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** $P(x)$.

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

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

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

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

The **recovery rate** $-\dot{P}/P$ 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 $-\dot{P}/P$ 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:$P(x)=1-\frac{V(x)-V_A}{V_{\max}-V_A}$

This gives $P=1$ on the attractor and $P\to 0$ 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, $\lambda_1\to 0^-$, 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: $P(x)=\frac{\tau(x)}{\tau_{\max}}$

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

The decay of $P(t)$ over time measures **memory loss**. Landauer’s principle connects information loss to entropy production: $\dot{P}/P\le-\frac{\dot{S}}{k_B\ln 2}$

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

Then $-\dot{P}/P$ $-\dot{P}/P$ is the rate of free energy minimisation, which slows near bifurcations.

Conclusion: For information-theoretic systems, P P 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** $\lambda_{\text{rec}}=-\dot{P}/P$ $\lambda_{\text{rec}}=-\dot{P}/P$ (just after a small disturbance) is a universal indicator:

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

As the system approaches a bifurcation, $\lambda_{\text{rec}}\to 0$ $\lambda_{\text{rec}}\to 0$. This is **critical slowing down**.
It shows up as rising lag-1 autocorrelation and variance (Scheffer et al., 2009).
So P P 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:Pcons=1for all times$P_{\text{cons}}=1\text{for all times}$

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)$P(x)$ 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 $-\dot{P}/P = \langle \lambda_1 \rangle$** 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 P 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 P** 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** $-\dot{P}/P-\dot{P}/P$ 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.

Metronome, Memory, and the Threefold Anchor: A Relational Account of Time [F] (2026)

Abstract

This paper presents a relational view of time based on the attractor framework.

We argue that two very different kinds of attractors work together to create what we call time:

- **Conservative attractors** (electrons, neutrinos, protons)

act as metronomes. They provide a steady, repeatable rhythm — a ruler for measuring duration.
- **Dissipative attractors** (living cells, minds, societies) act as memory. They accumulate irreversible changes, giving time its direction.

Time is not a mysterious substance. It is the coupling between these three fundamental metronomes and the irreversible flow of memory. What binds all dissipative systems — from a bacterium to a brain to a galaxy — is the continuous recycling of the same three eternal metronomes.

This view offers a conceptual account of how clocks work, why time has an arrow, and how aging, entropy, and history fit together.

The dance of time has three metronomes and a memory.

---

1. Two Classes of Persistence, Two Roles for Time

In the attractor framework, everything that persists does so by resisting disturbance. We identify two distinct types of persistent structures, each giving rise to a different aspect of time.

1.1 Conservative Attractors — The Metronome

Conservative attractors are protected by physical conservation laws (charge, baryon number, energy). They are:

- **Eternal** — they do not age or decay (or are effectively stable on all observable timescales).
- **Time-symmetric at the level of intrinsic persistence** — their existence as attractors is symmetric under time

reversal, though some interactions (weak force) violate CP and thus T.
- **Type-identical** – every electron has the same Compton frequency; every neutrino mass eigenstate has an invariant (though not yet precisely measured) frequency.

Because of these properties, conservative attractors serve as reference standards for duration – metronomes. The international definition of the second is literally a fixed number of such ticks.

1.2 Dissipative Attractors – Memory

Dissipative attractors (cells, minds, ecosystems, societies) are different:

- They require a continuous flow of energy and must export entropy.
- Their dynamics are irreversible – you cannot return to a past microstate without enormous cost.
- This irreversibility creates a directional arrow: before and after, past and future.
- They accumulate memory – irreversible state changes that persist and affect future behaviour.

Memory = irreversible accumulated state change (inscription). Examples: synaptic plasticity, scars, fossil records, cultural archives, radioactive decay (the daughter nucleus retains a record of the parent's disintegration).

---

2. The Three Metronomes: Our Most Fundamental Clocks

The Standard Model contains many particles, but only three classes are absolutely or effectively stable and serve as fundamental metronomes. The photon is not a metronome — it has zero rest mass, hence no rest-frame Compton frequency. It is a mode of propagation, not a standalone persistent entity.

Class / Particle	**Symbol**	**Key Property**	**Role as Metronome**
Electron	e^-	lightest charged lepton	Compton frequency ~1.24 × 10^{20} Hz
Neutrino mass eigenstates (collectively)	ν_1, ν_2, ν_3	neutral, tiny masses	Compton frequencies (mass-dependent); effectively stable
Proton	p	lightest baryon	Compton frequency ~2.27 × 10^{23} Hz; no observed decay

These three classes form what the framework calls the *eternal skeleton* — the collection of conservative structures that persist without decay and provide the stable background against which dissipative change occurs.

Stability notes

- Proton decay has never been observed; lower limit on half-life > 10^{34} years — effectively eternal. The proton is composite, but its stability derives from baryon number conservation, not merely nuclear binding energy.
- Neutrinos oscillate between flavours, but the underlying mass eigenstates are stable on cosmological timescales. Their exact Compton frequencies are not yet known to metrological precision — only mass-squared differences have been measured — but they are theoretically invariant.

These three metronomes do not need energy input to persist. Their frequencies are invariant (known for electron and proton; theoretically invariant for neutrinos). Any clock based on one agrees with any other after accounting for relativity, as confirmed by atomic clock comparisons.

---

3. Time as the Coupling Between Metronomes and Memory

Time is not a primitive substance. It is the relationship between the metronome ensemble and dissipative memory.

- The three metronomes provide a metric — an invariant ruler for “how much” duration has passed.
- Memory provides direction — which events are past, which are future.
- Without metronomes, change would be unmeasurable — no ruler.
- Without memory, change would be reversible and directionless — no before/after.

Both are necessary for what we operationally call time.

As a working placeholder, let the rate of memory inscription be $dM/dt=f(M,\nu)$, where ν is a characteristic metronome frequency and M is the current accumulated memory state. Two limiting cases anchor the idea:

- As $\nu\to 0$ — no metronome — duration becomes undefined. Change occurs but cannot be quantified as a metric interval. This is the “no ruler” condition.
- As dissipation $\to 0$ — no memory — M remains constant. Change leaves no trace, so there is no before/after. This is the “no arrow” condition.

What binds all dissipative systems — a bacterial cell, a human brain, a galaxy, a social institution — is the continuous **recycling of the same three eternal metronomes**. Every dissipative system operates by exchanging electrons, protons, and neutrinos with its environment. The metronomes are the invariant substrate; the memory is the transient pattern. The coupling is the recycling.

Thus, time is not merely a coordinate; it is the ongoing, irreversible reconfiguration of eternal components into transient, memory-bearing structures.

The three metronomes are time-symmetric at the level of intrinsic persistence. The arrow of time comes from dissipative systems that accumulate history. Time is the coupling between these two regimes.

4. Thermodynamic Information Theory and Persistence

The persistence functional P(x)$P(x)$ measures how deep an attractor basin is — formally, the depth of the basin in the system’s phase space (the energy or Lyapunov function value required to escape the basin). Higher PP means a more stable attractor.

- In a dissipative attractor, maintaining memory requires continuous energy export to counteract thermal noise.
- Landauer’s principle: erasing one bit costs at least kBTln2$k_B T \ln 2$ of free energy. Retaining memory against thermal fluctuations requires energy input.

We interpret P(x)$P(x)$ as a measure of information retention: systems with higher PP preserve mutual information between past and present for longer. The decay rate −Ṗ/P$-\dot{P}/P$ relates to entropy production, connecting the attractor framework to non-equilibrium thermodynamics.

---

5. Consequences and Applications

- **Clocks** – Atomic clocks derive stability from electron transitions. The three metronomes guarantee cross-calibration.
- **Aging** – Biological aging is the accumulation of irreversible memory, measured against metronomes like circadian rhythms.
- **Critical slowing down** – As a system approaches a bifurcation, $-\dot{P}/P-\dot{P}/P$ decreases, providing early-warning signals (rising autocorrelation, variance) in physiology, ecology, and social systems.
- **Hysteresis in beliefs** – Fantasy attractors exhibit hysteresis – the path of belief change differs when accumulating vs. removing evidence. The hysteresis loop area quantifies memory.[1]
- **Cosmological time** – The cosmic microwave background is a memory of the early universe (here “memory” is metaphorical). Atomic clocks measure the duration since those imprints were formed.

[1] *Fantasy attractor*: in the attractor framework, a dissipative structure (typically a belief system) with abnormally low corrective permeability, resistant to updating despite counter-evidence.

---

6. Relation to the Broader Attractor

Framework

The metronome-memory distinction is a special case of the conservative vs. dissipative attractor dichotomy. It sharpens the “eternal skeleton / transient dance” metaphor.

The three metronomes are the most fundamental layer of the eternal skeleton — the collection of conservative structures that persist without decay and provide the stable background against which dissipative change occurs.

The framework does not claim that time is “made of” attractors. It claims that the measurement and experience of time rely on the interaction of these two persistence regimes. Because every dissipative system continuously recycles the same eternal metronomes, all such systems are materially unified across space and time. That unity is what makes a universal, relational time possible.

---

7. Open Questions and Refinements

- **Formalising P(x)$P(x)$** — Rigorous derivation for deterministic (Lyapunov), stochastic (escape time), and information-theoretic (surprisal) cases.
- **Coupling equations** — Specify dM/dt=f(M,ν)$dM/dt=f(M,\nu)$. Can it be tested empirically?
- **Category clarity** — Conservative attractors span strict symmetry-protected invariants (elementary particles) and emergent approximate invariants (clocks). Future work should stratify these.
- **Falsifiability** — Concrete falsifiers: a persistent system without dissipation, or a social attractor that never updates despite counter-evidence.
- **Relation to other relational accounts** — Converges with Barbour (1999) and Rovelli (1996). The difference: the

present framework identifies the two required poles (conservative metronomes providing metric invariance; dissipative memory providing direction) and grounds both in attractor dynamics.

8. Conclusion

Time is not a primitive. It is the relational coupling between:

- the three fundamental conservative attractor classes — electron, neutrino mass eigenstates (collectively), and proton — which provide invariant metric structure (the metronome), and
- dissipative systems that accumulate irreversible state inscription (memory).

What binds all dissipative systems — from a bacterium to a brain to a galaxy — is the continuous recycling of the same three eternal metronomes. The metronomes are the invariant substrate; memory is the transient pattern; time is the coupling.

This account respects how physics measures time, explains the arrow via entropy and information persistence, and offers transferable concepts across neuroscience, ecology, sociology, and AI.

The dance has three metronomes and a memory.

References

Barbour, J. (1999). *The End of Time.* Oxford University Press.
Rovelli, C. (1996). Relational quantum mechanics. *International Journal of Theoretical Physics*, 35(8), 1637–1678.

Suggested citation: Galida, R. S. (2026). Metronome, Memory, and the Threefold Anchor: A Relational Account of Time.

Barbour, J. (1999). *The End of Time.* Oxford University Press.

Rovelli, C. (1996). Relational quantum mechanics. *International Journal of Theoretical Physics*, 35(8), 1637–1678.

Suggested citation: Galida, R. S. (2026). Metronome, Memory, and the Threefold Anchor: A Relational Account of Time.