defines a Lyapunov function $D(x)$.

Proof: If V is a Lyapunov function, then $V' = \nabla V \cdot f \leq 0$. Define $C = -V'$. Then $C \geq 0$, $C=0$ on the attractor, and $D = \int C = V(x) - V(\phi_T(x))$. Conversely, if $\nabla D \cdot f = -C$, then $D' = -C \leq 0$, so D is a Lyapunov function. $\square$

Interpretation: The persistence cost framework is mathematically equivalent to classical Lyapunov stability theory. For the connection to contraction analysis, see Lohmiller & Slotine (1998). For control Lyapunov functions, see Freeman & Kokotovic (1996). Entropy production is one physically meaningful realization of the cost function CC . For a detailed treatment of Lipschitz continuity of $D^\infty D^\infty$ under a Lipschitz-flow hypothesis, see Galida (2026a), Proposition 4.

4. Entropy Production as Persistence Cost

4.1 Entropy Balance

For an open system, the entropy balance equation is: $\frac{dS_{\text{system}}}{dt} = \sigma - \Phi$

where $\sigma \geq 0$ is the entropy production rate (always non-negative by the second law) and Φ is the entropy export rate to the environment. For foundational treatments of stochastic thermodynamics and entropy production, see Seifert (2012) and Sekimoto (2010).

For a system in a steady state: $\frac{dS_{\text{system}}}{dt} = 0 \implies \sigma = \Phi$

4.2 Excess Entropy Production

Define the **steady-state entropy production rate** σ_{ss} as the rate when the system is at its attractor.

Define the **excess entropy production rate**: $\sigma_{\text{excess}}(x) = \sigma(x) - \sigma_{ss}$

Assumption (Excess Entropy Decay): For all trajectories in the

basin, there exist constants $C < \infty$ and $\mu > 0$ such that:

$$\sigma_{\text{excess}}(\phi^t(x)) \leq C e^{-\mu t} \sigma_{\text{excess}}(x)$$

$$\sigma_{\text{excess}}(\phi^t(x)) \leq C e^{-\mu t} \sigma_{\text{excess}}(x)$$

for all $t \geq 0$. This ensures $D^\infty(x) < \infty$ and is the standard hypothesis under which the persistence functional and its associated bounds are well-defined, consistent with Galida (2026a, 2026b). The decay rate μ may be domain-specific and is empirically measurable.

Note on generalization: The exponential decay assumption is adopted here to ensure finiteness of D^∞ and to maintain consistency with the prior papers in this series. Generalization to L^1 integrable decays (e.g., algebraic) is a priority for future work.

4.3 The Entropy Persistence Functional

Definition 3 (Cumulative Excess Entropy Functional): For a finite horizon $T > 0$:

$$D^T(x) = \int_0^T \sigma_{\text{excess}}(\phi^t(x)) dt$$

$$D^T(x) = \int_0^T \sigma_{\text{excess}}(\phi^t(x)) dt$$

For trajectories that converge to the attractor:

$$D^\infty(x) = \int_0^\infty \sigma_{\text{excess}}(\phi^t(x)) dt$$

$$D^\infty(x) = \int_0^\infty \sigma_{\text{excess}}(\phi^t(x)) dt$$

Interpretation: The persistence functional is the total excess entropy generated during reconfiguration and recovery.

4.4 Corrective Permeability

Definition 4 (Corrective Permeability):

$$\kappa = \inf_{x \in B} \delta(x) D^\infty(x)$$

$$\kappa = \inf_{x \in B} \delta(x) D^\infty(x)$$

where $\delta(x) = d(x, A)$ is the distance to the attractor.

Interpretation: κ is the **minimum excess entropy cost per unit distance**. It measures the efficiency of reconfiguration: a system that returns with minimal excess entropy generation has

high κ ; a system that generates excess entropy has low κ .

4.5 Basin Depth

Proposition 1 (Properties of Basin Depth): Define $B = D^\infty(\text{saddle}) - D^\infty(A)$, where saddle is the lowest point on the basin boundary (the separatrix between attractors). For the connection to large-deviation theory and escape rates, see Freidlin & Wentzell (2012). Then:

1. $B \geq 0$, with equality iff the basin has no barrier (i.e., the boundary coincides with the attractor).
2. For gradient systems $\dot{x} = -\nabla V(x)$, $B = V(\text{saddle}) - V(A)$ (the classical energy barrier).
3. B is invariant under smooth coordinate changes (coordinate invariance).
4. B depends on the chosen persistence cost functional C ; different costs yield different barriers.

Proof: (1) follows from non-negativity of D^∞ . (2) follows from the transport equation $\nabla \cdot f = -C$ and the identity $f = -\nabla V$. (3) follows from the invariance of the integral under diffeomorphisms. (4) is self-evident.

5. Domain-Specific Realizations

5.1 Physical Systems: Thermodynamic

Excess Entropy

For a thermodynamic system, $S(x) = k_B \log \Omega(x)$, where $\Omega(x)$ is the number of microstates. For an isolated system, $\sigma_{ss} = 0$ (equilibrium), so $\sigma_{excess} = S - \sigma_{ss} = S$. $\kappa = \inf_x \delta(x) S(A) - S(x) = \inf_x S(A) - S(x) \delta(x)$

Example: A gas returning to equilibrium after compression. The entropy generated is $\Delta S = nR \log(V_f/V_i)$.

5.2 Biological Systems: Metabolic Excess Entropy

For a biological system, $S(x)$ is the metabolic entropy. The baseline σ_{ss} is the resting metabolic rate (homeostasis). The excess is: $\sigma_{excess} = \text{metabolic rate} - \text{resting metabolic rate}$. $\kappa = \inf_x \delta(x) \int_0^\infty \sigma_{excess}(\phi_t(x)) dt = \inf_x \int_0^\infty \sigma_{excess}(\phi_t(x)) dt \delta(x)$

Example: A cell returning to homeostasis after a nutrient shock. The excess entropy generated is the metabolic cost of restoring homeostasis above baseline. For the dissipative-structures framework underlying biological self-organization, see Nicolis & Prigogine (1989).

5.3 Cognitive Systems: Free Energy Dissipation

For a cognitive system, variational free energy $F = -\log p(y|x) + \text{DKL}[q(\cdot) \| p(\cdot|x)]$ is adopted here as one candidate persistence functional. We do not claim variational free energy is uniquely correct; it is adopted as the most developed existing candidate persistence functional for cognitive systems. Other candidates (Bayesian surprise, expected free energy, predictive information) are possible; this paper focuses

on FF due to its established role in the free-energy principle (Friston, 2010). For the thermodynamics of information and its connection to free-energy minimization, see Parrondo, Horowitz & Sagawa (2015) and Sagawa & Ueda (2008).

The baseline σ_{ss} is the baseline neural dissipation rate (resting brain activity). The excess is: $\sigma_{\text{excess}} = F' - F'_{ss}$
 $\kappa = \inf_x \delta(x) \int_0^\infty \sigma_{\text{excess}}(\phi_t(x)) dt$
 $\kappa = x \inf_x \int_0^\infty \sigma_{\text{excess}}(\phi_t(x)) dt \delta(x)$

Example: A cognitive system updating its beliefs after a prediction error. The excess entropy generated is the free energy dissipated during belief updating above baseline.

5.4 Social Systems: Coordination Excess Entropy

For a social system, define the aggregate social entropy production rate as: $\sigma_{\text{social}}(t) = \sum_i (S'_i(t) - S'_{i\text{rest}})$
 $\sigma_{\text{social}}(t) = \sum_i (S'_i(t) - S'_{i\text{rest}})$

where $S'_i(t)$ is the total entropy production rate of individual i , and $S'_{i\text{rest}}$ is the individual's baseline entropy production rate in a resting, minimally socially constrained state. This is measured via physiological proxies such as basal metabolic rate, resting allostatic load, or cortisol baseline (McEwen, 1998; Sterling & Eyer, 1988).

Interpretation: σ_{social} measures the excess dissipation attributable to social constraints: the additional entropy generated by coordination, communication, conflict, norm enforcement, and institutional friction.

Non-Negativity: Unlike total entropy production $S'_i \geq 0$ (which follows from the second law), σ_{social} is not guaranteed to be non-negative. Division of labor, infrastructure, and specialization may *reduce* an individual's

metabolic burden relative to a solitary baseline. The hypothesis is that during recovery from social disruption, $\sigma_{\text{social}} \geq 0 \sigma_{\text{social}} \rightarrow 0$; in steady-state, $\sigma_{\text{social}} \rightarrow 0 \sigma_{\text{social}} \rightarrow 0$. This is an empirical claim, not a theorem.

The baseline $\sigma_{ss} \sigma_{ss}$ is the steady-state social entropy production rate (well-coordinated society). The excess is: $\sigma_{\text{excess}} = \sigma_{\text{social}} - \sigma_{ss} \sigma_{\text{excess}} = \sigma_{\text{social}} - \sigma_{ss}$
 $\kappa = \inf_{\delta} \int_0^{\infty} \sigma_{\text{excess}}(\phi t(x)) dt \kappa = x \inf_{\delta} \int_0^{\infty} \sigma_{\text{excess}}(\phi t(x)) dt \delta(x)$

Example: A society recovering from a shock (economic crisis, political upheaval). The excess entropy generated is the coordination cost of restructuring above baseline. A harmonious society has $\sigma_{\text{excess}} = 0 \sigma_{\text{excess}} = 0$; a turbulent society has $\sigma_{\text{excess}} > 0 \sigma_{\text{excess}} > 0$; a chronically turbulent society may have settled into a new attractor with a higher $\sigma_{ss} \sigma_{ss}$. This illustrates the framework’s central distinction: the attractor is the state of minimum entropy generation for that class of system.

6. The Unified Framework

6.1 Summary Table

Domain	Entropy Functional	Baseline $\sigma_{ss} \sigma_{ss}$	Excess $\sigma_{\text{excess}} \sigma_{\text{excess}}$	Recovery Rate $\kappa \kappa$
Physical	Thermodynamic entropy	0 (equilibrium)	$S - S^*$	$\inf_{\delta} \delta \Delta S \inf_{\delta} \Delta S \delta$
Biological	Metabolic entropy	Resting metabolic rate	Metabolic rate – resting	$\inf_{\delta} \delta \int \sigma_{\text{excess}} dt \inf_{\delta} \int \sigma_{\text{excess}} dt \delta$
Cognitive	Free energy	Baseline neural dissipation	$F - F_{ss} F - F_{ss}$	$\inf_{\delta} \delta \int \sigma_{\text{excess}} dt \inf_{\delta} \int \sigma_{\text{excess}} dt \delta$
Social	Social entropy production	Steady-state social dissipation	$\sigma_{\text{social}} - \sigma_{ss} \sigma_{\text{social}} - \sigma_{ss}$	$\inf_{\delta} \delta \int \sigma_{\text{excess}} dt \inf_{\delta} \int \sigma_{\text{excess}} dt \delta$

6.2 The Universal Structure

Every domain follows the same mathematical structure:

Component	Expression
Excess entropy production	$\sigma_{\text{excess}}(x) = \sigma(x) - \sigma_{ss}$ $\sigma_{\text{excess}}[\cdot](x) = \sigma(\cdot)(x) - \sigma_{ss}[\cdot]$
Cumulative cost	$D^\infty(x) = \int_0^\infty \sigma_{\text{excess}}(\phi_t(x)) \, dt$ $D^\infty[\cdot](x) = \int_0^\infty \sigma_{\text{excess}}[\cdot](\phi_t[\cdot](x)) \, dt$
Recovery rate	$\kappa = \inf_{\delta} \delta(x) / D^\infty(x)$ $\kappa = \inf_{\delta} \delta[\cdot](x) / D^\infty[\cdot](x)$
Basin depth	$B = D^\infty(\text{saddle})$ $B = D^\infty[\cdot](\text{saddle})$
Transport equation	$\nabla D \cdot f = -\sigma_{\text{excess}}$ $\nabla D \cdot f = -\sigma_{\text{excess}}[\cdot]$

6.3 The Low-Energy Attractor Benchmark (Proposed Hypothesis)

We propose the following benchmark as an additional hypothesis: **the attractor is the state of minimum entropy generation for that class of system.**

Domain	Attractor	Entropy Generation at Attractor
Physical	Equilibrium	$\sigma = 0$ $\sigma = 0$
Biological	Homeostasis	$\sigma = \sigma_{ss} > 0$ $\sigma = \sigma_{ss}[\cdot] > 0$ (resting metabolism)
Cognitive	Settled Belief	$\sigma = \sigma_{ss} > 0$ $\sigma = \sigma_{ss}[\cdot] > 0$ (baseline neural dissipation)
Social	Coordinated Order	$\sigma = \sigma_{ss} > 0$ $\sigma = \sigma_{ss}[\cdot] > 0$ (baseline institutional friction)

Interpretation:

1. **For equilibrium systems** (gases, isolated systems), the attractor is the state where entropy generation reaches zero – the system has nowhere lower to go.

2. **For dissipative systems** (cells, brains, societies), the attractor is the state where entropy generation reaches its lowest *non-zero* steady-state value – the minimum entropy generation the system can sustain while maintaining its functional organization.

Important caveats:

- This is a proposed benchmark, not a derived theorem.
- For cognitive systems in particular, minimizing *entropy production rate* (a thermodynamic quantity) and minimizing *free energy/surprise* (the actual claim in the free-energy principle) are distinct minimization principles. The framework does not establish a bridge between them; this is an open question.
- The benchmark is an empirical hypothesis that requires domain-specific validation.

In all cases, the attractor is the **lowest entropy-generating state *that system can have while remaining itself***.

7. Testable Predictions

7.1 Core Prediction

Prediction: The recovery rate κ is inversely proportional to the excess entropy generated during reconfiguration: $\kappa \propto 1/D \propto \kappa \propto D^{-1}$

Falsification: If a system returns to its attractor with high excess entropy generation but high recovery rate, the prediction is falsified.

7.2 Secondary Prediction

Prediction: Systems that maintain their attractor with minimal excess entropy generation are more “efficient.” Systems that generate excess entropy are “inefficient” or “stressed.”

Falsification: If an inefficient system has lower excess entropy generation than an efficient system, the prediction is falsified.

7.3 Domain-Specific Predictions

Domain	Prediction	Falsification
Physical	$\kappa\kappa$ correlates with thermal efficiency	$\kappa\kappa$ high but efficiency low
Biological	$\kappa\kappa$ correlates with metabolic efficiency	$\kappa\kappa$ high but metabolic cost high
Cognitive	$\kappa\kappa$ correlates with learning efficiency	$\kappa\kappa$ high but learning cost high
Social	$\kappa\kappa$ correlates with institutional efficiency	$\kappa\kappa$ high but coordination cost high

8. Experimental Design

8.1 Physical Systems

- **System:** Gas in a piston
- **Perturbation:** Compression
- **Measurement:** Excess entropy generation (heat measurement) and recovery time
- **Test:** Correlation between $\kappa\kappa$ and $1/D_{\infty}1/D_{\infty}$

8.2 Biological Systems

- **System:** Cell culture
- **Perturbation:** Nutrient shock
- **Measurement:** Metabolic rate above resting (oxygen consumption) and recovery time
- **Test:** Correlation between $\kappa\kappa$ and metabolic cost

8.3 Cognitive Systems

- **System:** Human participants in a learning task
- **Perturbation:** Prediction error
- **Measurement:** Free energy dissipation above baseline (EEG complexity, pupil dilation) and belief updating rate
- **Test:** Correlation between $\kappa\kappa$ and free energy dissipation

8.4 Social Systems

- **System:** Institutional response to shocks
 - **Perturbation:** Economic or political crisis
 - **Measurement:** Social entropy production above baseline (allostatic load, cortisol, institutional friction) and recovery time
 - **Test:** Correlation between $\kappa\kappa$ and social entropy production
-

9. Open Questions

Question	Status	Difficulty
Q1: Uniqueness of $S(x)$	Are there multiple valid entropy functionals for a given domain?	Hard
Q2: Variational principle	Is there a universal variational principle that yields $S(x)$?	Hard
Q3: Social second law	Does $\sigma_{\text{social}} \geq 0$ always hold during recovery?	Very Hard
Q4: Cross- level entropy	How does entropy generation at one level relate to entropy generation at another?	Hard
Q5: Measurement	Can we measure excess entropy generation in cognitive and social systems directly?	Moderate
Q6: Unification	Can all domain-specific entropy functionals be derived from a single universal functional?	Very Hard

10. Conclusion

Every dissipative system maintains its attractor through continuous reconfiguration. Reconfiguration requires work; work generates excess entropy. The recovery rate κ – corrective permeability – is the rate at which a system reconfigures to return to its attractor after perturbation. We have proposed that κ is a measure of excess entropy generation rate.

We developed an abstract persistence cost framework and proved its equivalence to Lyapunov theory. We then identified entropy production as a physical realization of this cost, deriving:

$$\kappa = \inf_{\delta} \int_0^\infty \sigma_{\text{excess}}(\phi_t(x)) dt \quad \kappa = \inf_{\delta} \int_0^\infty \sigma_{\text{excess}}$$

$$(\phi t(x))dt\delta(x)$$

where $\sigma_{\text{excess}} = \sigma - \sigma_{\text{ss}}$ is the excess entropy production rate above the system's steady-state baseline – thermodynamic entropy for physical systems, metabolic entropy for biological systems, free energy dissipation for cognitive systems, and social entropy production for social systems.

We proposed a unified benchmark: the attractor is the state of minimum entropy generation for that class of system – zero for equilibrium systems, non-zero steady-state for dissipative systems. This provides a unified criterion for identifying attractors across domains: an attractor is a state from which the system cannot reduce its entropy generation further without losing its defining structure or function.

This unifies physical, biological, cognitive, and social systems. In each domain, persistence requires reconfiguration; reconfiguration generates excess entropy; κ measures the entropy cost of that reconfiguration. The framework is grounded in the second law of thermodynamics and non-equilibrium steady-state thermodynamics, not analogy.

Social Application: The framework provides a thermodynamic interpretation of social dynamics: harmony is a low-entropy attractor state; turbulence is a high-entropy state generated by excess dissipation during reconfiguration. The recovery rate κ measures how efficiently a society transitions from turbulence back to harmony – that is, how quickly it reduces its excess entropy production to zero.

11. Limitations

This paper establishes an abstract persistence cost framework with a proposed thermodynamic realization. Several limitations

should be explicitly acknowledged:

1. **Uniqueness.** Entropy production is not proved to be the unique persistence cost. Many positive functionals $C(x)$ satisfy $\nabla D \cdot f = -C \nabla D \cdot f = -C$. The identification of entropy production as the canonical cost is a physically motivated hypothesis, not a mathematical theorem.
 2. **Scope.** The framework does not imply that all domains obey thermodynamics literally. The cognitive and social realizations are proposed hypotheses requiring empirical validation.
 3. **Decay assumption.** Exponential decay of ϕ_{excess} is a sufficient assumption to ensure finiteness of $D_{\infty}[\phi]$, not a necessary one. Generalization to L^1 integrable decays (e.g., algebraic) is a priority for future work.
 4. **Basin depth.** Basin depth $B = D_{\infty}(\text{saddle}) = D_{\infty}[\phi](\text{saddle})$ is defined in terms of the persistence cost functional. Its relationship to classical energy barriers is established only for gradient systems.
 5. **Empirical validation.** The predictions of the framework – particularly the inverse relationship between κ and $D_{\infty}[\phi]$ – remain to be tested empirically across domains.
 6. **Low-energy attractor benchmark.** The benchmark proposed in §6.3 is a hypothesis, not a derived theorem. For cognitive systems, it risks conflating thermodynamic entropy production with free-energy minimization – distinct principles whose relationship remains open.
-

References

- Aubin, J. P. (1991). *Viability Theory*. Birkhäuser.
- Boltzmann, L. (1877). "Über die Beziehung zwischen dem zweiten Hauptsatz der mechanischen Wärmetheorie und der Wahrscheinlichkeitsrechnung." *Wiener Berichte*, 76, 373-435.
- Clausius, R. (1865). "Über verschiedene für die Anwendung bequeme Formen der Hauptgleichungen der mechanischen Wärmetheorie." *Annalen der Physik*, 125(7), 353-400.
- Freeman, R. A., & Kokotovic, P. V. (1996). *Robust Nonlinear Control Design: State-Space and Lyapunov Techniques*. Birkhäuser.
- Freidlin, M. I., & Wentzell, A. D. (2012). *Random Perturbations of Dynamical Systems* (3rd ed.). Springer.
- Friston, K. (2010). "The free-energy principle: a unified brain theory?" *Nature Reviews Neuroscience*, 11(2), 127-138.
- Galida, R. (2026a). "The Persistence Functional: A Candidate Formal Foundation for the Attractor Framework." *Fantasy Attractor*.
- Galida, R. (2026b). "Deriving Corrective Permeability from the Cumulative Deviation Functional." *Fantasy Attractor*.
- Jaynes, E. T. (1957). "Information Theory and Statistical Mechanics." *Physical Review*, 106(4), 620-630.
- Khalil, H. K. (2002). *Nonlinear Systems* (3rd ed.). Prentice Hall.
- Kondepudi, D., & Prigogine, I. (1998). *Modern Thermodynamics: From Heat Engines to Dissipative Structures*. Wiley.
- Lohmiller, W., & Slotine, J. J. E. (1998). "On contraction analysis for non-linear systems." *Automatica*, 34(6), 683-696.

Lyapunov, A. M. (1892). *The General Problem of the Stability of Motion*.

McEwen, B. S. (1998). "Stress, Adaptation, and Disease: Allostasis and Allostatic Load." *Annals of the New York Academy of Sciences*, 840(1), 33-44.

Nicolis, G., & Prigogine, I. (1989). *Exploring Complexity: An Introduction*. W. H. Freeman.

Parrondo, J. M. R., Horowitz, J. M., & Sagawa, T. (2015). "Thermodynamics of information." *Nature Physics*, 11(2), 131-139.

Prigogine, I. (1947). *Étude Thermodynamique des Phénomènes Irréversibles*. Dunod.

Prigogine, I., & Nicolis, G. (1977). *Self-Organization in Non-Equilibrium Systems*. Wiley.

Sagawa, T., & Ueda, M. (2008). "Second law of thermodynamics with discrete quantum feedback control." *Physical Review Letters*, 100(8), 080403.

Seifert, U. (2012). "Stochastic thermodynamics, fluctuation theorems and molecular machines." *Reports on Progress in Physics*, 75(12), 126001.

Sekimoto, K. (2010). *Stochastic Energetics*. Springer.

Shannon, C. E. (1948). "A Mathematical Theory of Communication." *Bell System Technical Journal*, 27(3), 379-423.

Sterling, P., & Eyer, J. (1988). "Allostasis: A New Paradigm to Explain Arousal Pathology." In *Handbook of Life Stress, Cognition and Health*, 629-649.

Suggested citation: Galida, R. S. (2026). Excess Entropy Production as a Candidate Universal Cost of Persistence: A

Deriving Corrective Permeability from the Cumulative Deviation Functional; Robert Galida (June 2026) [F]

Abstract

The attractor framework defines κ (corrective permeability) as the rate at which a system returns to its attractor after perturbation. Historically, κ has been treated as an empirical parameter – fitted to data rather than derived from first principles. This paper derives κ from the framework's foundational object: the cumulative deviation functional $DT(x) = \int_0^T \delta(\phi^t(x)) dt$ $DT_\infty(x) = \int_0^\infty \delta(\phi^t_\infty(x)) dt$, where $\delta(x) = d(x, A)$.

We define: $\kappa = \inf_{x \in B} \frac{\delta(x)}{D_\infty(x)}$ $\kappa = \inf_{x \in B} \frac{\delta(x)}{D_\infty(x)}$

We prove that for linear systems $x' = -Ax$ with A symmetric positive definite, this definition recovers the slowest eigenvalue $\lambda_{\min}(A)$ – the conventional notion of corrective permeability. We establish a sharp universal persistence bound $D_\infty(x) \leq \delta(x)/\kappa$, show homogeneity and scale invariance of the variational ratio, and demonstrate consistency with Koopman

spectral theory and resolvent poles for finite-dimensional linear systems. A comparison theorem links κ to classical exponential stability constants. A Hamilton-Jacobi-type transport equation for $D^\infty D^\infty$ is derived. A finite-horizon estimator $\kappa_T = \inf_x \delta(x) D_T(x) \kappa T = \inf_x \int_0^T D\tau(x) \delta(x)$ is provided with exponential convergence under explicit assumptions.

The derivation is rigorous for linear systems and testable. Open questions for nonlinear, multiscale, and stochastic systems are identified.

Keywords: corrective permeability, cumulative deviation functional, attractor framework, Koopman operator, trajectory functional

1. Introduction

The attractor framework has been applied across physics, biology, cognition, and social systems. Its central variable – corrective permeability κ – measures the rate at which a system returns to its attractor after perturbation. Historically, κ has been defined empirically as $\kappa = 1/\tau$, where τ is a measured recovery time constant.

This paper derives κ from a single foundational object: the cumulative deviation functional $D_T(x)$. Within the present framework, κ is defined variationally rather than introduced as an empirical fitting parameter. We show that κ is a consequence of the trajectory geometry – specifically, the ratio of initial distance to total cumulative deviation.

The derivation is rigorous for linear systems, connects to established theory (Koopman operators, resolvent poles), and provides a finite-horizon estimator for empirical use. Open

questions for nonlinear and stochastic systems are identified.

2. The Cumulative Deviation Functional

Let X be a metric space with distance function $\|\cdot\|$. Let $\phi_t(x)$ be the flow of a dynamical system starting from state $x \in X$ at time $t=0$. Let $A \subseteq X$ be an attractor set (a compact, invariant set to which trajectories converge). Let B be the basin of attraction of A .

Define the distance from a point to the attractor: $\delta(x) = d(x, A) = \inf_{a \in A} \|x - a\|$

Definition 1 (Cumulative Deviation Functional): For a finite horizon $T > 0$, define: $DT(x) = \int_0^T \delta(\phi_t(x)) dt$

For $T \rightarrow \infty$, define: $D^\infty(x) = \int_0^\infty \delta(\phi_t(x)) dt$

Proposition 1 (Finiteness of D^∞): Assume there exist constants $C < \infty$ and $\mu > 0$ such that: $\delta(\phi_t(x)) \leq Ce^{-\mu t} \delta(x)$

for all $x \in B$. Then $D^\infty(x) < \infty$ for every $x \in B$.

Proof: $D^\infty(x) = \int_0^\infty \delta(\phi_t(x)) dt \leq \int_0^\infty Ce^{-\mu t} \delta(x) dt = C\mu \delta(x) < \infty$

Properties (from Galida, 2026a):

Property	Statement
Non-negativity	$DT(x) \geq 0$
Monotonicity	$DT_2(x) \geq DT_1(x)$ for $T_2 \geq T_1$

Property	Statement
Additivity	$DT+S(x)=DT(x)+DS(\phi^T(x))$ $DT+S_\square(x)=DT(x)+DS_\square(\phi^{T_\square}(x))$
Instantaneous growth	$dDT(x)=\delta(\phi^T(x))dTD$ $dDT_\square(x)=\delta(\phi^{T_\square}(x))dDT_\square$
Occupation measure	$DT(x)=\int \delta(y) d\mu^T(y)$ $DT_\square(x)=\int \delta(y) d\mu^{T_\square}(y)$, where $\mu^T, \mu^{T_\square}$ is the occupation measure

3. Derivation of Corrective Permeability (κ)

3.1 Variational Definition

Definition 2 (Corrective Permeability): $\kappa = \inf_{x \in B} \frac{\delta(x)}{D^\infty(x)}$ $\kappa = \inf_{x \in B} \frac{\delta(x)}{D^\infty_\square(x)}$

Interpretation: κ is the *effective* recovery rate – the smallest ratio of initial distance to total cumulative deviation. It serves as a global measure of the slowest recovery mode in the basin.

Remark on κ : The definition allows $\kappa=0$ if $D^\infty(x)$ diverges or if the ratio $\delta(x)/D^\infty_\square(x)$ can be made arbitrarily small. Throughout the remainder of this paper, we assume hypotheses (such as the exponential stability in Proposition 1) that guarantee $\kappa>0$.

Remark on attainment: The infimum in the definition of κ need not be attained; minimizing sequences may exist without a minimizing state. For linear systems, the infimum is attained on the slow eigenspace.

3.2 Homogeneity and Scale Invariance

Theorem 1 (Homogeneity and Scale Invariance): Suppose the flow satisfies $\phi_t(\alpha x) = \alpha \phi_t(x)$ and $\phi_t(\alpha x) = \alpha \phi_t(x)$ for all t and all $\alpha > 0$, and the distance function satisfies $\delta(\alpha x) = \alpha \delta(x)$.

Then: $D^\infty(\alpha x) = D^\infty(x)$ and $D^\infty(\alpha x) = D^\infty(x)$.

Proof: $D^\infty(\alpha x) = \int_0^\infty \delta(\phi_t(\alpha x)) dt = \int_0^\infty \delta(\alpha \phi_t(x)) dt = \alpha \int_0^\infty \delta(\phi_t(x)) dt = \alpha D^\infty(x)$
 $D^\infty(\alpha x) = \int_0^\infty \delta(\phi_t(\alpha x)) dt = \int_0^\infty \delta(\alpha \phi_t(x)) dt = \alpha \int_0^\infty \delta(\phi_t(x)) dt = \alpha D^\infty(x)$

Corollary: For linear systems, the infimum over all $x \neq 0$ reduces to an infimum over the unit sphere: $\kappa = \inf_{x \neq 0} \frac{\delta(x)}{D^\infty(x)} = \inf_{\|x\|=1} \delta(x)$

3.3 Sharp Universal Persistence Bound

Theorem 2 (Sharp Universal Persistence Bound): For any $x \in B_A$, $D^\infty(x) \leq \delta(x) \kappa$ and $D^\infty(x) \geq \delta(x) / \kappa$.

Moreover, the constant $1/\kappa$ is optimal: it is the *smallest* constant such that this inequality holds for all x in the basin.

Proof: By definition of κ as the infimum of $\delta(x)/D^\infty(x)$, we have $\delta(x)/D^\infty(x) \geq \kappa$ for all x . Rearranging gives: $D^\infty(x) \leq \delta(x) / \kappa$.

Optimality follows from Theorem 3: for the slow eigenvector v_1 , $D^\infty(v_1) = \delta(v_1) / \kappa$, so no smaller constant can work. $\square$

3.4 Consistency with Linear Systems

Consider a linear system $\dot{x} = -Ax$, with A symmetric positive definite. Let its eigenvalues be $0 < \lambda_1 \leq \lambda_2 \leq \dots \leq \lambda_n$, with corresponding orthonormal eigenvectors $v_1, v_2, \dots, v_n$.

The flow is $\phi_t(x) = e^{-At}x$. The attractor is $A = \{0\}$, and the distance to the attractor is $\delta(x) = \|x\|$.

Theorem 3 (Linear Consistency): For $\dot{x} = -Ax$ with A symmetric positive definite, $\inf_{x \neq 0} \frac{D^\infty(x)}{\|x\|} = \lambda_{\min}(A)$.

Proof:

Since A is symmetric positive definite, e^{-At} is symmetric positive definite with eigenvalues $e^{-\lambda_i t}$. Hence its operator norm is $\|e^{-At}\| = e^{-\lambda_1 t}$. For any $x \neq 0$, $D^\infty(x) = \int_0^\infty \|e^{-At}x\| dt \leq \int_0^\infty \|x\| e^{-\lambda_1 t} dt = \|x\| \lambda_1^{-1}$.

Therefore: $\|x\| D^\infty(x) \geq \lambda_1^{-1} \|x\|^2$.

To show equality is achieved, take $x = v_1$ (the eigenvector corresponding to λ_1). Then: $\|e^{-At}v_1\| = \|v_1\| e^{-\lambda_1 t}$.

and: $D^\infty(v_1) = \int_0^\infty \|v_1\| e^{-\lambda_1 t} dt = \|v_1\| \lambda_1^{-1}$.

Thus: $\|v_1\| D^\infty(v_1) = \lambda_1^{-1} \|v_1\|^2$.

Hence: $\inf_{x \neq 0} \frac{D^\infty(x)}{\|x\|} = \lambda_1^{-1}$.

Corollary: For linear systems, the variational definition of κ recovers the slowest eigenvalue – the conventional notion of corrective permeability.

3.5 Transport Equation

Theorem 4 (Transport Equation): Assume the vector field f is C^1 , the flow ϕ_t is C^1 , and D^∞ is continuously differentiable on $B \subset A$. Then: $\nabla D^\infty(x) \cdot f(x) = -\delta(x)$ $\nabla D^\infty(x) \cdot f(x) = -\delta(x)$

Proof: From the definition: $D^\infty(\phi_s(x)) = D^\infty(x) - Ds(x) D^\infty(\phi_s(x)) = D^\infty(x) - Ds(x)$

Differentiating with respect to s at $s=0$: $\frac{d}{ds} D^\infty(\phi_s(x))|_{s=0} = -\delta(x) \frac{d}{ds} D^\infty(\phi_s(x))|_{s=0} = -\delta(x)$

By the chain rule: $\nabla D^\infty(x) \cdot f(x) = -\delta(x) \nabla D^\infty(x) \cdot f(x) = -\delta(x)$ $\square$

Interpretation: This is a first-order transport equation, $f \cdot \nabla D = -\delta$, which belongs to the broader Hamilton-Jacobi family but lacks a Hamiltonian in the usual sense. It may serve as a foundation for numerical computation and further theoretical development.

3.6 Local vs. Global Interpretation

The variational definition $\kappa = \inf_x \delta(x) D^\infty(x)$ is **global** – it is the slowest recovery rate over the entire basin. This is not necessarily the same as the local recovery rate near the attractor (the slowest eigenvalue of the linearization). For linear systems, they coincide. For nonlinear systems, they may differ if transient excursions produce slower effective recovery than the local linearization predicts.

This distinction is important: κ is a global invariant of the

basin, not merely a local property of the attractor. The relationship between the global κ and the local Lyapunov exponent is an open question (see §6).

3.7 Non-Symmetric Linear Systems

For a general linear system $\dot{x} = Ax$ (where A is stable, i.e., all eigenvalues have negative real parts), the same principle holds in the diagonalizable case. The slowest mode corresponds to the eigenvalue with the largest real part (closest to zero).

Conjecture: An analogous result holds for non-normal linear systems under additional assumptions on the semigroup, such as a uniformly exponentially stable semigroup satisfying suitable norm bounds. This remains an open question.

3.8 Comparison with Exponential Stability

Theorem 5 (Comparison with Exponential Stability): Suppose the system satisfies the exponential stability bound: $\delta(\phi_t(x)) \leq C e^{-\mu t} \delta(x)$

for all $x \in B$, with constants $C < \infty$ and $\mu > 0$. Then: $\kappa \geq \mu C$

Proof: From the stability bound: $D^\infty(x) = \int_0^\infty \delta(\phi_t(x)) dt \leq \int_0^\infty C e^{-\mu t} \delta(x) dt = C \mu \delta(x)$

Therefore: $\delta(x) D^\infty(x) \geq \mu C \delta(x)^2 \geq C \mu$

Taking the infimum over x : $\kappa = \inf_x \delta(x) D^\infty(x) \geq \mu C = \inf_x \delta(x)^2 \geq C \mu$

Interpretation: The variational constant κ is bounded below by the exponential stability constant μ/C .

4. Connections to Existing Theory

4.1 Koopman Operator

The Koopman operator K_t acts on observables as: $(K_t f)(x) = f(\phi_t(x))$

For linear systems $\dot{x} = -Ax$, the Koopman eigenvalues are $e^{-\lambda_i t}$. The dominant nontrivial eigenvalue (largest less than 1) is $e^{-\lambda_1 t}$, corresponding to the slowest decay rate.

For finite-dimensional linear systems, $\rho = e^{-\lambda_{\min} t}$, and therefore: $-\frac{1}{t} \log \rho = \lambda_{\min} = \kappa$

Thus, under the hypotheses of Theorem 3, the variational constant equals the exponential decay rate associated with the dominant Koopman eigenvalue.

4.2 Resolvent Poles

For finite-dimensional stable linear systems, the resolvent $(sI + A)^{-1}$ has poles at $s = -\lambda_i$. The pole closest to the imaginary axis is $s = -\lambda_1$.

Since Theorem 3 identifies $\kappa = \lambda_{\min}$, and the resolvent poles are $s_i = -\lambda_i$, we obtain: $\kappa = \min_i \operatorname{Re}(s_i)$

for finite-dimensional linear systems.

5. Finite-Horizon Estimation

In practice, we can only measure finite trajectories. Define the finite-horizon estimator: $\kappa_T = \inf_{x \in K} \delta(x) D_T(x)$ $\kappa_T = \inf_{x \in K} \delta(x) D_T(x)$

where $K \subset B_K \subset B$ is compact and $K \cap A = \emptyset$.

Proposition 2 (Finite-Horizon Estimation): Assume:

1. The flow $\phi_t(x)$ is jointly continuous in (t, x) .
2. $\delta(x)$ is continuous.
3. The exponential stability bound $\delta(\phi_t(x)) \leq C e^{-\mu t} \delta(x)$ holds uniformly for all $x \in K$, with $\mu > 0$.

Then the variational constant κ (from Definition 2) satisfies $\kappa \geq \mu/C$ by Theorem 5, and: $\kappa_T \rightarrow \kappa$ as $T \rightarrow \infty$

with error: $|\kappa_T - \kappa| = O(e^{-\mu T})$

Proof: For any $x \in K$, the tail bound gives: $|D_\infty(x) - D_T(x)| = \int_T^\infty \delta(\phi_t(x)) dt \leq C e^{-\mu T} \delta(x) \mu |D_\infty(x) - D_T(x)| = \int_T^\infty \delta(\phi_t(x)) dt \leq \mu C e^{-\mu T} \delta(x)$

Since $\delta(x)$ is bounded on the compact set K , let $M = \sup_{x \in K} \delta(x) < \infty$. Then: $|D_\infty(x) - D_T(x)| \leq C M e^{-\mu T} \mu |D_\infty(x) - D_T(x)| \leq \mu C M e^{-\mu T}$

The right-hand side is independent of x and tends to zero as $T \rightarrow \infty$. Hence $D_T \rightarrow D_\infty$ uniformly on K .

Moreover, since K is compact and $K \cap A = \emptyset$, continuity of δ gives $\inf_{x \in K} \delta(x) > 0$. Since $D_T(x)$ is continuous (by assumptions 1–2) and monotonically non-

decreasing in T (from §2), for any fixed finite $T_0 > 0$ $T_0 > 0$, $D_\infty(x) \geq D_{T_0}(x) D_\infty(x) \geq D_{T_0}(x)$, and $D_{T_0} D_{T_0}(x)$ is continuous and strictly positive on K . A continuous, strictly positive function on a compact set has a positive infimum: $m = \inf_{x \in K} D_{T_0}(x) > 0$ $m = \inf_{x \in K} D_{T_0}(x) > 0$

Thus: $\inf_{x \in K} D_\infty(x) \geq m > 0$ $\inf_{x \in K} D_\infty(x) \geq m > 0$

Uniform convergence of D_{T_0} to D_∞ on K therefore implies uniform convergence of $\delta(x)/D_{T_0}(x)$ to $\delta(x)/D_\infty(x)$. Consequently, the infima converge. $\square$

6. Open Questions

Question	Status	Difficulty
Q1: Nonlinear systems	Does $\inf \delta D_\infty \inf D_\infty \delta$ equal the local Lyapunov exponent?	Hard
Q2: Local vs. global consistency	Does $\lim_{x \rightarrow A} \delta(x) D_\infty(x) = \kappa \lim_{x \rightarrow A} D_\infty(x) \delta(x) = \kappa$ hold for general nonlinear systems?	Hard
Q3: Non-normal systems	Does the infimum equal the slowest eigenvalue for non-normal A ?	Moderate
Q4: Multiple timescales	Does the infimum isolate the slowest timescale?	Hard
Q5: Stochastic systems	How does noise affect the finite-horizon estimator?	Hard
Q6: Multiple attractors	How does κ behave in basins with multiple attractors?	Moderate

7. Conclusion

This paper derives corrective permeability κ from the cumulative deviation functional $DT(x)DT^\square(x)$. The variational definition: $\kappa = \inf_x \delta(x) D^\infty(x) \kappa = x \inf_x D^\infty^\square(x) \delta(x)$

is shown to recover the slowest eigenvalue for linear systems, consistent with the conventional empirical definition $\kappa = 1/\tau$. A sharp universal persistence bound $D^\infty(x) \leq \delta(x)/\kappa$ $D^\infty^\square(x) \leq \delta(x)/\kappa$ is established. A comparison theorem links κ to classical exponential stability constants. A Hamilton-Jacobi-type transport equation for D^∞ is derived. Connections to Koopman theory and resolvent theory are established for finite-dimensional linear systems. A finite-horizon estimator $\kappa T \kappa T^\square$ is provided with exponential convergence under explicit assumptions.

Key contribution: Within the present framework, κ is defined variationally rather than introduced as an empirical fitting parameter – at least for the class of systems analyzed here.

Next steps: Extend the derivation to nonlinear systems (Q1–Q2), non-normal systems (Q3), multiple timescales (Q4), and stochastic dynamics (Q5).

References

Crandall, M. G., Ishii, H., & Lions, P. L. (1992). “User’s Guide to Viscosity Solutions of Second Order Partial Differential Equations.” *Bulletin of the American Mathematical Society*, 27(1), 1-67.

Evans, L. C. (2010). *Partial Differential Equations*. American

Mathematical Society.

Galida, R. (2026a). "The Persistence Functional: A Candidate Formal Foundation for the Attractor Framework." *Fantasy Attractor*.

Hale, J. K. (1988). *Asymptotic Behavior of Dissipative Systems*. American Mathematical Society.

Hirsch, M. W., Smale, S., & Devaney, R. L. (2004). *Differential Equations, Dynamical Systems, and an Introduction to Chaos* (2nd ed.). Elsevier Academic Press.

Khalil, H. K. (2002). *Nonlinear Systems* (3rd ed.). Prentice Hall.

Koopman, B. O. (1931). "Hamiltonian Systems and Transformations in Hilbert Space." *Proceedings of the National Academy of Sciences*, 17(5), 315-318.

Lyapunov, A. M. (1892). *The General Problem of the Stability of Motion*. (English translation: 1992, Taylor & Francis).

Mezić, I. (2005). "Spectral Properties of Dynamical Systems, Model Reduction and Decompositions." *Nonlinear Dynamics*, 41(1-3), 309-325.

Pazy, A. (1983). *Semigroups of Linear Operators and Applications to Partial Differential Equations*. Springer.

Vidyasagar, M. (1993). *Nonlinear Systems Analysis* (2nd ed.). Prentice Hall.

Suggested citation: Galida, R. S. (2026). Deriving Corrective Permeability from the Cumulative Deviation Functional. *Fantasy Attractor*.

The Persistence Functional: A Mathematical Measure of Attractor Resilience

Author: Robert Galida

Date: May 2026

Abstract

The attractor framework says that **persistence under disturbance** is the basic mark of reality.

To turn this idea into a formal science, we introduce the **persistence functional** $P(x)P(x)$.

$P(x)P(x)$ is a single number that measures:

- How deep a state is inside an attractor basin.
- How quickly it returns after a knock.

We define $P(x)P(x)$ for three different kinds of systems:

1. **Deterministic dissipative systems** – here PP is linked to Lyapunov exponents and basin stability.
2. **Stochastic systems** – here PP is linked to escape time and quasipotential.
3. **Information-theoretic systems** – here PP is linked to negative free energy or mutual information.

The **recovery rate** $-P'/P - P'/P$ is a universal sign of **critical slowing down** – a warning that a system is about to tip.

We also discuss limitations: resilience may depend on direction (“anisotropic”), and multiple timescales may need **vector** or **tensor** persistence. We list open mathematical problems.

This paper is a **roadmap**, not a finished theory.

1. Introduction

In the attractor framework, **persistence under disturbance** is central. But we have not had a single number to say *how persistent* a state is.

The **persistence functional** $P(x)$ aims to fill that gap.

What $P(x)$ should do:

- $P(x) > 0$ for states inside an attractor basin.
- For a **conservative attractor** (like a free electron), P is maximal (normalised to 1).
- For a **dissipative attractor**, P drops after a disturbance and then recovers.

The recovery rate $-P'/P - P'/P$ equals:

- the negative of the largest Lyapunov exponent (for deterministic systems)
 - the inverse return time (for stochastic systems)
 - the rate of information loss (for informational systems)
-
- P falls as the system approaches a **bifurcation**, giving early warning.

We do **not** give one universal formula. Instead, we give

a **family** of definitions, each suited to a different type of system, all united by the same purpose – measuring resilience.

2. Deterministic Dissipative Systems

Consider a smooth system $\dot{x}=f(x)$ with a stable attractor A and its basin $B(A)$.

A natural candidate for $P(x)$ uses a **Lyapunov function** $V(x)$ – a kind of energy that always decreases inside the basin ($\dot{V}<0$).

We define: $P(x) = 1 - V(x) - V_{\max} - V_A$

This gives $P=1$ on the attractor and $P \rightarrow 0$ at the basin boundary.

Near the attractor, the recovery rate is related to the **largest Lyapunov exponent** λ_1 : $-\dot{P}/P \approx -\lambda_1$

When the system approaches a tipping point, $\lambda_1 \rightarrow 0^-$, so the recovery rate slows down – this is **critical slowing down**.

Conclusion: For deterministic systems, P can be built from a Lyapunov function. The recovery rate equals the negative of the largest Lyapunov exponent.

3. Stochastic Systems

When noise is present, persistence is about how long it takes to escape from the basin.

The **mean first passage time** $\tau(x)$ – the average time to leave – is a natural measure.

We define: $P(x) = \tau(x) / \tau_{\max}$ $P(x) = \tau(x) / \tau_{\max}$

where $\tau_{\max}$ is the value at the attractor.

For weak noise, $\tau(x)$ grows exponentially with the **quasipotential** $U(x)$ (Freidlin–Wentzell theory): $\tau(x) \sim e^{U(x)/\epsilon}$

So: $P(x) \propto e^{-(U_{\max} - U(x))/\epsilon}$

The recovery rate is the inverse of the return time. As a tipping point is approached, the return time diverges, and the recovery rate goes to zero. This again gives **critical slowing down** – rising variance and autocorrelation.

Conclusion: For stochastic systems, P is proportional to the mean exit time (or the exponential of the quasipotential). This connects persistence to large deviation theory.

4. Information-Theoretic Systems

For systems where information matters (neural, cognitive, social), we can define persistence using **mutual information** between past and future.

Let $I_{\text{past}, \text{future}}$ be the **predictive information**. Then: $P(t) = I(\text{past}; \text{future at time } t)$ or $P = e^{-\text{surprisal}}$

The decay of $P(t)$ over time measures **memory loss**.

Landauer's principle connects information loss to entropy production: $P'/P \leq -S' / k_B \ln 2$ $P'/P \leq -k_B \ln 2 S'$

Alternatively, in the **free energy principle** (Friston), the negative free energy $-F$ acts like a Lyapunov function. We can set: $P = e^{-F/kT}$ or $P = -F$

Then $-P'/P - P'/P$ is the rate of free energy minimisation, which slows near bifurcations.

Conclusion: For information-theoretic systems, PP can be defined via mutual information decay or negative free energy, linking persistence to entropy production and predictive coding.

5. Unifying Recovery Rate and Critical Slowing Down

Across all types of systems, the **recovery rate** $\lambda_{\text{rec}} = -P'/P$ (just after a small disturbance) is a universal indicator:

- **Deterministic dissipative:** $\lambda_{\text{rec}} = -\lambda_1$ (absolute value of the largest Lyapunov exponent)
- **Stochastic:** λ_{rec} = inverse of the return time, related to the quasipotential's curvature
- **Information-theoretic:** λ_{rec} = rate of free energy minimisation or information loss

As the system approaches a bifurcation, $\lambda_{\text{rec}} \rightarrow 0$. This is **critical slowing down**.

It shows up as rising lag-1 autocorrelation and variance (Scheffer et al., 2009).

So PP and its recovery rate give early warnings.

6. Normalisation for Conservative

Attractors

For a perfect **conservative attractor** (e.g., an electron in its ground state, no decay), the persistence functional should be constant and maximal: $P_{\text{cons}} = 1$ for all times t .

No recovery rate is defined (or it is zero). This anchors the scale.

For **emergent approximate conservative systems** (like atomic clocks), PP is very close to 1 and decays extremely slowly.

7. Limitations – Scalar Collapse and Anisotropic Resilience

A single scalar $P(x)$ may not be enough for systems where resilience is **anisotropic** – that is, recovery speed depends on the direction of the perturbation.

High-dimensional systems can have **multiple timescales** (fast and slow modes). A scalar average can miss important structure.

Future work may need:

- **Vector persistence** – a list of recovery rates along different directions.
- **Tensor persistence** – a metric that captures the full shape of the basin.
- **Persistence manifold** – the geometry of the basin in state space.

We accept this limitation. The scalar PP is a useful first approximation for systems with isotropic resilience or for early-warning applications where a single number is enough.

For complex systems, a multidimensional generalisation is an open research problem.

8. Open Mathematical Problems

1. **Derive $P(x)$ from first principles** for a given class of systems (e.g., from a variational principle).
 2. **Prove that $-P'/P = \lambda_1$** for a wide class of dissipative systems.
 3. **Extend the definition to systems with multiple attractors and chaotic basins** (where basin stability is fractal).
 4. **Establish a rigorous relationship between PP and the mutual information decay rate** for non-equilibrium processes.
 5. **Formulate a universal persistence functional** that works across all regimes – or prove it's impossible.
 6. **Test the predictive power of PP** in controlled experiments (e.g., ecological microcosms, neural cultures, social media sentiment).
 7. **Develop vector/tensor persistence** for anisotropic resilience.
-

9. Conclusion

The persistence functional $P(x)$ gives a mathematical language for attractor resilience.

We have given **operational definitions** for three regimes:

- **Deterministic dissipative** → Lyapunov / basin stability

- **Stochastic** → escape time / quasipotential
- **Information-theoretic** → mutual information / free energy

The **recovery rate** $-P'/P - P'/P$ unifies critical slowing down across all these domains.

We have explicitly noted **limitations** (scalar collapse, anisotropy) as open problems.

This paper is a **roadmap**, not a final theory. The framework now has a quantitative step.

Suggested citation: Galida, R. S. (2026). *The Persistence Functional: Towards a Mathematical Measure of Attractor Resilience (Reader-Friendly Version)*. Fantasy Attractor.