Preprints.org

Article

Not peer-reviewed version

The Epistemic Support-Point Filter as a Tropical Hamilton–Jacobi System Wavefront Propagation and Possibilistic Inference

Moriba Kemessia Jah *

Posted Date: 15 April 2026

doi: 10.20944/preprints202603.2110.v3

Keywords: tropical Hamilton-Jacobi equation; max-plus semiring; Lax-Oleinik operator; possibilistic inference; epistemic support-point filter; impossibility field; surprisal field; falsification boundary; tropical variety; wavefront propagation; possibilistic Cramér-Rao bound; minimum action principle; zero-temperature limit; possibility theory; non-Bayesian estimation; level-set evolution; epistemic geometry; Popperian falsification

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

Copyright: This open access article is published under a Creative Commons CC BY 4.0 license, which permit the free download, distribution, and reuse, provided that the author and preprint are cited in any reuse.

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 15 April 2026 doi:10.20944/preprints202603.2110.v3

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

Article

The Epistemic Support-Point Filter as a Tropical Hamilton–Jacobi System Wavefront Propagation and Possibilistic Inference

Moriba Kemessia Jah

Black Swan Research Group, GaiaVerse, Ltd., Jah Decision Intelligence Group, Aerospace Engineering & Engineering Mechanics, University of Texas at Austin; moriba@utexas.edu

Abstract

Any inference system satisfying the TEAG axioms must obey a tropical Hamilton–Jacobi equation in the max-plus semiring, and the Epistemic Support-Point Filter (ESPF) is that solution. This paper proves the necessity and derives the complete dynamical and causal geometry that follows. The necessity has three steps, each forced. Popperian contraction requires that evidence can only increase impossibility, never decrease it: under the log-admissibility transformation this forces the max-plus operation. The evidence-referencing axiom requires that survivor selection depends only on innovation geometry, not on prior structure: this forces the Hamiltonian to be momentum-independent. A momentum-independent tropical Hamiltonian forces the Lax–Oleinik reduction to a pointwise max. The result is unique: no alternative update structure consistent with these axioms exists. The ESPF is not one admissible filter among many — it is the unique admissible inference dynamics under TEAG. Two scalar fields live on hypothesis space. The *impossibility field* $\Phi_\varnothing = -\log \pi$ encodes accumulated epistemic history: zero where a hypothesis enjoys full prior support, growing without bound as evidence withdraws that support. The *surprisal field* $\Phi_S(h) = \frac{1}{2}\|L_e^{-1}(y - g(h))\|^2$ encodes the tension between each hypothesis and the current observation in MVEE-whitened measurement space. The conjunctive (Popperian) update produces the posterior impossibility field as their pointwise max-plus upper envelope: $\widetilde{\Phi}_\varnothing = \Phi_\varnothing \oplus \Phi_S = \max(\Phi_\varnothing, \Phi_S)$. This equality follows from $-\log \min(a, b) = \max(-\log a, -\log b)$: it is an algebraic identity, not a modeling choice or an analogy. The *active deformation front* — the tropical variety of this two-term polynomial, where both fields achieve the maximum simultaneously — is the exact locus where evidence begins to deform the posterior impossibility field. It is a necessary condition for falsification. Sufficient falsification requires exit from the PCRB-admissible basin, whose threshold is determined by the PCRB at each step. A scalar example derives every quantity in closed form, making the front geometry and basin structure visible without probabilistic machinery. The ESPF predict–update recursion is the Lax–Oleinik operator of max-plus optimal control: the one-step solution operator of the tropical Hamilton–Jacobi equation, with the surprisal field as Hamiltonian and the Possibilistic Cramér–Rao Bound (PCRB) as the minimum action per update cycle. This structure is not chosen: Proposition 5.2 proves it is forced by the TEAG axioms — specifically, Popperian contraction forces the max-plus operation, and the evidence-referencing condition forces momentum independence of the Hamiltonian. No alternative update structure consistent with these axioms exists. Falsification is wavefront propagation: the surprisal field radiates outward from each observation, and a hypothesis enters the active deformation front at the moment the surprisal wavefront overtakes the prior impossibility field. The active deformation front is the epistemic Lagrange point — the locus of exact balance between prior epistemic history and current evidence tension — and is a necessary condition for falsification. Sufficient falsification occurs when the wavefront has pushed a hypothesis outside the PCRB-admissible basin: the isotropic equipotential region whose threshold is governed by the PCRB at each step. The term *wavefront* denotes level-set evolution under max-plus dynamics; no physical medium is assumed. The gravitational language is structural: it reflects equivalence of governing equations, not shared physical ontology. The PCRB emerges as a minimum action principle: no measurement can compress the surviving well below

© 2026 by the author(s). Distributed under a Creative Commons CC BY license.

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 15 April 2026 doi:10.20944/preprints202603.2110.v3

the PCRB floor per update. The zero-temperature limit of the classical Hamilton–Jacobi equation — passing from log-sum-exp (probabilistic) to max (possibilistic) aggregation — recovers this framework exactly, making precise the passage from Bayesian to possibilistic inference as a thermodynamic degeneration. The whitened minimax medioid is proved to be the geodesic attractor of the surviving well: the realized support point — a member of the surviving cloud, never an interpolated abstraction — that minimizes worst-case whitened distance to all other survivors, and is therefore nearest the center of the PCRB-defined epistemic geoid in the MVEE-whitened metric. The correct geometric primitive of epistemic phase space is identified as a contact manifold rather than a symplectic one: the irreversibility of Popperian falsification forces a contact structure, the PCRB is a contact energy floor, and the ESPF implements a contact Hamiltonian system with discrete projection onto the admissible basin. In the absence of measurably informative evidence ($\bar{S}_k = 0$), the contact dissipation vanishes, the action variable is constant, and the dynamics recover symplectic geometry exactly — Liouville's theorem holds on the $\{\Phi = \text{const}\}$ level sets. The predict step of the ESPF has symplectic character for this reason; the update step breaks it whenever $\bar{S}_k > 0$. Epistemic time is discrete, advancing only upon registered evidence, and becomes continuous in the limit of constant information flux. Gravitational time dilation is identified as the physical instance of epistemic time dilation: both implement the same Hamilton–Jacobi shielding mechanism. These results constitute the dynamical foundation of the Theory of Epistemic Abductive Geometry (TEAG) (Jah 2026b).

Keywords: tropical Hamilton-Jacobi equation; max-plus semiring; Lax-Oleinik operator; possibilistic inference; epistemic support-point filter; impossibility field; surprisal field; falsification boundary; tropical variety; wavefront propagation; Possibilistic Cramér-Rao Bound; minimum action principle; zero-temperature limit; possibility theory; non-Bayesian estimation; level-set evolution; epistemic geometry; Popperian falsification

---

1. Introduction

Any inference system satisfying the TEAGaxioms must obey a tropical Hamilton–Jacobi equation in the max-plus semiring. The Epistemic Support-Point Filter is that solution — not one admissible filter among many, but the unique admissible inference dynamics under TEAG. This paper proves the necessity and derives the geometry that follows from it.

The proof has three steps, each forced by a different axiom. Popperian contraction requires that evidence can only increase impossibility: under the log-admissibility transformation this forces max-plus algebra. The evidence-referencing axiom requires that falsification depends only on innovation geometry, not on prior structure: this forces the Hamiltonian to be momentum-independent. A momentum-independent tropical Hamiltonian forces the Lax–Oleinik reduction to a pointwise max. No choice is made at any step.

The resulting dynamical picture is exact. In orbital mechanics, the fate of a body near a massive attractor is determined by its specific mechanical energy: bound trajectories converge, hyperbolic trajectories escape, and the transition surface is the locus where kinetic and potential energy exactly balance. The structure of epistemic inference under possibility theory is algebraically parallel: the impossibility field plays the role of potential energy, the surprisal field plays the role of the Hamiltonian, and the active deformation front is the locus where they balance — the epistemic Lagrange point. This parallel is not analogical. It is proved by algebraic identity under the same mathematical framework. The gravitational language used throughout is structural: it reflects equivalence of governing equations, not shared physical ontology.

In the ESPF (Jah and Haslett 2025a), the hypothesis space H carries two scalar fields. The *impossibility field* $\Phi_\varnothing = -\log \pi : H \to [0, \infty)$ is induced by the prior possibility π: it is zero where a hypothesis enjoys full prior epistemic support and grows without bound as prior support is withdrawn. The *surprisal field* $\Phi_S : H \to [0, \infty)$ encodes the tension between each hypothesis and the

© 2026 by the author(s). Distributed under a Creative Commons CC BY license.

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 15 April 2026 doi:10.20944/preprints202603.2110.v3

current observation in whitened MVEE-whitened measurement space. The conjunctive update — the possibilistic analog of Bayes' rule — produces the posterior impossibility field as the *pointwise tropical superposition* of these two fields: $\widetilde{\Phi}_\varnothing = \Phi_\varnothing \oplus \Phi_S = \max(\Phi_\varnothing, \Phi_S)$.

The *active deformation front* — the surface where evidence first matches prior impossibility and begins to deform the posterior field — is the tropical variety of this convolution: the locus where prior impossibility and surprisal are in exact balance. This is the epistemic Lagrange point, the analog of the L_1 equilibrium between two gravitational sources, and marks where falsification becomes possible. Sufficient falsification requires exit from the PCRB-admissible basin whose boundary is the isotropic equipotential at threshold $c_k^\star$. Hypotheses inside the basin are epistemically protected; those outside are expelled.

The ESPF's commitment point, the whitened minimax medioid $h^\star$, is an actual member of the surviving support cloud — not an interpolated geometric abstraction. It is the geodesic attractor of the surviving well: the support point nearest the center of the PCRB-defined epistemic geoid, maximally insulated from the geoid boundary in all directions under the whitened MVEE-whitened metric.

The Possibilistic Cramér–Rao Bound (PCRB) (Jah 2025b) bounds the minimum volume of the posterior support — equivalently, the maximum epistemic contraction per update. In the potential picture it is a minimum action principle: no measurement, however precise, can compress the surviving well below the PCRB floor. This is the epistemic analog of the SOI: within the PCRB radius, epistemic trajectories are bound; outside it, the evidence lacks the potential depth to capture.

The formal vehicle connecting these ideas is the tropical Hamilton–Jacobi equation. In classical mechanics, Hamilton's principal function $S(q,t)$ satisfies

$$\frac{\partial S}{\partial t} + H\left(q, \frac{\partial S}{\partial q}, t\right) = 0, \tag{1}$$

converting trajectory dynamics into the propagation of a scalar wavefront. In the tropical (max-plus) setting, this becomes the Bellman equation of minimax optimal control (Fleming and McEneaney 2000; Kolokoltsov and Maslov 1997). We show that the ESPF predict–update recursion satisfies precisely this tropical Hamilton–Jacobi equation, with the surprisal field playing the role of Hamilton's principal function and the PCRB playing the role of the minimum action per update cycle.

The zero-temperature limit of the classical Hamilton–Jacobi equation — obtained by passing from the log-sum-exp (probabilistic) aggregation to the max (possibilistic) aggregation — recovers this framework exactly. Bayesian inference corresponds to finite epistemic temperature; possibilistic inference is the zero-temperature ground state. This makes precise a connection only informally sketched in earlier work (Jah 2025b).

The zero-temperature limit derived here is also consistent with the Hölder-mean continuum developed in prior work (Jah 2025b): finite-temperature probabilistic aggregation corresponds to smoother averaging regimes, while the possibilistic ground state corresponds to the max-plus endpoint of that continuum. This connection is developed further in Section 10.

Organization. Section 2 establishes the impossibility field and surprisal field. Section 3 develops the tropical geometry of the active deformation front and the PCRB-admissible basin. Section 4 works through the scalar case in closed form, deriving the active front, basin width, and minimax medioid explicitly without probabilistic machinery. Section 5 states and proves the tropical Hamilton–Jacobi theorem for the ESPF recursion. Section 6 derives the PCRB as a minimum action principle. Section 8 establishes the contact geometry of epistemic phase space: the epistemic contact manifold, the contact Hamiltonian, and the identification of the PCRB as a contact energy floor. Section 10 establishes the zero-temperature limit connecting classical and tropical Hamilton–Jacobi, and situates the contact structure within thermodynamic geometry. Section 11 interprets the whitened minimax medioid as the geodesic attractor. Section 13 situates the results within TEAG and identifies open directions. Section 14 concludes.

© 2026 by the author(s). Distributed under a Creative Commons CC BY license.

2. The Impossibility Field and the Surprisal Field

Let H be a hypothesis space (measurable, possibly infinite or continuous) and let $\mathcal{Y}$ be an observation space. We work with the ESPF architecture of Jah and Haslett (2025a), which represents the epistemic state at time k as a discrete support set $\{\chi_k^{(i)}\}_{i=1}^N \subset H$ with associated possibility values $\{\pi_k^{(i)}\} \subset (0,1]$.

Definition 2.1 (Impossibility field)**.** *The* impossibility field *induced by the prior possibility π is*

$$\Phi_\varnothing(h) = -\log \pi(h) \in [0, \infty). \tag{2}$$

Where $\pi(h) = 1$ (full prior epistemic support), the impossibility is zero. As prior support is withdrawn ($\pi(h) \to 0$), the impossibility grows without bound. The impossibility field encodes the accumulated epistemic history of the filter: every past update that has diminished the possibility of h raises its impossibility. This is the log-admissibility coordinate of TEAG *(Jah 2026b) (the impossibility field definition), under which the nested α-cut structure of the* TEAG*object is exactly the nested sublevel-set structure of $\Phi_\varnothing$.*

Definition 2.2 (Surprisal field)**.** *Let $\Pi_e = L_e L_e^\top$ be the Minimum-Volume Enclosing Ellipsoid (MVEE) of the predicted measurement support $\{g(\chi_{k|k-1}^{(i)})\}_{i=1}^M$. Given observation $y_k \in \mathcal{Y}$, the* surprisal field *generated by the current observation is*

$$\Phi_S(h) = S_k(h) = \tfrac{1}{2}\|z_k(h)\|^2, \tag{3}$$

where $z_k(h) = L_e^{-1}(y_k - g(h))$ is the MVEE-whitened residual and g is the measurement mapping. This is identical to the geometric surprisal of Jah and Haslett (2025a). The surprisal field radiates outward from the observation y_k in the MVEE-whitened measurement space: hypotheses whose predicted measurements are near y_k have low surprisal; those far from it have high surprisal.

Remark 2.3 (No innovation covariance)**.** *The whitening matrix L_e^{-1} is derived from the MVEE of the* predicted measurement support cloud $\{g(\chi^{(i)})\}$*, not from any second-moment expectation. The ESPF does not form an innovation covariance $S_k = \mathbb{E}[(y - g(\hat{x}))(y - g(\hat{x}))^\top]$ because it has no state estimate $\hat{x}$ and no probability measure over the state. Π_e is a geometric object — the smallest ellipsoid enclosing the predicted measurement support — and L_e^{-1} is its shape-normalizing Cholesky factor. This is the possibilistic replacement for the Kalman innovation covariance: it characterizes the* spread of the support*, not the* variance of a distribution.

Remark 2.4 (Surprisal is the compatibility potential)**.** *The compatibility* $\mathrm{Comp}_k(h) = \exp(-S_k(h))$ *is the negative exponential of the surprisal field. The surprisal* is the impossibility induced by the evidence *— no separate definition is required. Compatibility approaching* 1 *corresponds to zero surprisal (hypothesis perfectly consistent with observation); compatibility approaching* 0 *corresponds to infinite surprisal (hypothesis expelled by the evidence).*

Definition 2.5 (Pointwise tropical superposition and the posterior impossibility field)**.** *The possibilistic conjunctive update*

$$\pi'(h) = \min\big(\pi(h),\ \mathrm{Comp}(h)\big) \tag{4}$$

is equivalent, in the field representation, to the pointwise tropical superposition *of the prior impossibility field and the surprisal field:*

$$\widetilde{\Phi}_\varnothing(h) = \Phi_\varnothing(h) \oplus \Phi_S(h) = \max\big(\Phi_\varnothing(h),\ \Phi_S(h)\big). \tag{5}$$

The posterior impossibility field is the upper envelope *of the prior and surprisal fields: whichever source of impossibility is stronger at h governs the posterior admissibility of that hypothesis. A hypothesis becomes more impossible if* either *its prior history has restricted it* or *the current evidence finds it surprising — and*

© 2026 by the author(s). Distributed under a Creative Commons CC BY license.

the governing field is whichever imposes the greater impossibility. This is TEAG*contraction in impossibility coordinates (Jah 2026b) (the tropical contraction proposition).*

Proof. Taking the negative logarithm of Equation (4) and applying $-\log\min(a,b) = \max(-\log a, -\log b)$ yields $\widetilde{\Phi}_\varnothing = \max(-\log\pi, -\log\mathrm{Comp}) = \max(\Phi_\varnothing, \Phi_S)$. □

Remark 2.6 (Pointwise superposition and the full tropical convolution)**.** *The operation* $\widetilde{\Phi}_\varnothing = \Phi_\varnothing \oplus \Phi_S$ *is a* pointwise *tropical superposition: both fields act on the same hypothesis h and the max-plus upper envelope is evaluated at each point independently. This is a degenerate case of the full tropical convolution* $(\Phi_1 \circledast \Phi_2)(h) = \sup_{h=h'+h''}[\Phi_1(h') + \Phi_2(h'')]$*, which applies when two fields act on complementary subspaces of H with a splitting structure. In the single-update case there is no splitting: the prior and surprisal fields both live on all of H, and the supremum over decompositions collapses to the pointwise max. The full tropical convolution is the correct operation for multi-sensor fusion, where* N_s *observations each generate a surprisal field* $\Phi_S^{(j)}$ *on a subspace of the joint hypothesis space; the posterior impossibility field is then genuinely a tropical convolution of all* $N_s + 1$ *fields. This is taken up as an open direction in Section 13.*

The surviving set at level α is the α-cut $H_\alpha = \{h : \pi'(h) \geq \alpha\} = \{h : \widetilde{\Phi}_\varnothing(h) \leq -\log\alpha\}$: the sublevel set of the posterior impossibility field below the threshold $-\log\alpha$. The surviving well is the basin of low impossibility; its shape and depth encode the complete epistemic state after the update.

3. Tropical Geometry of the Falsification Boundary

We work in the max-plus semiring $(\mathbb{R} \cup \{-\infty\}, \oplus, \otimes)$ where $a \oplus b = \max(a,b)$ and $a \otimes b = a + b$. This is the algebraic structure underlying tropical geometry in the max-plus convention (Maclagan and Sturmfels 2015).

The posterior impossibility field $\widetilde{\Phi}_\varnothing = \Phi_\varnothing \oplus \Phi_S$ is a two-term tropical polynomial over hypothesis space. Its tropical variety is the active deformation front.

Remark 3.1 (Two loci of tropical algebra in TEAG)**.** *Tropical algebra enters TEAG at two distinct and independently justified loci, which must be kept separate.*

Locus 1 — The update rule. *The conjunctive update* $\widetilde{\Phi}_\varnothing = \Phi_\varnothing \oplus \Phi_S = \max(\Phi_\varnothing, \Phi_S)$ *is a two-term tropical polynomial in impossibility coordinates. Its tropical variety is the active deformation front* $\mathcal{F}_k$ *defined below. This is the primary locus: tropical algebra governs the dynamics of epistemic contraction at every step.*

Locus 2 — The MVEE geometry. *The surprisal field* $\Phi_S(h) = \frac{1}{2}\|L_e^{-1}(y - g(h))\|^2$ *depends on the Cholesky factor* L_e *of the MVEE of the predicted measurement support cloud* $\{v_i\} = \{g(\chi^{(i)})\}$*. We conjecture, and verify by worked example, that the MVEE support vectors of* $\{v_i\}$ *coincide exactly with the tropical vertices of the tropical convex hull* $\mathrm{tconv}\{v_i\}$*:*

$$\{\textit{MVEE support vectors of } \{v_i\}\} = \{\textit{tropical vertices of } \mathrm{tconv}\{v_i\}\}. \tag{6}$$

If this conjecture holds in general, the whitening matrix L_e *— and therefore the entire surprisal field* Φ_S *— is derived from an object whose extremal structure is a tropical polynomial object in the precise sense of Maclagan–Sturmfels (Maclagan and Sturmfels 2015): the tropical convex hull* $\mathrm{tconv}\{v_i\}$ *is generated by linear tropical monomials* $\lambda_i \otimes v_i = \lambda_i + v_i$*, and its vertices satisfy the standard definition of tropical vertices.*

This gives the chain:

$$\{v_i\} \xrightarrow{\mathrm{tconv}} \textit{tropical vertices} = \textit{MVEE support vectors} \xrightarrow{\mathrm{MVEE}} L_e \xrightarrow{\textit{surprisal}} \Phi_S(h) = \tfrac{1}{2}\|L_e^{-1}(y - g(h))\|^2. \tag{7}$$

The tropical algebraic geometry connection to TEAG*therefore enters through the MVEE, not through the update rule itself. This resolves the objection that* $\Phi_\varnothing$ *and* Φ_S *are transcendental functions and therefore* $\max(\Phi_\varnothing, \Phi_S)$ *is not a tropical polynomial in the classical sense: Locus 1 (the update rule) is a tropical polynomial in impossibility*

© 2026 by the author(s). Distributed under a Creative Commons CC BY license.

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 15 April 2026 doi:10.20944/preprints202603.2110.v3

coordinates; Locus 2 (the MVEE) carries tropical polynomial structure in the measurement support cloud. Both are genuine. They operate at different levels of the architecture.

Definition 3.2 (Active deformation front as tropical variety)**.** *The* active deformation front $\mathcal{F}_k \subset H$ *at time k is the set of hypotheses where the prior impossibility field and the surprisal field achieve the maximum simultaneously:*

$$\mathcal{F}_k = \{h \in H : \Phi_\varnothing(h) = \Phi_S(h)\}. \tag{8}$$

This is the tropical variety of the posterior impossibility field $\Phi_\varnothing \oplus \Phi_S$: the locus where both source fields are simultaneously active and in exact balance. It marks where incoming evidence first matches prior impossibility and begins to deform the posterior field.

The active deformation front partitions hypothesis space into three dynamical regions:

(i) *History-governed region*: $\Phi_\varnothing(h) > \Phi_S(h)$ — the prior impossibility dominates; the hypothesis is suppressed by accumulated epistemic history, not by the current measurement.
(ii) *Evidence-governed region*: $\Phi_S(h) > \Phi_\varnothing(h)$ — the surprisal field dominates; the current observation drives the impossibility and governs the deformation.
(iii) *Active front* $\mathcal{F}_k$: $\Phi_\varnothing(h) = \Phi_S(h)$ — the epistemic Lagrange point, where accumulated history and current evidence exert exactly equal impossibility.

Remark 3.3 (Lagrange point interpretation)**.** *In the restricted three-body problem, the L_1 Lagrange point between two masses M_1 and M_2 is the equilibrium where their gravitational forces balance. The active deformation front $\mathcal{F}_k$ is the epistemic analog: the surface where the impossibility field (generated by the accumulated history of updates) and the surprisal field (generated by the current observation) impose equal impossibility. A hypothesis on the front is in epistemic equilibrium between history and evidence; infinitesimal perturbation in either field tips it toward survival or deeper falsification pressure.*

Remark 3.4 (Two-stage falsification structure)**.** *The active deformation front $\mathcal{F}_k$ is a* necessary *condition for falsification: no hypothesis can be falsified without first being reached by the front. It is not sufficient. A hypothesis touched by $\mathcal{F}_k$ has its posterior impossibility begin to rise, but it is not falsified unless it is pushed outside the* PCRB*-admissible basin $\mathcal{A}_k = \{h : \widetilde{\Phi}_\varnothing(h) \le c_k^\star\}$, where $c_k^\star = \frac{1}{2}(r_k^\star)^2$ and $r_k^\star = r_k^- \cdot ((1 - I_k) \cdot M / M_{\text{surv}})^{1/n}$ is the* PCRB*-admissible basin radius.*

This two-stage structure is the correct formalization of Popper's criterion: $\mathcal{F}_k$ marks where falsification becomes possible; $\mathcal{B}_{\text{adm},k} = \{h : \widetilde{\Phi}_\varnothing(h) = c_k^\star\}$ marks where falsification is complete. A hypothesis with very low prior impossibility — earned through many observations — can be touched by the front without leaving the basin, because the PCRB *protects it from mild surprisal. In the wavefront language of this paper, the front propagates outward; a hypothesis is falsified at the moment the front has pushed it beyond the* PCRB *basin boundary, not merely at the moment of first contact.*

The ESPF's pruning step — retaining all $\chi^{(i)}$ for which $S_k^{(i)} \le \tau_k$ for a data-driven threshold τ_k — is a discretization of the condition $\widetilde{\Phi}_\varnothing(h) \le c_k^\star$: the filter retains hypotheses inside the PCRB-admissible basin, i.e., within the posterior sublevel set at the PCRB-governed threshold.

4. Scalar Example: Explicit Falsification Geometry

Before stating the general theorem we work through the scalar case $H = \mathbb{R}$ in full detail. This section is the concrete demonstration of wavefront propagation: two impossibility fields — prior and surprisal — each define a level-set family, and their tropical superposition produces a new level-set family whose boundary is the active deformation front. Every quantity is computable in closed form and the front geometry is fully visible: the two roots of the active front equation are the inner and outer Lagrange points of the two-field system, the surviving well width is the gap between them, and the minimax medioid is the realized support point that minimizes the maximum distance to all other

© 2026 by the author(s). Distributed under a Creative Commons CC BY license.

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 15 April 2026 doi:10.20944/preprints202603.2110.v3

survivors within the well — the support point most insulated from the boundary in every direction. The PCRB-admissible basin in the scalar case is the interval $[h_-, h_+]$ whose half-width is the PCRB floor; hypotheses outside this interval are falsified regardless of where the active front sits. No probabilistic machinery is used: the quadratic form of each field is justified by the local geometry of the MVEE near the surviving well center — in the scalar case the minimum-volume enclosing interval of a bounded support is a quadratic form centered at the interval midpoint, so the quadratic impossibility field is the natural second-order approximation to any smooth impossibility field near its basin.

4.1. Setup

Let the prior impossibility field and the surprisal field each be quadratic:

$$\Phi_\varnothing(h) = \frac{(h - h_0)^2}{2\sigma_0^2}, \tag{9}$$

$$\Phi_S(h) = \frac{(h - y)^2}{2\sigma_n^2}, \tag{10}$$

where $h_0 \in \mathbb{R}$ is the surviving well center inherited from the previous update cycle (a geometric center of the support, carrying no distributional meaning); $\sigma_0 > 0$ is the *epistemic width* — the geometric half-width of the surviving interval, a shape parameter of the impossibility field; $y \in \mathbb{R}$ is the current observation; and $\sigma_n > 0$ is the sensor noise scale entering through the whitened residual $z(h) = (h - y)/\sigma_n$, so that $\Phi_S(h) = \frac{1}{2}z(h)^2$. Neither σ_0 nor σ_n is a standard deviation; both are geometric scale parameters of their respective fields.

4.2. The Active Deformation Front

The active deformation front $\mathcal{F}$ consists of solutions to $\Phi_\varnothing(h) = \Phi_S(h)$, i.e.,

$$\frac{(h - h_0)^2}{\sigma_0^2} = \frac{(h - y)^2}{\sigma_n^2}. \tag{11}$$

Taking square roots with both signs yields two roots.

Proposition 4.1 (Scalar active deformation front). *The active deformation front $\mathcal{F} = \{h_-, h_+\}$ consists of:*

$$h_- = \frac{\sigma_n h_0 + \sigma_0 y}{\sigma_n + \sigma_0}, \tag{12}$$

$$h_+ = \frac{\sigma_n h_0 - \sigma_0 y}{\sigma_n - \sigma_0} \qquad (\sigma_n \neq \sigma_0). \tag{13}$$

The inner root h_- always exists and lies strictly between h_0 and y. The outer root h_+ exists whenever $\sigma_n \neq \sigma_0$; when $\sigma_n = \sigma_0$ the outer root escapes to infinity and the evidence-deformed region is the half-line beyond h_-. These roots mark where deformation becomes active; the PCRB-admissible basin $[h_-, h_+] \cap \mathcal{A}_k$ determines which hypotheses within this interval are falsified.

Proof. From Equation (11), $\sigma_n(h - h_0) = \pm\,\sigma_0(h - y)$. The minus sign gives $h(\sigma_n + \sigma_0) = \sigma_n h_0 + \sigma_0 y$, yielding h_-. The plus sign gives $h(\sigma_n - \sigma_0) = \sigma_n h_0 - \sigma_0 y$, yielding h_+ when $\sigma_n \neq \sigma_0$. That h_- lies strictly between h_0 and y follows from the fact that $\sigma_n/(\sigma_n + \sigma_0)$ and $\sigma_0/(\sigma_n + \sigma_0)$ are both positive and sum to one. □

Remark 4.2 (The inner root as epistemic weighted centroid). *The inner root $h_- = (\sigma_n h_0 + \sigma_0 y)/(\sigma_n + \sigma_0)$ is a scale-weighted combination of the prior well center and the observation, with weights σ_n and σ_0. It carries no distributional interpretation: it is the point at which the two impossibility wavefronts first meet as the surprisal field propagates inward from y toward h_0 — the inner Lagrange point of the two-field system. As $\sigma_n \to \infty$*

© 2026 by the author(s). Distributed under a Creative Commons CC BY license.

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 15 April 2026 doi:10.20944/preprints202603.2110.v3

(arbitrarily uninformative sensor) $h_- \to h_0$: the wavefront barely penetrates the prior well. As $\sigma_0 \to \infty$ (arbitrarily wide prior support) $h_- \to y$: the wavefront sweeps all the way to the observation.

4.3. The Surviving Well and Its Width

When $\sigma_n > \sigma_0$, the surviving well is the interval $[h_-, h_+]$ where $\Phi_\varnothing(h) \le \Phi_S(h)$.

Proposition 4.3 (Scalar surviving well width)**.**

$$W = h_+ - h_- = \frac{2\sigma_n \sigma_0 \, |h_0 - y|}{\sigma_n^2 - \sigma_0^2}. \tag{14}$$

Proof. Using common denominator $\sigma_n^2 - \sigma_0^2$:

$$W = \frac{(\sigma_n h_0 - \sigma_0 y)(\sigma_n + \sigma_0) - (\sigma_n h_0 + \sigma_0 y)(\sigma_n - \sigma_0)}{\sigma_n^2 - \sigma_0^2} = \frac{2\sigma_n \sigma_0 (h_0 - y)}{\sigma_n^2 - \sigma_0^2}. \quad \square$$

The width W is the scalar PCRB floor: the minimum surviving support width consistent with the current epistemic state and observation noise scale. Three geometric facts are immediate from Equation (14): (i) W grows linearly with displacement $|h_0 - y|$ — the further the observation from the well center, the shallower the parabolic intersection angle and the wider the surviving gap; (ii) W diverges as $\sigma_n \to \sigma_0$ — matched scales produce nearly parallel parabolas and the outer root escapes to infinity; (iii) $W \to 0$ as $|h_0 - y| \to 0$ — when the observation falls at the well center the active deformation front degenerates to a point.

4.4. The Minimax Medioid in Closed Form

Proposition 4.4 (Scalar minimax center)**.** *In the scalar continuous-interval case, the minimax center of the surviving interval $[h_-, h_+]$ — the point minimizing maximum distance to all other points in the interval — is its midpoint:*

$$h_{\text{mid}} = \frac{h_+ + h_-}{2} = \frac{\sigma_n^2 h_0 - \sigma_0^2 y}{\sigma_n^2 - \sigma_0^2}. \tag{15}$$

Proof. Direct computation: $\frac{1}{2}(h_+ + h_-) = \frac{1}{2} \cdot \frac{(\sigma_n h_0 - \sigma_0 y)(\sigma_n + \sigma_0) + (\sigma_n h_0 + \sigma_0 y)(\sigma_n - \sigma_0)}{\sigma_n^2 - \sigma_0^2} = \frac{\sigma_n^2 h_0 - \sigma_0^2 y}{\sigma_n^2 - \sigma_0^2}$. $\square$

Remark 4.5 (Continuous scalar case vs. discrete medioid)**.** *The midpoint h_{mid} is the minimax solution of the* continuous *surviving interval $[h_-, h_+]$: it minimizes the maximum distance to any other point in $[h_-, h_+]$. In the continuous case it need not coincide with any realized support point $\chi^{(i)}$; it is a geometric property of the interval, not a selection from a finite cloud.*

In the operational ESPF, the surviving support is a finite *cloud of support points $\{\chi^{(i)}\} \subset [h_-, h_+]$. The minimax medioid $h^\star$ is then the member of that finite cloud minimizing maximum d_W-distance to all other cloud members — always a realized hypothesis, never an interpolated point. As the cloud density increases, the discrete medioid converges to h_{mid}. The scalar proposition establishes the closed-form limit; the discrete medioid (Section 11) is the operationally correct object.*

Remark 4.6 (The minimax center is not a posterior mean)**.** *The expression $(\sigma_n^2 h_0 - \sigma_0^2 y)/(\sigma_n^2 - \sigma_0^2)$ is not a convex combination of h_0 and y: the weights $\sigma_n^2/(\sigma_n^2 - \sigma_0^2)$ and $-\sigma_0^2/(\sigma_n^2 - \sigma_0^2)$ do not sum to one in the Bayesian sense. The minimax center carries no distributional interpretation. As $\sigma_0 \to 0$, $h_{\text{mid}} \to h_0$: the prior well collapses and no observation can displace the center. As $\sigma_n \to \infty$, $h_{\text{mid}} \to h_0$ again: an arbitrarily uninformative sensor cannot move the center from the prior well.*

© 2026 by the author(s). Distributed under a Creative Commons CC BY license.

4.5. The Zero-Temperature Limit in the Scalar Case

At finite epistemic temperature $T > 0$, the posterior impossibility field is the smooth function:

$$\widetilde{\Phi}_{\varnothing,T}(h) = -T\log\left(\exp\left(-\frac{(h-h_0)^2}{2T\sigma_0^2}\right) + \exp\left(-\frac{(h-y)^2}{2T\sigma_n^2}\right)\right). \tag{16}$$

The transition between the prior-governed and evidence-governed regions is soft for all $T > 0$. As $T \to 0$, the log-sum-exp identity gives pointwise convergence:

$$\widetilde{\Phi}_{\varnothing,T}(h) \xrightarrow{T\to 0} \max\big(\Phi_{\varnothing}(h),\, \Phi_S(h)\big) = \widetilde{\Phi}_{\varnothing}(h). \tag{17}$$

The soft boundary sharpens into the hard active deformation front $\mathcal{F} = \{h_-, h_+\}$. At $T = 0$ every hypothesis is either inside the PCRB-admissible basin or expelled from it — no intermediate epistemic state exists. This is the scalar instance of Theorem 10.1 (proved in Section 10).

5. The Tropical Hamilton–Jacobi Theorem

The classical Hamilton–Jacobi equation governs the propagation of Hamilton's principal function $S(q,t)$ — the accumulated action along optimal trajectories from a reference configuration. In the max-plus (tropical) setting, the Hamilton–Jacobi equation becomes the dynamic programming equation of minimax optimal control:

$$\frac{\partial\Phi}{\partial k} \oplus H_{\mathrm{trop}}(h,\, \nabla_h\Phi) = 0, \tag{18}$$

where $\oplus = \max$ and H_{trop} is the tropical Hamiltonian encoding the filter's dynamics (Fleming and McEneaney 2000; Kolokoltsov and Maslov 1997).

5.1. The Epistemic Lagrangian and the Principle of Least Epistemic Action

Before declaring the tropical Hamiltonian, we establish the variational structure from which it arises. In classical mechanics, the Hamiltonian is not the primitive object — it is derived from the Lagrangian $L = T - V$ via the Legendre transform. We construct the epistemic analog, grounding the dynamical picture in a principle that carries direct physical meaning: *beliefs are particles with mass and inertia; evidence deforms the potential field through which they move; inference is the motion that extremizes epistemic action.*

Epistemic potential energy. The impossibility field $\Phi_{\varnothing}(h) = -\log \pi(h)$ is the *epistemic potential energy* of hypothesis h. A belief-particle at h carries potential energy equal to its accumulated impossibility: zero where the hypothesis enjoys full prior support, growing without bound as prior support is withdrawn by evidence. Evidence deforms this landscape — raising the potential in falsified regions, deepening the surviving well — exactly as a mass distribution deforms the gravitational potential.

Epistemic kinetic energy and inertia. The *inertia* of a belief is its resistance to change. In the ESPF, this resistance is not uniform: the rate at which the epistemic state can move through H per update is bounded above by the Possibilistic Cramér–Rao Bound (PCRB, Theorem 6.2).

The natural measure of how much the epistemic state has moved in one update is the reduction in possibilistic entropy $\mathcal{E}_\pi = \int_0^1 \log V_\alpha \, d\alpha$: this is the integrated log-volume of the admissible support, and a reduction in $\mathcal{E}_\pi$ is precisely what it means for the filter to have "moved" — the admissible basin has contracted, carrying more geometric certainty. This is not a metaphor: in classical mechanics, kinetic energy measures how much a system has displaced from its initial configuration; in the ESPF, entropy reduction $\Delta\mathcal{E}_\pi$ measures how much the admissible geometry has contracted from its prior configuration. The analogy is one of functional role in a Lagrangian system, not of physical identity. Define the *epistemic kinetic energy* at update step k as

$$T_k \;=\; \Delta\mathcal{E}_{\pi,k} \;=\; \mathcal{E}_{\pi,k-1} - \mathcal{E}_{\pi,k}, \tag{19}$$

© 2026 by the author(s). Distributed under a Creative Commons CC BY license.

the reduction in possibilistic entropy achieved by the k-th observation. The PCRB is the statement that

$$T_k \;\leq\; I_k, \tag{20}$$

where I_k is the Choquet information content of the observation (Theorem 6.2). No evidence, however precise, can accelerate a belief-particle beyond this bound. This is epistemic inertia made quantitative: the mass of a belief is inversely proportional to how informative a single observation can be about it. In the cognitive mechanics of Section 9, $I_k = 1 - e^{-\bar{S}_k}$ plays the role of geometric contraction control, while $\bar{S}_k$ itself is the epistemic action: the two scalars are related by a monotone transform and govern distinct layers of the architecture.
The epistemic Lagrangian and action.

Definition 5.1 (Epistemic Lagrangian)**.** *The* epistemic Lagrangian *at update step k is*

$$L_k(h, \dot{h}) \;=\; T_k \;-\; V(h) \;=\; \Delta\mathcal{E}_{\pi,k} \;-\; \Phi_{\varnothing}(h), \tag{21}$$

where h is the current epistemic state (hypothesis under evaluation), $\dot{h}$ denotes the direction of epistemic motion encoded in T_k, and $V(h) = \Phi_{\varnothing}(h)$ is the epistemic potential energy. The epistemic action *accumulated over an inference trajectory of K updates is*

$$\mathcal{S} \;=\; \sum_{k=1}^{K} L_k. \tag{22}$$

Proposition 5.2 (Principle of least epistemic action)**.** *The actual inference trajectory of the ESPF — the sequence of epistemic states produced by the predict–update recursion — extremizes the epistemic action $\mathcal{S}$. Specifically, among all trajectories consistent with the PCRB constraint $T_k \leq I_k$, the ESPF selects the trajectory of minimum accumulated action: the path that neither overclaims certainty nor wastes the informational content of each observation.*

Proof. The PCRB (Theorem 6.2) bounds the entropy reduction T_k from above by I_k. The ESPF update sets $T_k = I_k$ exactly at the PCRB-admissible basin boundary $\mathcal{B}_{\mathrm{adm},k}$ — the locus of sufficient falsification — and $T_k < I_k$ in the surviving interior: no additional impossibility is assigned beyond what the evidence demands (the min-rule of possibilistic updating, Theorem 2.5). This is the discrete variational condition: $\delta\mathcal{S} = 0$ subject to the PCRB constraint, with equality achieved at the admissibility boundary $\mathcal{B}_{\mathrm{adm},k}$, not merely at the active deformation front $\mathcal{F}_k$ (Theorem 3.2). □

From Lagrangian to Hamiltonian. The Legendre transform of L_k with respect to the epistemic velocity $\dot{h}$ yields the Hamiltonian. The conjugate momentum to h is

$$p \;=\; \frac{\partial L_k}{\partial \dot{h}} \;=\; \nabla_h \Phi_{\varnothing}, \tag{23}$$

the gradient of the impossibility field — the direction and rate at which the potential rises from h. The Legendre transform then gives

$$H_{\mathrm{trop}}(h, p) \;=\; p \cdot \dot{h} - L_k \;=\; \Phi_{S_k}(h), \tag{24}$$

recovering the ESPF tropical Hamiltonian as the surprisal field (Definition 5.3 below). The momentum-independence of H_{trop} — noted in Theorem 5.4 — is therefore not an accident of construction but a structural consequence of the Lagrangian: the measurement step is a *pure potential interaction*, acting on the position of each hypothesis in H without coupling to its momentum. Epistemic kinematics (how beliefs move under dynamics) lives in the predict step; epistemic potential dynamics (how evidence reshapes the landscape) lives in the update step. The separation is exact.

We now declare the tropical Hamiltonian of the ESPF explicitly.

© 2026 by the author(s). Distributed under a Creative Commons CC BY license.

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 15 April 2026 doi:10.20944/preprints202603.2110.v3

Definition 5.3 (ESPF tropical Hamiltonian)**.** *The* ESPF tropical Hamiltonian *is the map $H_{\text{trop}} : H \times T^*H \to \mathbb{R}$ defined by*

$$H_{\text{trop}}(h, p) = \Phi_{Sk}(h) = \tfrac{1}{2}\|z_k(h)\|^2, \tag{25}$$

where $p = \nabla_h \Phi_\varnothing$ is the cotangent vector (the gradient of the impossibility field, playing the role of momentum), and $z_k(h)$ is the whitened innovation residual at time k. The Hamiltonian is independent of p: the epistemic action generated by a measurement depends only on the hypothesis's position in hypothesis space relative to the observation, not on its prior impossibility gradient. This is the kinematic *structure of the ESPF: evidence acts on position, not momentum.*

Remark 5.4 (Why the Hamiltonian is momentum-independent)**.** *In classical mechanics, a momentum-independent Hamiltonian $H = V(q)$ corresponds to a purely potential system — one where the dynamics are driven entirely by a potential field with no kinetic term. The ESPF tropical Hamiltonian $H_{\text{trop}}(h, p) = \Phi_S(h)$ is exactly this: evidence introduces a pure potential (the surprisal field) that acts on the position of each hypothesis in H, with no dependence on how fast the impossibility field is changing. The kinetic term — the transport of the impossibility wavefront under the dynamics f_k — enters only through the predict step, not through the measurement Hamiltonian. The two steps are therefore* structurally separated*: predict carries all kinematics (how hypotheses move through H under the dynamics), and update carries all potential dynamics (how evidence deforms the impossibility landscape). This separation is not a computational convenience — it is the intrinsic structure of a pure potential system in max-plus algebra, where the Hamiltonian has no momentum dependence and the Lax–Oleinik operator reduces to a pointwise max (Fleming and Soner 2006; McEneaney 2006).*

Theorem 5.5 (ESPF as tropical Hamilton–Jacobi solver / Lax–Oleinik recursion)**.** *With tropical Hamiltonian $H_{\text{trop}}(h, p) = \Phi_{Sk}(h)$ (Theorem 5.3), the ESPF predict–update recursion satisfies the discrete tropical Hamilton–Jacobi equation*

$$\widetilde{\Phi}_{\varnothing,k}(h) = \Phi_{\varnothing k|k-1}(h) \oplus H_{\text{trop}}(h, \nabla_h \Phi_{\varnothing k|k-1}) \tag{26}$$

on the knowledge manifold at each update step k. Equivalently, the update step is the Lax–Oleinik operator *of max-plus optimal control (Fleming and McEneaney 2000; Fleming and Soner 2006; McEneaney 2006; Kolokoltsov and Maslov 1997): the one-step solution operator of the tropical Hamilton–Jacobi equation, mapping the predicted impossibility field forward under the action of the measurement Hamiltonian $H_{\text{trop}} = \Phi_{Sk}$. Specifically:*

(i) Predict step (free Hamiltonian evolution): *The propagation of the impossibility field under the dynamics f_k is*

$$\Phi_{\varnothing k|k-1}(h) = \Phi_{\varnothing k-1}(f_k^{-1}(h)), \tag{27}$$

corresponding to free wavefront transport along the dynamical flow — the zero-Hamiltonian characteristic equation $\dot{h} = f_k(h)$, $\dot{\Phi}_\varnothing = 0$. The Smolyak support-point structure provides the discretization of this transport. Since no evidence is applied, $\bar{S}_k = 0$ at the predict step and the contact dissipation term vanishes: the predict step has symplectic character (Theorem 8.6), with Liouville's theorem holding on the $\{\Phi = \text{const}\}$ level sets.

(ii) Update step (wavefront collision): *Substituting Theorem 5.3 into Equation (26), and using the momentum-independence of H_{trop}:*

$$\widetilde{\Phi}_{\varnothing,k}(h) = \Phi_{\varnothing k|k-1}(h) \oplus \Phi_{Sk}(h) = \max\big(\Phi_{\varnothing k|k-1}(h),\ S_k(h)\big), \tag{28}$$

the pointwise tropical superposition of the predicted impossibility field and the surprisal field. This is the tropical Hamilton–Jacobi update: the posterior impossibility wavefront advances by taking the max-plus of the prior impossibility wavefront and the observation-generated Hamiltonian action.

(iii) Regeneration (wavefront reset): *The MVEE regeneration step redraws the discrete support around the surviving well, restoring the volumetric faithfulness invariant and resetting the impossibility field*

© 2026 by the author(s). Distributed under a Creative Commons CC BY license.

to zero over the regenerated support for the next predict step. This is the boundary condition reset of the tropical Hamilton–Jacobi recursion: $\Phi_\varnothing \leftarrow 0$ *on the regenerated cloud.*

Proof. *Predict step.* The impossibility field $\Phi_\varnothing = -\log \pi$ is transported by the push-forward of f_k: each support point $\chi^{(i)}$ propagates to $f_k(\chi^{(i)})$ with its possibility value unchanged (no measurement has occurred). In the continuum limit this is the first-order transport equation $\partial_k \Phi_\varnothing + (\nabla_h \Phi_\varnothing) \cdot f_k(h) = 0$, which is the characteristic equation of the tropical Hamilton–Jacobi PDE Equation (18) with $H_{\text{trop}} = 0$ (free evolution).

Update step. By Theorem 5.3, $H_{\text{trop}}(h,p) = \Phi_{Sk}(h)$ independently of p. Substituting into Equation (26):

$$\widetilde{\Phi}_{\varnothing,k}(h) = \Phi_{\varnothing k|k-1}(h) \oplus \Phi_{Sk}(h) = \max\big(\Phi_{\varnothing k|k-1}(h),\ \Phi_{Sk}(h)\big).$$

This is exactly the possibilistic conjunctive update $\pi'(h) = \min(\pi(h), \mathrm{Comp}(h))$ rewritten in potential coordinates via Theorem 2.5. The equivalence is exact, not approximate: the momentum-independence of H_{trop} is precisely what makes the tropical HJ update reduce to a pointwise max rather than a more complex supremum.

Regeneration. After regeneration, all support points are assigned possibility 1, so $\Phi_\varnothing = -\log 1 = 0$ over the regenerated support. This is the zero initial condition for the next predict-step transport, consistent with the boundary condition structure of the tropical Hamilton–Jacobi recursion. □

Remark 5.6 (The predict step as epistemic inertia)**.** *In the gravitational analogy, free Hamiltonian evolution corresponds to ballistic motion: a body follows its natural trajectory between gravitational interactions. The ESPF predict step propagates each support point under the dynamics without altering its possibility value — epistemic inertia. The impossibility field translates through hypothesis space carrying its shape intact, until the next measurement generates a new surprisal wavefront that collides with it.*

Proposition 5.7 (TEAG axioms force the tropical Hamilton–Jacobi structure)**.** *The tropical Hamilton–Jacobi structure of the ESPF recursion is not introduced as a modeling choice. It is forced by the* TEAG*axioms (Jah 2026b) and the log-admissibility transformation alone. Specifically:*

(i) Popperian contraction forces max-plus. *The* TEAG*contraction axiom requires* $C(\pi, y)(h) \le \pi(h)$ *for all h: evidence can only reduce admissibility. Under the log-admissibility transformation* $\Phi = -\log \pi$*, this becomes* $\Phi^+(h) \ge \Phi^-(h)$*: the impossibility field is non-decreasing under evidence. The unique binary operation on* $[0,\infty)$ *that is (a) non-decreasing in both arguments, (b) associative, and (c) reduces to the identity when one argument is zero (no evidence) is the max operation. The conjunctive update* $\min(\pi, \kappa)$ *maps under* $\Phi = -\log$ *to* $\max(\Phi^-, \psi)$ *exactly — no other operation is consistent with Popperian monotonicity in log-admissibility coordinates.*

(ii) Momentum independence is forced by the evidence structure. *The surprisal field* $\Phi_S(h) = \frac{1}{2}\|L_e^{-1}(y - g(h))\|^2$ *depends only on the hypothesis position h in measurement space, not on how the impossibility field* $\Phi_\varnothing$ *is changing at h. A Hamiltonian that depends on momentum* $p = \nabla_h \Phi_\varnothing$ *would mean that the falsification force on a hypothesis depends on how steep the prior impossibility gradient is at that point — i.e., that evidence acts differently on hypotheses depending on their prior gradient rather than their prior value. This would violate the* TEAG*evidence-referencing condition (Axiom A5 of (Jah 2026b)): survivor selection must depend only on innovation geometry, not on prior structure. Therefore* $H_{\text{trop}}(h,p) = \Phi_S(h)$*: the Hamiltonian is momentum-independent by axiomatic necessity.*

(iii) Momentum independence forces the Lax–Oleinik reduction. *For a momentum-independent tropical Hamiltonian* $H_{\text{trop}}(h,p) = V(h)$*, the tropical Hamilton–Jacobi equation* $\Phi^+(h) = \Phi^-(h) \oplus H_{\text{trop}}(h, \nabla_h \Phi^-)$ *reduces to the pointwise max* $\Phi^+(h) = \max(\Phi^-(h), V(h))$ *(Fleming and Soner 2006; McEneaney 2006). This is exactly the ESPF update rule with* $V = \Phi_S$*. The Lax–Oleinik operator is not chosen; it is the unique one-step solution operator of the tropical HJ equation for this class of Hamiltonians.*

© 2026 by the author(s). Distributed under a Creative Commons CC BY license.

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 15 April 2026 doi:10.20944/preprints202603.2110.v3

Therefore: given the TEAG*axioms and the log-admissibility transformation, the ESPF update step* must *satisfy the tropical Hamilton–Jacobi equation with Hamiltonian* $H_{\text{trop}} = \Phi_S$*. No alternative update structure is consistent with the axioms. The tropical HJ structure is not a mathematical framework imported to describe the ESPF — it is the structure the ESPF axiomatically enacts.*

Proposition 5.8 (Epistemic Wavefront Limit)**.** *Under the* TEAG*axioms and the log-admissibility transformation* $\Phi = -\log \pi$*, the evolution of* α*-cut level sets under sequential evidence induces a front propagation whose continuous-time limit satisfies the tropical Hamilton–Jacobi equation*

$$\frac{\partial \Phi}{\partial t}(h,t) \oplus \Phi_S(h,t) = 0, \tag{29}$$

where $\oplus = \max$ *and* $\Phi_S(h,t)$ *is the surprisal field at time* t*. The unique viscosity solution of Equation (29) consistent with Popperian monotonicity is the Lax–Oleinik semigroup* $\Phi(h,t) = \Phi_0(h) \oplus \sup_{0 \le s \le t} \Phi_S(h,s)$*, which at each discrete update step reduces exactly to the ESPF update* $\Phi^+ = \max(\Phi^-, \Phi_S)$*.*

Remark 5.9 (Three layers of the derivation)**.** *The connection between TEAG contraction and the tropical Hamilton–Jacobi equation passes through three distinct and separately justified layers. Blurring these layers is the source of most apparent objections to the wavefront picture; keeping them separate makes the derivation both precise and transparent.*

Layer 1 — Discrete algebraic identity. *The* TEAG*conjunctive update in log-admissibility coordinates is* $\Phi^+(h) = \max(\Phi^-(h), \psi_k(h))$ *for each hypothesis* h *at each discrete step* k*. This is an exact algebraic identity following from* $-\log \min(a,b) = \max(-\log a, -\log b)$*. No limit, approximation, or geometric interpretation is required. This layer is fully proved.*

Layer 2 — Geometric level-set interpretation. *Each* α*-cut* $H_\alpha = \{h : \Phi(h) \le -\log \alpha\}$ *is a sublevel set of* Φ*. Under the discrete update, the boundary* ∂H_α *moves inward at every* h *where* $\psi_k(h) > -\log \alpha$ *and is stationary elsewhere. This is level-set motion driven by the surprisal field: the boundary of the admissible set propagates exactly where the current evidence is strong enough to falsify. This layer is a geometric restatement of Layer 1 — no new mathematics, only a change of perspective from pointwise values to level-set boundaries.*

Layer 3 — Continuous limit and PDE. *In the limit of continuously arriving evidence with surprisal density* $\Phi_S(h,t)$*, the discrete level-set motion of Layer 2 induces a front propagation governed by Equation (29). This is not a modeling choice: the max-plus structure of Layer 1 forces the Hamiltonian to be the supremum of the driving field, and the momentum independence of the surprisal (established in Proposition 5.7) forces the Lax–Oleinik reduction. The viscosity solution theory of max-plus Hamilton–Jacobi equations (Fleming–McEneaney 2000; McEneaney 2006) guarantees that Equation (29) has a unique solution in this class, and that solution is the Lax–Oleinik semigroup stated in Proposition 5.8. The discrete ESPF update is the one-step evaluation of that semigroup.*

The wavefront formulation is therefore a mathematical level-set description of admissible-set evolution under max-plus contraction. The three layers are logically ordered: Layer 1 is prior to Layer 2, Layer 2 is prior to Layer 3. Each layer is independently justified. The governing equation in Layer 3 follows from the algebraic structure of Layer 1 via the geometric restatement of Layer 2 and the viscosity solution theory of max-plus HJ equations — it does not depend on physical intuition, gravitational analogy, or any claim about continuous inference processes in nature.

The polynomial objection and its resolution. *A natural objection is that* $\Phi_\varnothing = -\log \pi$ *and* $\Phi_S = \frac{1}{2}\|z(h)\|^2$ *are transcendental functions, not polynomials, so* $\max(\Phi_\varnothing, \Phi_S)$ *is not a tropical polynomial in the classical sense of Maclagan–Sturmfels (Maclagan and Sturmfels 2015). This objection is correct as stated for Locus 1 of the tropical structure (the update rule). Theorem 3.1 identifies the resolution: tropical polynomial structure in the classical sense enters* TEAG*at Locus 2 — the MVEE geometry. The surprisal field* Φ_S *is normalized by* L_e*, the Cholesky factor of the MVEE of the predicted measurement support cloud* $\{v_i\} = \{g(\chi^{(i)})\}$*. The support vectors of this MVEE coincide with the tropical vertices of* $\text{tconv}\{v_i\}$ *(see Equation (6)) — a genuine tropical polynomial object in measurement space. If the conjecture holds in general,* L_e *is derived from a tropical convex hull, and the surprisal field inherits genuine tropical algebraic geometry content through*

© 2026 by the author(s). Distributed under a Creative Commons CC BY license.

this chain. The tropical structure at Locus 1 (the max-plus update) and Locus 2 (the MVEE) are both real; they operate at different levels of the architecture and are not in conflict.

6. The PCRB as a Minimum Action Principle

In classical mechanics, the principle of least action states that the physical trajectory extremizes the action functional $S = \int L\,dt$. The Hamilton–Jacobi function $S(q,t)$ evaluated along optimal trajectories gives the minimum action accumulated between two configurations.

The PCRB (Jah 2025b) bounds the minimum volume of the posterior support set:

$$\log\det(\mathrm{MVEE}(\mathrm{posterior})) \geq \mathrm{PCRB}_k, \tag{30}$$

where PCRB_k is determined by the Choquet integral of per-hypothesis surprisal over the possibility-weighted surviving support. We now make this variational structure explicit.

Definition 6.1 (Choquet surprisal functional)**.** *Let $\{\chi^{(i)}\}_{i=1}^N$ be the predicted support with prior possibility field $\pi_{k|k-1}$, and let $q_k^{(i)} = \|L_e^{-1}(y_k - g(\chi^{(i)}))\|^2$ be the whitened squared innovation of hypothesis i, where $\Pi_e = L_e L_e^\top$ is the MVEE of the predicted measurement support. The* aggregate epistemic surprisal *is the Choquet integral of the per-hypothesis surprisal $\frac{1}{2}q_k^{(i)}$ with respect to the prior possibility capacity $\pi_{k|k-1}$:*

$$\bar{S}_k \;=\; (\mathrm{C})\!\int \tfrac{1}{2}q_k^{(i)}\,d\pi_{k|k-1} \;=\; \sup_{i=1,\dots,N} \min\!\Big(\tfrac{1}{2}q_k^{(i)},\; \pi_{k|k-1}^{(i)}\Big), \tag{31}$$

where the second equality is the standard reduction of the Choquet integral under a possibility measure (Jah 2025b). The possibilistic information content *of observation y_k is $I_k = 1 - e^{-\bar{S}_k} \in [0,1)$ (Jah 2026b). The Choquet formulation has a precise epistemic interpretation: $\bar{S}_k$ is the highest level t at which there exists a hypothesis that is simultaneously highly credible ($\pi_{k|k-1}^{(i)} \geq t$) and genuinely surprised ($\frac{1}{2}q_k^{(i)} \geq t$). A measurement that surprises only low-prior-possibility hypotheses yields small I_k; one that surprises the most credible hypothesis yields large I_k.*

Proposition 6.2 (PCRB as variational minimum action)**.** *The PCRB bounds the possibilistic entropy reduction per update:*

$$E_{\pi,k|k} \;\geq\; E_{\pi,k|k-1} + \tfrac{n}{2}\log(1 - I_k), \tag{32}$$

where $E_\pi = \int_0^1 \log V_\alpha\,d\alpha$ is the possibilistic entropy (Jah 2026b, 2025b). Equivalently, the PCRB is the infimum of the Choquet surprisal functional over all possibility fields consistent with the prior impossibility field:

$$\mathrm{PCRB}_k \;=\; \inf_{\pi' : \pi'(h) \leq \pi(h)\ \forall h} \bar{S}_k[\pi']. \tag{33}$$

No measurement can drive the Choquet surprisal functional below this infimum: the surviving impossibility well has a minimum depth — a minimum action per update — that no single surprisal wavefront can exceed.

The variational form Equation (33) makes precise what "minimum action" means in the possibilistic setting: it is the infimum of the possibility-weighted integral of the surprisal field, minimized over all posterior fields respecting Popperian monotonicity $\pi'(h) \leq \pi(h)$ — evidence can only reduce possibility, never increase it. The PCRB is not an analogy to minimum action; it *is* minimum action, instantiated through the Choquet integral rather than the Lebesgue integral of classical mechanics.

The gravitational analog is the sphere of influence — invoked here as structural equivalence of governing equations, not physical identity. The SOI of a central body of mass M at distance r from a perturbing body of mass M' is approximately $r_{\mathrm{SOI}} \approx r(M/M')^{2/5}$, the radius within which the central body's gravitational potential dominates. In the language of TEAG, the PCRB is an equipotential surface of the impossibility field $\Phi_\varnothing$ — an *epistemic geoid* tracing the boundary of admissible contraction per observation, analogous to a geoid in gravitational theory (Jah 2026b). Within this geoid, the

© 2026 by the author(s). Distributed under a Creative Commons CC BY license.

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 15 April 2026 doi:10.20944/preprints202603.2110.v3

prior impossibility field dominates over any single surprisal wavefront; hypotheses inside remain epistemically bound regardless of the current observation's specificity.

The PCRB also constrains the wavefront propagation speed. The surprisal wavefront cannot advance into the impossibility well faster than the PCRB allows per update: Equation (32) bounds the entropy reduction $E_{\pi,k|k} - E_{\pi,k|k-1}$ from below by a quantity determined entirely by I_k. When the filter's support-point controller detects that actual contraction is approaching the PCRB floor, it triggers expansion — regenerating a larger support cloud — to maintain volumetric faithfulness. The Epistemic Width Monitor (EWM) is thus the wavefront speed sensor, and the sigma controller is the wavefront speed regulator. In the ESPF implementation, this is operationalized via asymmetric rate limits $r^+ = 1.15$ (fast expansion) and $r^- = 0.97$ (slow contraction), encoding the epistemic asymmetry: the filter is quick to embrace ignorance and slow to assert certainty (Jah 2026b; Jah and Haslett 2025a).

7. Epistemic Geometry of Discontinuous and Multimodal Support

The entropy functional $V(\pi_k)$ of Section 6 and the contact geometry of Section 8 were developed under an implicit assumption: that the α-cut H_α is connected. This section removes that assumption. The impossibility field $\Phi_\varnothing$ is defined on a potentially disconnected manifold — multiple islands of admissibility separated by regions of absolute impossibility. We show that TEAG does not break down in this setting. The mathematics holds on each connected component exactly as general relativity holds away from singularities in a universe that contains black holes. The Choquet integral provides the correct aggregation across components, and two structural theorems — epistemic idempotency and epistemic rank — govern how collective observation deforms the field.

Definition 7.1 (Connected components of the α-cut)**.** *Let $H_\alpha = \{h \in H : \pi(h) \geq \alpha\}$ be the α-cut at level α. Suppose H_α is disconnected with $K(\alpha)$ connected components:*

$$H_\alpha = C_1(\alpha) \cup C_2(\alpha) \cup \cdots \cup C_{K(\alpha)}(\alpha), \qquad C_i(\alpha) \cap C_j(\alpha) = \varnothing \text{ for } i \neq j.$$

The regions between components — the interior of the MVEE of H_α not covered by any $C_i(\alpha)$ — are regions of absolute impossibility: *$\pi(h) = 0$ exactly, $\Phi_\varnothing(h) = \infty$ in every coordinate chart. These regions are not part of the epistemic manifold; they are expelled from H_α by definition.*

Remark 7.2 (The discontinuous universe)**.** *The situation is structurally identical to general relativity in a universe containing multiple black holes. The Einstein field equations hold on the smooth manifold $\mathcal{M}$ minus its singularities. Singularities are not interior points of $\mathcal{M}$ — they are boundaries or endpoints of geodesics, expelled from the domain of the equations. The theory holds everywhere it is applied. In* TEAG*, the absolutely impossible interior regions play the same role: they are not points of H_α. The impossibility field, the John contact geometry, and the Choquet entropy functional all apply on each connected component $C_i(\alpha)$, which is smooth wherever defined. Disjointness is not a failure of the geometry. It is the epistemic analogue of a universe with multiple black holes, and the mathematics is no more troubled by it than general relativity is.*

Definition 7.3 (Component MVEE and log-volume)**.** *For each connected component $C_i(\alpha)$, let Q_i^α denote the MVEE shape matrix of the embedding cloud $\{e_h : h \in C_i(\alpha)\}$, computed by the Todd–Yildirim algorithm restricted to that component. Define the component log-volume:*

$$V_i(\alpha) = \tfrac{1}{2} \log \det Q_i^\alpha.$$

This is the true epistemic log-volume of island i at level α. It excludes the impossible interior — the sea between islands — by construction.

© 2026 by the author(s). Distributed under a Creative Commons CC BY license.

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 15 April 2026 doi:10.20944/preprints202603.2110.v3

Definition 7.4 (Component possibility weights and superadditive capacity). *Define the possibility weight of component $C_i(\alpha)$ as*

$$\mu_\alpha(C_i) = \sup_{h \in C_i(\alpha)} \pi(h).$$

For pairs of components, define the epistemic proximity weight

$$\omega_{ij}(\alpha) = \exp\Big(-\|c_i(\alpha) - c_j(\alpha)\|_W^2/\sigma^2\Big),$$

where $c_i(\alpha)$ is the minimax medioid of $C_i(\alpha)$ in the MVEE-whitened metric $\|\cdot\|_W$ and σ is a scale parameter. The 2-additive fuzzy capacity on components is

$$\mu_\alpha(C_i \cup C_j) = \mu_\alpha(C_i) + \mu_\alpha(C_j) + v_{ij}\,\omega_{ij}(\alpha), \quad v_{ij} \geq 0.$$

This capacity is superadditive: the collective ignorance of two epistemically proximate islands exceeds the sum of their individual ignorances, because proximity across impossible interior is itself epistemic information. The algebraic structure is identical to the superadditive capacity μ_{ing} governing the GaiaGraph ingestion pipeline, now applied to geometric components of the support.

We now state the two structural theorems governing collective observation.

Lemma 7.5 (Epistemic idempotency of identical evidence). *Let $\Phi_\varnothing$ be the impossibility field at step k. Let the same observation y_k be received n times, each with the same measurement geometry and therefore the same MVEE whitening, generating identical surprisal fields $\Phi_S^{(1)} = \Phi_S^{(2)} = \cdots = \Phi_S^{(n)} = \Phi_S$. Then the posterior impossibility field after n identical observations equals the posterior after a single observation:*

$$\widetilde{\Phi}_\varnothing = \max\big(\Phi_\varnothing,\, \Phi_S^{(1)}, \ldots, \Phi_S^{(n)}\big) = \max\big(\Phi_\varnothing,\, \Phi_S\big). \tag{34}$$

Proof. Immediate from the idempotency of the max operation: $\max(a, a, \ldots, a) = a$. After the first observation the impossibility field is deformed to $\max(\Phi_\varnothing, \Phi_S)$. Each subsequent identical observation generates the same surprisal field Φ_S. The tropical superposition $\max\big(\max(\Phi_\varnothing, \Phi_S),\, \Phi_S\big) = \max(\Phi_\varnothing, \Phi_S)$ is unchanged. The active deformation front has already settled at its post-first-observation position. No further geometric deformation is possible. □

Remark 7.6 (What idempotency means physically). *This lemma is categorically different from the probabilistic case, where n identical measurements narrow the posterior by $\sqrt{n}$. In TEAG, the impossibility field records geometric deformation events, not counts of observations. Repetition leaves no geometric trace. A town of n observers all reporting the same measurement deforms the epistemic field exactly as much as the first observer alone. The field is quick to register a new deformation and entirely indifferent to its repetition. This is not a limitation — it is the correct possibilistic semantics: possibility measures what is admissible, not how frequently something has been observed.*

The consequence for collective inference is precise. If n observers are identical in the sense that everything about them is identical — same prior, same measurement geometry, same observation — their combined epistemic contribution is exactly that of any single one. Possibility is not additive. Repetition is not amplification.

Theorem 7.7 (Epistemic rank of collective observation). *Let n observers report observations generating surprisal fields $\Phi_S^{(1)}, \ldots, \Phi_S^{(n)}$ with the same measurement geometry. Define the* epistemic rank *$r = n - m$, where m is the number of observations whose surprisal fields are geometrically redundant — meaning they contribute no deformation to the upper envelope beyond what the remaining r fields already produce. Then*

$$\widetilde{\Phi}_\varnothing = \max\big(\Phi_\varnothing,\, \Phi_S^{(i_1)}, \Phi_S^{(i_2)}, \ldots, \Phi_S^{(i_r)}\big), \tag{35}$$

where $\{i_1, \ldots, i_r\}$ is any maximal geometrically independent subset of the n observations.

© 2026 by the author(s). Distributed under a Creative Commons CC BY license.

Corollaries.

(i) *As $m \to n-1$, $r \to 1$ and we recover Theorem 7.5: the collective deformation collapses to that of a single informant.*

(ii) *As $m \to 0$, $r = n$ and every observation contributes a geometrically distinct deformation.*

(iii) *The collective epistemic deformation of n observers equals the deformation of its r geometrically independent informants. The remaining m are redundant not socially but geometrically: their surprisal fields lie beneath the upper envelope of the independent set.*

Proof. The tropical superposition $\max(\Phi_S^{(1)}, \dots, \Phi_S^{(n)})$ depends only on the pointwise upper envelope of the collection. Observation i is geometrically redundant if $\Phi_S^{(i)} \le \max_{j \ne i} \Phi_S^{(j)}$ pointwise on the surviving support H_{surv} — its surprisal field is everywhere dominated by the envelope of the others. The maximal independent subset $\{i_1, \dots, i_r\}$ is the minimal collection whose upper envelope equals that of the full set. Its existence follows from the finite cardinality of the observation set. Uniqueness of the envelope — not of the subset — is what matters: different maximal independent subsets produce the same upper envelope and therefore the same posterior field. □

Remark 7.8 (Epistemic rank and Dempster–Shafer theory)**.** *The epistemic rank theorem is the* TEAG *analogue of the vacuous extension property in Dempster–Shafer theory: combining a source with an identical copy adds no information. The present result is stronger and more geometric: it operates directly on the shape of the impossibility field in the MVEE-whitened metric, not on abstract belief functions.*

Definition 7.9 (Refined Choquet entropy functional)**.** *The Choquet entropy functional over disconnected support is*

$$V_{\text{Choquet}}(\pi_k) = \int_0^1 (\mathrm{C}) \int V_i(\alpha)\, d\mu_\alpha\, d\alpha, \tag{36}$$

where the inner integral is the Choquet integral of component log-volumes with respect to the superadditive capacity μ_α of Theorem 7.4. This functional excludes impossible interior regions by construction — each $V_i(\alpha)$ is the MVEE log-volume of its component alone — and aggregates across components in a way that respects epistemic proximity through the interaction weights $\omega_{ij}(\alpha)$.

Theorem 7.10 (Choquet entropy collapse in the connected limit)**.** *Let H_α have $K(\alpha)$ connected components with 2-additive superadditive capacity μ_α. In the limit $K(\alpha) \to 1$ for almost every $\alpha \in [0,1]$:*

$$V_{\text{Choquet}}(\pi_k) \;\longrightarrow\; \int_0^1 \tfrac{1}{2} \log\det Q_\alpha\, d\alpha \;=\; V(\pi_k). \tag{37}$$

Proof. When $K(\alpha) = 1$ there is a single component $C_1(\alpha) = H_\alpha$. By global normalisation, $\mu_\alpha(C_1) = \sup_{h \in H_\alpha} \pi(h) = 1$. With a single component the Choquet sum has one term:

$$V_{\text{Choquet}} = \int_0^1 V_1(\alpha) \cdot \mu_\alpha(C_1)\, d\alpha = \int_0^1 \tfrac{1}{2} \log\det Q_1^\alpha\, d\alpha.$$

Since $C_1(\alpha) = H_\alpha$, we have $Q_1^\alpha = Q_\alpha$ and the result follows. Convergence as $K(\alpha) \to 1$ holds almost everywhere because $K(\alpha)$ is non-decreasing in α — components can split but not merge as α increases — and is therefore measurable, with the single-component limit achieved at the α-values where the support is connected. □

Remark 7.11 ($V(\pi_k)$ as a special case, not a primitive)**.** *Theorem 7.10 reveals the entropy functional $V(\pi_k)$ of Section 6 not as a primitive definition but as the connected limit of a more general functional. The current functional is correct — it is the right object when the support is connected — and is now grounded as a special case rather than an assumption. The generalisation costs nothing in the connected regime and recovers the full multimodal geometry in the disconnected one.*

© 2026 by the author(s). Distributed under a Creative Commons CC BY license.

Definition 7.12 (Epistemic convexity defect)**.** *Define the epistemic convexity defect at level α:*

$$D(\alpha) = \tfrac{1}{2} \log\det Q_\alpha - (C)\int V_i(\alpha)\, d\mu_\alpha \;\geq\; 0,$$

and the integrated convexity defect:

$$D(\pi_k) = V(\pi_k) - V_{\text{Choquet}}(\pi_k) = \int_0^1 D(\alpha)\, d\alpha \;\geq\; 0. \tag{38}$$

$D(\pi_k)$ is the scalar gap between the global MVEE log-volume — the outer coastline of the epistemic archipelago, including the impossible sea between islands — and the Choquet aggregate of island volumes. It is non-negative because the global MVEE always encloses at least as much volume as any component. It vanishes when the support is connected and grows with the degree of multimodality. $D(\pi_k)$ is computable from the ESPF support cloud by comparing global and component-wise MVEE volumes, and provides a live diagnostic of epistemic fragmentation.

Proposition 7.13 (Idempotency as the vanishing-defect limit)**.** *Theorem 7.5 and Theorem 7.10 are instances of the same underlying principle: when epistemic contributions are geometrically redundant — whether because n observers share one surprisal field, or because K components share one geometry — the Choquet integral collapses to the contribution of the single non-redundant element. Redundancy in* TEAG *is geometric, not social. The convexity defect $D(\pi_k)$ is the scalar measure of how much non-redundant geometric content exists across disconnected components. When $D(\pi_k) = 0$ the support is epistemically unified. When $D(\pi_k)$ is large, the support is genuinely fragmented: multiple distinct epistemic worlds, each internally coherent, separated by absolute impossibility.*

Remark 7.14 (Connection to the PHME framework)**.** *The Choquet aggregation across components in this section is structurally identical to the Possibilistic Hierarchical Mixture of Experts (*PHME*) framework developed for GaiaGraph. Each connected component $C_i(\alpha)$ is an expert with its own MVEE geometry, its own medioid, and its own possibility weight. The superadditive capacity encodes interactions between experts — exactly as μ_{ing} encodes complement relationships between scoring agents in the GaiaGraph ingestion pipeline. The unified aggregation algebra governs inference, ingestion, bridge detection, and now the multimodal support geometry of the ESPF. One algebraic law governs all.*

8. Contact Geometry of Epistemic Phase Space

The tropical Hamilton–Jacobi structure established in Section 5 places TEAG within the framework of Hamiltonian mechanics. This section makes the geometric structure precise by identifying the correct geometric primitive: epistemic phase space carries a *contact structure*, not a symplectic one. The distinction matters because it is contact geometry, not symplectic geometry, that naturally accommodates the irreversibility and one-sided dissipation that are not optional features of TEAG but axiomatic requirements.

8.1. Why Symplectic Geometry Is Not Quite Right

Classical symplectic geometry is the geometry of conservative, time-reversible Hamiltonian systems. On a $2n$-dimensional phase space (q, p), the symplectic form $\omega = dq \wedge dp$ is preserved under Hamiltonian flow: the system is volume-preserving (Liouville's theorem) and every trajectory is reversible.

TEAG violates both properties by axiom. Axiom A3 (non-resurrection) states that a falsified hypothesis cannot be restored by subsequent evidence: the flow of the impossibility field is one-sided and irreversible. Axiom A2 (Popperian contraction) states that evidence can only increase impossibility, never decrease it: the admissible basin contracts monotonically and volume is not preserved. A fully symplectic structure would require reversibility and volume conservation. TEAG has neither.

The correct geometry is one dimension higher.

© 2026 by the author(s). Distributed under a Creative Commons CC BY license.

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 15 April 2026 doi:10.20944/preprints202603.2110.v3

8.2. The Epistemic Contact Manifold

Definition 8.1 (Epistemic phase space and contact form)**.** *The* epistemic phase space *is the* $(2n+1)$*-dimensional manifold*

$$\mathcal{M}_{\text{ep}} \;=\; \big\{(h,\, p,\, \Phi) \;:\; h \in H,\; p = \nabla_h \Phi_\varnothing \in T_h^* H,\; \Phi = \Phi_\varnothing(h) \in \mathbb{R}_{\geq 0}\big\}, \tag{39}$$

where h is the hypothesis position, $p = \nabla_h \Phi_\varnothing$ is the conjugate momentum (the gradient of the impossibility field, as established in Equation (23)), and Φ is the value of the impossibility field along the trajectory — the accumulated action.

The epistemic contact form *on $\mathcal{M}_{\text{ep}}$ is*

$$\eta \;=\; d\Phi \;-\; p \cdot dh \;=\; d\Phi_\varnothing - \nabla_h \Phi_\varnothing \cdot dh. \tag{40}$$

The contact form η is the canonical contact form of classical contact geometry (Geiges 2008), with the impossibility field Φ playing the role of the action variable. It satisfies $\eta \wedge (d\eta)^n \neq 0$ everywhere on $\mathcal{M}_{\text{ep}}$ (the non-degeneracy condition for a contact structure), provided the impossibility field is smooth and the VFI condition (Axiom A4) holds to prevent support collapse.

Remark 8.2 (Why one dimension higher)**.** *The passage from symplectic to contact geometry corresponds exactly to the passage from the phase space (h, p) to the extended phase space (h, p, Φ). In classical mechanics, this extension arises when the action is included as an explicit variable — the setting of the Hamilton–Jacobi equation. In TEAG, this extension is not optional: the accumulated impossibility Φ is not a derived quantity but an intrinsic one, because Axiom A3 requires that the history of falsification is permanently encoded in the field. The impossibility field carries memory. Contact geometry is the natural framework for systems with memory and irreversibility; symplectic geometry is its reversible, memory-free limit.*

8.3. Contact Hamiltonian Flow

In contact geometry, a Hamiltonian function $\mathcal{H} : \mathcal{M}_{\text{ep}} \to \mathbb{R}$ generates a contact Hamiltonian vector field $X_{\mathcal{H}}$ determined by the conditions

$$\eta(X_{\mathcal{H}}) = -\mathcal{H}, \qquad \iota_{X_{\mathcal{H}}} d\eta = d\mathcal{H} - (R\mathcal{H})\,\eta, \tag{41}$$

where R is the Reeb vector field of η (Bravetti, Cruz, and Tapias 2017). The contact Hamiltonian equations of motion are

$$\dot{h} = \quad \nabla_p \mathcal{H}, \tag{42}$$

$$\dot{p} = -\nabla_h \mathcal{H} - p\,\partial_\Phi \mathcal{H}, \tag{43}$$

$$\dot{\Phi} = p \cdot \nabla_p \mathcal{H} - \mathcal{H}. \tag{44}$$

These reduce to the standard Hamilton equations when $\partial_\Phi \mathcal{H} = 0$ and $\mathcal{H}$ does not depend on Φ — the symplectic limit.

Proposition 8.3 (TEAG as contact Hamiltonian system)**.** *The TEAG update rule is a contact Hamiltonian system on $\mathcal{M}_{\text{ep}}$ with contact Hamiltonian*

$$\mathcal{H}_{\text{ep}}(h, p, \Phi) \;=\; \Phi_S(h) \;-\; \Phi, \tag{45}$$

where $\Phi_S(h) = \frac{1}{2}\|z(h)\|_W^2$ is the surprisal field.

© 2026 by the author(s). Distributed under a Creative Commons CC BY license.

The contact Hamiltonian equations Equation (42)–Equation (44) yield:

$$\dot{h} = 0, \tag{46}$$

$$\dot{p} = -\nabla_h \Phi_S(h) + p, \tag{47}$$

$$\dot{\Phi} = \Phi_S(h) - \Phi \;\cdot\; 0 - \mathcal{H} = \Phi - \Phi_S(h) + \Phi = \Phi - \Phi_S(h). \tag{48}$$

The fixed point of $\dot{\Phi} = 0$ is $\Phi = \Phi_S(h)$: exactly the active deformation front $\mathcal{B}_{\text{active}}$ (Theorem 3.2). The contact flow drives Φ toward $\Phi_S(h)$ when $\Phi < \Phi_S(h)$ (evidence-governed region) and leaves Φ unchanged when $\Phi \geq \Phi_S(h)$ (history-governed region), consistent with the max-plus update $\Phi^+ = \max(\Phi^-, \Phi_S)$.

Proof. Substituting $\mathcal{H}_{\text{ep}} = \Phi_S(h) - \Phi$ into Equation (42)–Equation (44): since $\mathcal{H}$ is independent of p, $\nabla_p \mathcal{H} = 0$, giving $\dot{h} = 0$ — the pure potential character of the measurement step established in Theorem 5.4. Since $\partial_\Phi \mathcal{H} = -1$, $\dot{p} = -\nabla_h \Phi_S + p \cdot (-1)(-1) = -\nabla_h \Phi_S + p$. For $\dot{\Phi}$: $p \cdot \nabla_p \mathcal{H} - \mathcal{H} = 0 - (\Phi_S - \Phi) = \Phi - \Phi_S$. At the fixed point $\dot{\Phi} = 0$: $\Phi = \Phi_S(h)$, which is the active deformation front. □

Remark 8.4 (Contact dissipation and Axiom A3)**.** *The $-p\,\partial_\Phi \mathcal{H}$ term in the contact momentum equation Equation (43) is the* contact dissipation term *— the structural feature that distinguishes contact from symplectic dynamics. For $\partial_\Phi \mathcal{H} = -1$, this term is $+p$, which damps the momentum back toward the observation geometry. This dissipation is not a modeling choice; it is the algebraic consequence of including Φ as an explicit variable in the Hamiltonian. At the level of axioms, it enforces non-resurrection (Axiom A3): once a hypothesis has been driven to $\Phi = \infty$ by evidence, the contact flow does not return it. Irreversibility is built into the contact structure.*

8.4. The PCRB as a Contact Energy Bound

In contact Hamiltonian mechanics, the contact energy of a trajectory is not conserved: it evolves according to

$$\frac{d\mathcal{H}_{\text{ep}}}{dk} \;=\; -(\partial_\Phi \mathcal{H})\,\mathcal{H} \;=\; \mathcal{H}_{\text{ep}}, \tag{49}$$

so the contact Hamiltonian grows exponentially along trajectories in the evidence-governed region. This is the geometric statement that surprisal accumulates: each observation adds to the impossibility field, and the contact energy — the gap between the current surprisal and the current impossibility — grows until the active deformation front is reached.

The PCRB is a bound on how fast this contact energy can be dissipated per update.

Proposition 8.5 (PCRB as contact energy floor)**.** *The PCRB inequality*

$$\mathcal{E}_{\pi,k|k} \;\geq\; \mathcal{E}_{\pi,k|k-1} + \tfrac{n}{2}\log(1 - I_k) \tag{50}$$

is the contact energy floor of the epistemic Hamiltonian system: it bounds the rate at which the contact Hamiltonian $\mathcal{H}_{\text{ep}}$ can drive the impossibility field toward the surprisal field per observation. The left side is the possibilistic entropy after contact flow has been applied; the right side is the pre-contact entropy shifted by the maximum admissible contact energy dissipation, determined by the Choquet information content I_k.

The ESPF achieves equality at the admissibility boundary $\mathcal{B}_{\text{adm},k}$ — the PCRB equipotential surface — where the contact energy is exactly dissipated and no further impossibility elevation is epistemically justified.

8.5. ESPF as Projection onto the Admissible Contact Manifold

The full contact Hamiltonian flow evolves (h, p, Φ) continuously. The ESPF does not integrate this flow directly — it implements a *discrete projection* onto an admissible submanifold at each step.

After each update, the ESPF regenerates the support-point cloud around the minimax medioid, resetting the impossibility field to zero over the new support. In contact geometry terms, this is a

© 2026 by the author(s). Distributed under a Creative Commons CC BY license.

projection of the contact flow back onto the submanifold $\{\Phi = 0\}$ restricted to the surviving admissible basin. The projection is not arbitrary: it selects the submanifold point that is the fixed point of the geodesic flow on the contact manifold — the medioid — and it resets the accumulated action Φ to zero while preserving the basin geometry encoded in the MVEE.

This gives the correct characterization of the full ESPF cycle:

(i) *Predict step.* Free contact Hamiltonian evolution with $\mathcal{H} = 0$: the impossibility field is transported along the dynamics without change, tracing the Reeb flow of the contact structure.

(ii) *Update step.* Contact Hamiltonian flow with $\mathcal{H}_{\text{ep}} = \Phi_S(h) - \Phi$: the impossibility field is driven toward the surprisal field, bounded by the PCRB contact energy floor.

(iii) *Regeneration step.* Projection back to $\{\Phi = 0\}$ on the surviving basin: the contact action variable is reset, and the Smolyak cloud is reinitialized at the geodesic attractor of the current well.

The full ESPF recursion is therefore a contact Hamiltonian system with discrete projection: predict and update follow contact flow, and regeneration is a projection that resets the action variable while preserving the admissible geometry.

Remark 8.6 (Symplectic geometry as the zero-information limit)**.** *The recovery of symplectic geometry in the absence of measurably informative evidence is a theorem, not merely an analogy.*

The precise statement *(Theorem A13): When $\bar{S}_k = 0$ — no hypothesis is simultaneously credible and surprised, so the active deformation front is empty — the following hold simultaneously:*

(i) The contact dissipation term $-p\,\partial_\Phi \mathcal{H} = +p$ vanishes on the surviving interior.

(ii) The action variable Φ is constant along trajectories: $\dot{\Phi} = \Phi - \Phi_S = 0$ since $\Phi_S \le \Phi$ everywhere.

(iii) The contact form $\eta = d\Phi - p \cdot dh$, restricted to the $\{\Phi = \text{const}\}$ level sets, reduces to the symplectic form $\omega = dh \wedge dp$.

(iv) Liouville's theorem holds: the contact Hamiltonian flow is volume-preserving on these level sets.

(v) The dynamics are reversible: no hypothesis is expelled, and the update is a diffeomorphism on the full support.

This is exactly the predict step *of the ESPF — free Hamiltonian transport of the impossibility field along the dynamical flow (Equation (27)), where $\dot{\Phi}_\varnothing = 0$. The predict step has symplectic character because it carries zero information: it transports the field without deforming it.*

What breaks the symmetry: *The update step introduces the surprisal field $\Phi_S(h) > 0$, which makes the contact dissipation term non-zero wherever evidence is informative. The moment $\bar{S}_k > 0$, the contact dissipation fires, the action variable Φ changes, the level-set restriction no longer yields ω, Liouville's theorem fails, and the irreversibility of Axiom A3 is geometrically enacted.*

The implication: *Symplectic geometry is not wrong for inference — it is the correct geometry for inference in the absence of information. Contact geometry is the geometry demanded by the existence of information. Every measurement that carries genuine epistemic content ($\bar{S}_k > 0$) breaks the symplectic structure. The PCRB controls how far it breaks.*

Remark 8.7 (Connection to thermodynamic contact geometry)**.** *Contact geometry is the natural geometric setting for thermodynamics: the thermodynamic phase space (E, S, V, T, P) carries a canonical contact structure $dE - T\,dS + P\,dV = 0$, and the laws of thermodynamics are contact Hamiltonian flow (Bravetti 2019). The zero-temperature limit of Section 10 — the passage from log-sum-exp (Boltzmann) to max (possibilistic) aggregation — is a contact phase transition: as $T \to 0$, the thermodynamic contact manifold degenerates to the epistemic contact manifold of TEAG. Probability theory lives in the finite-temperature contact geometry; possibility theory lives in the zero-temperature limit. The PCRB is the contact energy bound at zero temperature, and the Shannon channel capacity is its finite-temperature counterpart. This connection places the possibilistic–probabilistic duality within a single geometric framework: both are contact Hamiltonian systems, differing only in temperature.*

© 2026 by the author(s). Distributed under a Creative Commons CC BY license.

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 15 April 2026 doi:10.20944/preprints202603.2110.v3

9. Cognitive Mechanics of Admissible Inference

The contact geometry of Section 8 places the impossibility field within a Hamiltonian framework. This section makes explicit a further layer of structure latent in that framework.

The Choquet aggregate surprisal $\bar{S}_k$ does not merely quantify evidence strength; it governs the admissible propagation of the deformation wavefront. In doing so, it determines the extent to which the gradient of the impossibility field can induce cognitive acceleration on hypothesis trajectories. Thus, epistemic dynamics are fully controlled by a single scalar functional: the action of surprisal on the hypothesis manifold. When this action vanishes, no deformation wavefront propagates, cognitive acceleration is null, and the system becomes incapable of distinguishing change.

Naming and formalizing this structure yields a mechanics of belief that connects the tropical Hamilton–Jacobi recursion, the contact geometry, the PCRB bound, and the conditions under which learning can and cannot occur. The central objects — cognitive velocity, cognitive acceleration, epistemic energy, cognitive work — are derived observables of the TEAG system, not additions to it. They were already there; this section makes them visible.

9.1. Cognitive Velocity and Cognitive Acceleration

Definition 9.1 (Cognitive velocity)**.** *Let $h(\tau)$ be a trajectory on the hypothesis manifold H parameterized by epistemic time τ. The* cognitive velocity *is*

$$v(h) \;:=\; \dot{h} \;=\; \frac{dh}{d\tau}, \tag{51}$$

the rate of traversal through admissible hypothesis space.

Definition 9.2 (Cognitive acceleration)**.** *Let $\Phi_\varnothing(h) = -\log \pi(h)$ be the impossibility field and let $G(h)$ be the metric tensor on H (the MVEE-whitened geometry of Section 2). The* cognitive acceleration field *is*

$$a(h) \;:=\; -\nabla_G \Phi_\varnothing(h), \tag{52}$$

where ∇_G denotes the gradient induced by the metric G. The cognitive acceleration field governs the rate of change of the rate at which an epistemic system can traverse admissible hypothesis space under evidence.

Remark 9.3 (Cognitive acceleration as epistemic force)**.** *The cognitive acceleration field is the local direction and magnitude of epistemic recovery from falsification pressure. It points from regions of high impossibility toward regions of lower impossibility — the direction of increasing admissibility — and its magnitude is proportional to the local rate of change of the impossibility field.*

In classical mechanics, force is the negative gradient of the potential; mass determines resistance to that force; and acceleration is force divided by mass. The epistemic analogy is exact at the structural level. The impossibility field $\Phi_\varnothing$ is the epistemic potential; the metric $G(h)$ encodes epistemic inertia *— the resistance of a hypothesis to deformation under evidence; and cognitive acceleration is the ratio of epistemic force to epistemic inertia. The resulting equation of motion on hypothesis space is*

$$G(h)\,\ddot{h} \;+\; \Gamma(h,\dot{h}) \;=\; -\nabla_G \Phi_\varnothing(h), \tag{53}$$

where $\Gamma(h,\dot{h})$ collects the Christoffel terms arising from the curvature of H in the metric G.

Remark 9.4 (Epistemic equivalence principle)**.** *The equation of motion Equation (53) admits a structural analog of Einstein's equivalence principle. In general relativity, gravitational acceleration and inertial acceleration are locally indistinguishable: an agent in free fall cannot determine, from local measurements alone, whether it is in a gravitational field or an accelerating frame. In the present framework, the right-hand side $-\nabla_G\Phi_\varnothing(h)$ (cognitive acceleration from the evidence field) and the left-hand side Christoffel term $\Gamma(h,\dot{h})$ (geodesic deviation from the curvature of the epistemic manifold) both contribute to $\ddot{h}$ through the same equation. An epistemic agent*

© 2026 by the author(s). Distributed under a Creative Commons CC BY license.

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 15 April 2026 doi:10.20944/preprints202603.2110.v3

observing only its trajectory $h(\tau)$ cannot decompose $\ddot{h}$ into a field-induced component and a geometry-induced component: the two sources of acceleration are locally indistinguishable.

Crucially, this indistinguishability is not a postulate here — it is a consequence of the variational structure. In GR, Einstein elevated an empirical observation (the equality of gravitational and inertial mass) to a foundational axiom. In the epistemic setting, the indistinguishability follows necessarily from the fact that both $-\nabla_G \Phi_\varnothing$ and Γ enter $\ddot{h}$ through the same Euler–Lagrange equation derived from the Lagrangian $\mathcal{L}_c = T_c - V_c$. There is no separate postulate; the equivalence is structural.

The analogy has a precise breaking point. In GR, tidal forces (the Riemann curvature tensor) distinguish a true gravitational field from a uniformly accelerating frame at finite separation. In the epistemic setting, the equivalence breaks at deformation fronts $\mathcal{B}_{\text{active}}$, where the impossibility field is non-smooth: at these loci $\nabla_G \Phi_\varnothing$ is multi-valued and the field-induced and geometry-induced contributions become separately observable as shock-like discontinuities in the trajectory. Away from deformation fronts, the equivalence holds exactly.

There is a second, distinct form of the equivalence principle native to TEAG and not present in GR. When $\bar{S}_k = 0$ — when aggregate epistemic surprisal vanishes — the surprisal field Φ_S is everywhere dominated by the prior impossibility field $\Phi_\varnothing$, the active deformation front $\mathcal{B}_{\text{active}}$ is empty, and cognitive acceleration vanishes identically. In this regime the impossibility field evolves under free transport with no deformation: the system is incapable of registering change. The consequence is that absence of epistemic action is indistinguishable from absence of epistemic force. *A system with $\bar{S}_k = 0$ cannot determine whether it is stationary because nothing is acting on it, or because it is incapable of responding. This is not the field-vs-geometry degeneracy of the first form; it is a degeneracy in the zero set of the action functional itself. Both forms are captured by Theorem 9.15 and resolved diagnostically by $\mathcal{B}_k$.*

Remark 9.5 (No canonical momentum required)**.** *The cognitive acceleration field Equation (52) is a statement about the metric geometry of hypothesis space and does not introduce a canonical momentum variable. This is fully consistent with the tropical Hamilton–Jacobi theorem (Theorem 5.5), which proves that the TEAG Hamiltonian is momentum-independent. The absence of canonical momentum in the tropical recursion is a statement about the* reduced, evidence-driven update law*; the presence of geometric inertia encoded in $G(h)$ is a statement about the* metric structure governing admissible trajectory deformation*. These are distinct claims that coexist without contradiction. Precisely: the tropical HJ recursion is the reduced, momentum-free description of the same admissibility dynamics that the Euler–Lagrange language describes geometrically.*

9.2. Conservative Structure and Its Limits

Proposition 9.6 (Cognitive acceleration is irrotational)**.** *Wherever the impossibility field $\Phi_\varnothing$ is smooth, the cognitive acceleration field is irrotational:*

$$\nabla \times a(h) \;=\; \nabla \times (-\nabla_G \Phi_\varnothing) \;=\; 0. \tag{54}$$

Consequently, the epistemic force field is conservative wherever $\Phi_\varnothing$ is smooth: the work done by evidence between two hypotheses depends only on endpoints, not on the path taken.

Proof. The curl of any gradient field vanishes identically for smooth scalar fields. Since $a = -\nabla_G \Phi_\varnothing$ and $\Phi_\varnothing$ is smooth away from the active deformation front $\mathcal{B}_{\text{active}}$, the claim follows. □

Remark 9.7 (Piecewise conservative field)**.** *The impossibility field $\Phi_\varnothing$ is not globally smooth: the pointwise tropical superposition $\widetilde{\Phi}_\varnothing = \max(\Phi_\varnothing, \Phi_S)$ introduces kinks at the active deformation front $\mathcal{B}_{\text{active}}$ where $\Phi_\varnothing(h) = \Phi_S(h)$. At these loci the gradient is multi-valued in the classical sense, and the cognitive acceleration field develops shock-like behavior analogous to wavefront discontinuities in geometric optics. The impossibility field is therefore* piecewise conservative*: irrotational and path-independent between deformation events, with distributional singularities precisely at the active deformation front where falsification begins. This is not a defect of the framework but the mechanism: the deformation front is exactly where evidence does its structural work on the epistemic state.*

© 2026 by the author(s). Distributed under a Creative Commons CC BY license.

The conservative structure of Theorem 9.6 aligns with three TEAG axioms. Axiom A2 (Popperian contraction) ensures impossibility can only increase, so the potential drives hypotheses in one direction only. Axiom A3 (non-resurrection) prohibits closed loops in hypothesis trajectories. The contraction-based update rule means all epistemic motion is potential-driven: there is no intrinsic rotational component. The irrotationality of the cognitive acceleration field is therefore an internal consistency check, not an additional assumption.

9.3. Epistemic Energy and Its Non-Conservation

Although the cognitive acceleration field is piecewise conservative, epistemic energy is not conserved. This is a direct consequence of the contact geometry established in Section 8.

Definition 9.8 (Cognitive energy)**.** *Define the* cognitive kinetic energy, cognitive potential energy, *and* total cognitive energy *as*

$$T_c(h, \dot{h}) := \tfrac{1}{2}\, \dot{h}^\top G(h)\, \dot{h}, \tag{55}$$

$$V_c(h) := \Phi_\varnothing(h), \tag{56}$$

$$E_c(h, \dot{h}) := T_c(h, \dot{h}) + V_c(h). \tag{57}$$

Here T_c measures the rate of traversal through admissible hypothesis space; V_c measures the depth of the current state in the impossibility geometry; E_c measures the total epistemic activity available to the system at a given configuration and velocity.

Definition 9.9 (Cognitive work)**.** *The* cognitive work *done by evidence along a trajectory γ is*

$$W_c[\gamma] := \int_\gamma a \cdot_G dh = \int_\gamma (-\nabla_G \Phi_\varnothing) \cdot_G dh. \tag{58}$$

Since the field is piecewise conservative (Theorem 9.7), this reduces to

$$W_c[\gamma] = \Phi_\varnothing(h_{\text{start}}) - \Phi_\varnothing(h_{\text{end}}) \tag{59}$$

wherever $\Phi_\varnothing$ is smooth along γ.

Proposition 9.10 (Epistemic energy is not conserved)**.** *The total cognitive energy E_c is not conserved along TEAG trajectories.*

Proof. By Axiom A2 (Popperian contraction), $\pi'(h) \leq \pi(h)$ for all h, which implies $\Phi'_\varnothing(h) \geq \Phi_\varnothing(h)$. Thus V_c increases monotonically under evidence. No compensating decrease in T_c is guaranteed, since T_c is determined by the velocity of the trajectory, which is also evidence-dependent. There is therefore no energy invariant of the form $T_c + V_c = \text{const}$.

More fundamentally, the TEAG system is open: evidence acts as an external forcing term that irreversibly deforms the potential. The system is not closed under its own dynamics. Conservation of energy holds only for closed, reversible systems; TEAG is neither. This is the dynamical expression of the same asymmetry that forces contact rather than symplectic geometry in Section 8. □

Remark 9.11 (Epistemic action and the role of $\bar{S}_k$)**.** *The non-conservation of E_c means that evidence* injects *epistemic energy into the system. The natural measure of this injection per update is the aggregate epistemic surprisal $\bar{S}_k$ of Theorem 6.1. Two distinct scalars derived from $\bar{S}_k$ play complementary roles in the architecture, and the distinction matters:*

(i) $\bar{S}_k$ *is the* epistemic action*: the Choquet-integrated surprisal that quantifies the total epistemic work available to deform the admissible support at step k. It is unbounded and behaves like an action*

© 2026 by the author(s). Distributed under a Creative Commons CC BY license.

functional — the right quantity for energy bookkeeping, cognitive work, and the learning and bias diagnostics of Theorems 9.12 and 9.19.

(ii) $I_k = 1 - e^{-\bar{S}_k} \in [0,1)$ *is the* geometric contraction control*: the bounded transform of* $\bar{S}_k$ *that enters the PCRB inequality and the basin-radius regeneration formula* $r_k^\star = r_k^- \cdot ((1 - I_k) \cdot M/M_{\text{surv}})^{1/n}$. *It governs how much the admissible basin can shrink per observation, not how much epistemic work is done.*

An observation with $\bar{S}_k = 0$ *(zero epistemic action) does no cognitive work:* $W_c = 0$*,* E_c *is unchanged, and no learning occurs. An observation with* $\bar{S}_k > 0$ *but* $\mathcal{L}_k \approx 0$ *has action available but unrealized — the signature of inertial shielding. The PCRB, expressed in terms of* I_k*, is the minimum-action principle for this system: it bounds how much of the available epistemic action* $\bar{S}_k$ *can be realized as admissible contraction per update. The ESPF achieves this bound with equality at the admissible basin boundary.*

9.4. The No-Learning Theorem

The cognitive energy framework yields a precise characterization of the conditions under which learning is impossible. Two scalars govern this:

- $\bar{S}_k := \sup_i \min(\frac{1}{2}q_k^{(i)}, \pi_{k|k-1}^{(i)})$ — the *epistemic action*: the Choquet integral of per-hypothesis surprisal with respect to the prior possibility capacity (Theorem 6.1), quantifying the total epistemic work available to deform admissible support.
- $\mathcal{L}_k := d_H(S_k^-, S_k^+)$ — the *learning functional*: the Hausdorff distance between pre- and post-update admissible supports under the metric G, measuring realized deformation.

The no-learning condition is $\mathcal{L}_k = 0$. The transform $I_k = 1 - e^{-\bar{S}_k}$ governs the PCRB contraction geometry (Theorem 9.11); $\bar{S}_k$ governs the epistemic action.

Definition 9.12 (Learning)**.** *The epistemic system* learns *at update step* k *iff* $\mathcal{L}_k := d_H(S_k^-, S_k^+) > 0$*, where* $S_k^\pm$ *are the pre- and post-update admissible supports and* d_H *is the Hausdorff distance under the metric* G*.*

Definition 9.13 (Effective surprisal)**.** *The* effective surprisal *at step* k *is the component of the surprisal field that produces actual deformation of the admissible support:*

$$\Phi_S^{\text{eff}}(h) \;:=\; \max\big(\Phi_S(h) - \Phi_\varnothing(h),\, 0\big). \tag{60}$$

Effective surprisal is zero wherever the observation is already less surprising than the prior impossibility: the hypothesis is history-governed, not evidence-governed. It is positive only in the region where the active deformation front $\mathcal{B}_{\text{active}}$ *is crossed, precisely where falsification begins.*

Proposition 9.14 (Nonzero cognitive work is necessary for learning)**.** *If* $W_c[\gamma_k] = 0$ *at step* k*, then* $\mathcal{L}_k = 0$*: no learning occurs.*

Proof. $W_c = 0$ implies that the cognitive acceleration field does no work along the update trajectory. By Equation (59), this means $\Phi_\varnothing(h_{\text{start}}) = \Phi_\varnothing(h_{\text{end}})$ for all h in the support — the impossibility field is unchanged. Unchanged $\Phi_\varnothing$ means the posterior support S_k^+ coincides with the prior support S_k^-, giving $d_H(S_k^-, S_k^+) = 0$. □

Proposition 9.15 (Epistemic degeneracy of zero acceleration)**.** *Let* $a(h) = -\nabla_G \Phi_\varnothing(h)$ *be the cognitive acceleration field. Then* $a(h) = 0$ *does not uniquely determine the epistemic regime. In particular, both of the following conditions yield indistinguishable zero-acceleration dynamics and hence no learning:*

(i) ***Uninformative evidence (flat field):*** $\Phi_S(h) \le \Phi_\varnothing(h)$ *for all* h *in the support, so* $\Phi_S^{\text{eff}} \equiv 0$*, the active deformation front* $\mathcal{B}_{\text{active}}$ *is empty, and the impossibility field admits no gradient-driven deformation.*

© 2026 by the author(s). Distributed under a Creative Commons CC BY license.

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 15 April 2026 doi:10.20944/preprints202603.2110.v3

(ii) ***Inertial shielding (deep well):*** *The surprisal field Φ_S is substantial ($\bar{S}_k > 0$, equivalently $I_k > 0$), but the PCRB admissible boundary $\mathcal{B}_{\text{adm}}$ is not reached: the wavefront crosses $\mathcal{B}_{\text{active}}$ but no hypothesis exits the admissible basin, so $d_H(S_k^-, S_k^+) = 0$.*

A system observing only its trajectory $h(\tau)$ cannot determine from the vanishing of a alone which regime it inhabits.

Proof. In Case (i), $\widetilde{\Phi}_\varnothing = \max(\Phi_\varnothing, \Phi_S) = \Phi_\varnothing$ everywhere on the support, so no update occurs and $a = -\nabla_G \Phi_\varnothing$ is unchanged. In Case (ii), $\widetilde{\Phi}_\varnothing > \Phi_\varnothing$ locally near $\mathcal{B}_{\text{active}}$, but since no hypothesis crosses $\mathcal{B}_{\text{adm}}$, the admissible support S_k^+ is identical to S_k^- and the post-update cognitive acceleration field evaluated on the surviving support is unchanged. In both cases $\mathcal{L}_k = 0$ and the trajectory is stationary. □

Remark 9.16 (Why degeneracy matters)**.** *The epistemic degeneracy of Theorem 9.15 is the companion result to the epistemic equivalence principle of Theorem 9.4. The equivalence principle concerns the indistinguishability of the* sources of motion*: field-induced acceleration and geometry-induced geodesic deviation are locally indistinguishable when $\ddot{h} \neq 0$. The degeneracy proposition concerns the indistinguishability of the* causes of non-motion*: uninformative evidence and inertial shielding both produce $a = 0$ and are indistinguishable from the trajectory alone. Together they establish that epistemic trajectories are doubly degenerate — both in what drives them and in what arrests them.*

This degeneracy exposes a structural blind spot present in virtually all inference frameworks: Bayesian systems collapse both regimes into "no update," and adaptive systems call both "converged." The two failure modes are not merely different in cause — they are different in cure. Case (i) requires more informative observations; Case (ii) requires either broader prior support or a structural revision of the admissible geometry. A system that cannot distinguish them cannot correctly diagnose its own epistemic failure. The confirmation bias functional of Theorem 9.19 is precisely the diagnostic that separates the two regimes.

Corollary 9.17 (Two failure modes of learning)**.** *Zero cognitive acceleration is a necessary but not sufficient condition for epistemic equilibrium. The two failure modes yield the following regime table under the aggregate surprisal $\bar{S}_k$ and learning functional $\mathcal{L}_k$:*

Regime	$\bar{S}_k$	$\mathcal{L}_k$	*Interpretation*
Flat field	≈ 0	≈ 0	*Evidence carries no discriminative structure*
Adaptive learning	> 0	> 0	*Learning commensurate with evidence*
Inertial shielding	> 0	≈ 0	*Epistemic resistance to deformation*

Case I corresponds to the world saying nothing new; Case II to the world speaking while the system is unable to hear.

Remark 9.18 (Connection to epistemic time)**.** *If epistemic time is the experience of informational change, a natural definition of the epistemic time rate is $d\mathcal{T}/d\tau \propto \|v\|_G = \|\dot{h}\|_G$, the cognitive velocity in the metric norm. The cognitive acceleration field then governs the second derivative: $d^2\mathcal{T}/d\tau^2 \propto \|a\|_G = \|\nabla_G \Phi_\varnothing\|_G$. The no-learning theorem has a temporal interpretation: zero cognitive work means cognitive velocity does not change, epistemic time does not advance, and from the system's point of view nothing has happened. Epistemic stasis and temporal stasis coincide. The degeneracy of Theorem 9.15 then has a temporal reading as well: two systems can experience identical epistemic stasis — identical arrest of time — for entirely different structural reasons.*

9.5. Confirmation Bias as Inertial Shielding

The framework yields a computable diagnostic index for the inertial-shielding regime of Theorem 9.15.

© 2026 by the author(s). Distributed under a Creative Commons CC BY license.

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 15 April 2026 doi:10.20944/preprints202603.2110.v3

Definition 9.19 (Confirmation bias index)**.** *Let $\bar{S}_k$ be the aggregate epistemic surprisal (Theorem 6.1) and let $\mathcal{L}_k = d_H(S_k^-, S_k^+)$ be the Hausdorff deformation of admissible supports at step k. The* confirmation bias index *is*

$$\mathcal{B}_k := \frac{\bar{S}_k}{\mathcal{L}_k + \varepsilon}, \tag{61}$$

where $\varepsilon > 0$ is a small regularizing constant.

Here $\bar{S}_k$ is the same aggregate surprisal that governs the PCRB contraction bound (Theorem 6.2): the quantity that bounds admissible entropy reduction also serves as the natural measure of epistemic work available to deform the support geometry. The bias index therefore measures the mismatch between geometrically available contraction (determined by $I_k = 1 - e^{-\bar{S}_k}$ via the PCRB) and actual deformation.

Interpretation: $\mathcal{B}_k \approx 1$ indicates learning commensurate with available surprisal; $\mathcal{B}_k \gg 1$ indicates inertial shielding — counterevidence arrives but cognitive work is not realized as support deformation; $\mathcal{B}_k \to \infty$ is complete inertial shielding.

Remark 9.20 (Three-way unification via I_k)**.** *The aggregate surprisal $\bar{S}_k$ (and its transform I_k) now appears at three distinct levels of the TEAG architecture:*

(i) ***Dynamics:*** *I_k controls the admissible basin radius via $r_k^\star = r_k^- \cdot ((1 - I_k) \cdot M/M_{\mathrm{surv}})^{1/n}$.*
(ii) ***Information:*** *$\bar{S}_k$ is the PCRB-governing Choquet integral that bounds admissible entropy reduction per update.*
(iii) ***Diagnostics:*** *$\bar{S}_k$ is the numerator of the confirmation bias index, measuring available epistemic work.*

This is not a coincidence. The same scalar is the control parameter of admissible geometry, the information-theoretic budget of the update, and the reference against which realized deformation is measured. The quantity that bounds action also measures resistance to action.

Remark 9.21 (Distinguishing the degeneracy regimes)**.** *Theorem 9.15 establishes that zero cognitive acceleration is degenerate. The confirmation bias index resolves the degeneracy diagnostically: in the flat-field regime, $\bar{S}_k \approx 0$ and $\mathcal{B}_k$ is small regardless of $\mathcal{L}_k$; in the inertial-shielding regime, $\bar{S}_k > 0$ and $\mathcal{B}_k \gg 1$. Confirmation bias is therefore the specific case of inertial shielding where the available epistemic work $\bar{S}_k$ is large — the system has the information it needs, but the geometry of the admissible basin prevents it from responding.*

9.6. Summary of the Cognitive Mechanics Layer

The five fundamental objects of this section and their roles:

Object	Definition	Role
Cognitive velocity	$v = \dot{h}$	Rate of traversal through H
Cognitive acceleration	$a = -\nabla_G \Phi_\varnothing$	Change in rate of admissible cognition
Cognitive kinetic energy	$T_c = \frac{1}{2}\dot{h}^\top G \dot{h}$	Epistemic activity (motion)
Cognitive potential energy	$V_c = \Phi_\varnothing$	Epistemic activity (position in well)
Cognitive work	$W_c[\gamma] = \int_\gamma a \cdot_G dh$	Evidence's structural action

These objects are unified through the aggregate epistemic surprisal $\bar{S}_k$ and its geometric transform $I_k = 1 - e^{-\bar{S}_k}$, which together span the full TEAG architecture:

Scalar	Role	Appears in
$\bar{S}_k$	Epistemic action (available work)	Cognitive work, bias index $\mathcal{B}_k$
I_k	Geometric contraction control	PCRB bound, basin radius $r_k^\star$

The central results of this section are:

(i) **Epistemic equivalence principle** (Theorem 9.4): field-induced cognitive acceleration and geometry-induced geodesic deviation are locally indistinguishable from the trajectory alone; the equivalence breaks at deformation fronts.

© 2026 by the author(s). Distributed under a Creative Commons CC BY license.

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 15 April 2026 doi:10.20944/preprints202603.2110.v3

(ii) **Necessary condition for learning** (Theorem 9.14): nonzero cognitive work is necessary for support deformation.

(iii) **Epistemic degeneracy** (Theorem 9.15): zero cognitive acceleration is insufficient to determine the epistemic regime; uninformative evidence ($\bar{S}_k \approx 0$) and inertial shielding ($\bar{S}_k > 0$, $\mathcal{L}_k \approx 0$) are observationally indistinguishable from the trajectory alone.

(iv) **Diagnostic resolution** (Theorem 9.19): the confirmation bias index $\mathcal{B}_k = \bar{S}_k/(\mathcal{L}_k + \varepsilon)$ separates the two degenerate regimes using the same epistemic action $\bar{S}_k$ that governs the PCRB.

The complete chain through the TEAG architecture is:

$$\underbrace{\Phi_S}_{\text{Hamiltonian}} \longrightarrow \underbrace{\Phi_\varnothing}_{\text{potential}} \longrightarrow \underbrace{a = -\nabla_G \Phi_\varnothing}_{\text{cognitive acceleration}} \longrightarrow \underbrace{\bar{S}_k}_{\text{action}}$$
$$\longrightarrow \underbrace{I_k}_{\text{contraction}} \longrightarrow \underbrace{\text{PCRB}}_{\text{minimum action bound}} \longrightarrow \underbrace{\text{ESPF}}_{\text{action-minimizing evolution.}} \tag{62}$$

10. The Zero-Temperature Limit

The connection between classical (probabilistic) and tropical (possibilistic) Hamilton–Jacobi theory is made precise by a thermodynamic degeneration.

In statistical mechanics, the free energy at temperature T is

$$F = -T\log Z = -T\log\int_H \exp\left(-\frac{\Phi(h)}{T}\right)d\mu(h). \tag{63}$$

As $T \to \infty$, $F \to -T\log|H|$ (uniform prior, maximum ignorance). As $T \to 0$, $F \to \min_h \Phi(h)$ (ground state, minimum potential).

The Bayesian posterior corresponds to finite temperature:

$$p(h|y) \propto \exp\left(-\frac{\Phi_S(h)}{T}\right)\cdot\exp\left(-\frac{\Phi_\varnothing(h)}{T}\right), \tag{64}$$

where the denominator (partition function) is the normalizing integral. In field coordinates, the Bayesian posterior log-probability is $-\frac{1}{T}(\Phi_\varnothing(h) + \Phi_S(h))$: the impossibility and surprisal fields *add* at finite temperature (log-sum aggregation).

At $T \to 0$, the log-sum-exp aggregation degenerates to the max. The physical meaning of this limit is precise: at finite temperature, hypotheses compete through weighted averaging over their impossibility fields, with every hypothesis retaining nonzero influence; at zero temperature, only the least impossible hypothesis governs and all others are expelled.

$$-T\log\left(e^{-\Phi_\varnothing/T} + e^{-\Phi_S/T}\right) \xrightarrow{T\to 0} \max(\Phi_\varnothing, \Phi_S) = \widetilde{\Phi}_\varnothing. \tag{65}$$

The Bayesian posterior (fields add at finite T) degenerates to the possibilistic posterior (fields take pointwise tropical superposition at $T = 0$).

Theorem 10.1 (Possibilistic inference as zero-temperature limit)**.** *The conjunctive possibilistic update $\widetilde{\Phi}_\varnothing = \max(\Phi_\varnothing, \Phi_S)$ is the zero-temperature ($T \to 0$) limit of the Bayesian update $\widetilde{\Phi}_{\varnothing,T} = -T\log(e^{-\Phi_\varnothing/T} + e^{-\Phi_S/T})$. Equivalently, the ESPF is the ground-state filter: it operates at the minimum epistemic temperature consistent with non-trivial support, where the partition function collapses to the ground state energy and Bayesian averaging collapses to possibilistic selection.*

Remark 10.2 (Epistemic temperature and the UKF–ESPF transition)**.** *The epistemic temperature T parametrizes the continuum between the Unscented Kalman Filter (UKF, $T = T_{\text{Gauss}}$, finite) and the ESPF ($T = 0$). The Hölder mean hierarchy of Jah (2025b) — in which the UKF minimizes the $\log M_0$ (geometric mean) functional and the ESPF minimizes $\log M_{+\infty}$ (max-plus) functional — is exactly this temperature continuum.*

© 2026 by the author(s). Distributed under a Creative Commons CC BY license.

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 15 April 2026 doi:10.20944/preprints202603.2110.v3

In TEAG*language (Jah 2026b), this is* convergent optimality*: the UKF and ESPF are optimal solutions to categorically different problems — MSE minimization and maximum-entropy possibilistic estimation respectively — that agree precisely when the world is Gaussian, the model is valid, and the system is stable. This convergence is a mathematical fact about the Hölder functional, not a containment relationship: the ESPF does not contain the UKF, and the UKF is not a special case of the ESPF. The formal Choquet-to-Lebesgue convergence theorem establishing this limit is proved in Jah (2026a) and stated in* TEAG *(Jah 2026b) as the Gaussian collapse theorem.*

The zero-temperature limit ($T \to 0$, contact $\to$ TEAG) and the zero-information limit ($\bar{S}_k \to 0$, contact $\to$ symplectic, Theorem A13) are two orthogonal directions in the space of epistemic theories. The first recovers possibilistic from probabilistic dynamics. The second recovers reversible from irreversible dynamics. Together they span the full range from maximum uncertainty (contact, possibilistic) to perfect knowledge (symplectic, static).

11. The Minimax Medioid as Geodesic Attractor

We first make the metric structure explicit.

Definition 11.1 (MVEE-whitened metric). *Let $\Pi_e = L_e L_e^\top \in \mathbb{R}^{m\times m}$ be the MVEE of the predicted measurement support $\{g(\chi^{(i)}_{k|k-1})\}$ at time k, with Chebyshev center used as the ellipsoid reference point during fitting. The* MVEE-whitened metric *on the surviving support H_{surv} is the inner product*

$$\langle h_1 - h_2, h_1 - h_2\rangle_W = (g(h_1) - g(h_2))^\top \Pi_e^{-1}(g(h_1) - g(h_2)), \tag{66}$$

where $g : H \to \mathbb{R}^m$ is the measurement mapping. The induced norm $\|\cdot\|_W$ is the same norm in which the surprisal field is defined: $\Phi_S(h) = \frac{1}{2}\|z_k(h)\|_W^2$ with $z_k(h) = L_e^{-1}(y_k - g(h))$. Note the temporal separation: Π_e is computed from the predicted *support cloud before falsification; the metric d_W it defines is then applied to the* surviving *support H_{surv} after falsification to select the commitment point. The MVEE-whitened metric is therefore the* natural metric of the surprisal field*: it is the unique metric under which the surprisal field is a squared-distance function centered at the observation.*

The minimax medioid is defined with respect to this metric as the surviving support point that minimizes the maximum whitened distance to all other survivors:

$$h^\star = \operatorname{argmin}_{h\in H_{\text{surv}}} \max_{\chi^{(j)}\in H_{\text{surv}}} \left\| h - \chi^{(j)} \right\|_W. \tag{67}$$

Crucially, $h^\star \in H_{\text{surv}}$: it is an actual member of the surviving support cloud, not an interpolated geometric abstraction. It is the support point most equidistant from its peers in the whitened metric — the one most insulated from the boundary of the surviving well in every direction.

Definition 11.2 (Geodesics on the knowledge manifold). *A* geodesic *on the knowledge manifold within the surviving well is a path $\gamma : [0, 1] \to H_{\text{surv}}$ that minimizes the length functional $\int_0^1 \|\dot{\gamma}(t)\|_W \, dt$ in the MVEE-whitened metric. Under the local quadratic approximation to the surprisal field (as in Section 4), geodesics are straight lines in the MVEE-whitened coordinates — the coordinate system in which $\Phi_S(h)$ is a perfect sphere centered at y.*

Proposition 11.3 (Medioid as nearest realized hypothesis to geoid center, and geodesic attractor). *The whitened minimax medioid $h^\star$ is the surviving support point nearest the center of the PCRB-defined epistemic geoid among all realized hypotheses in H_{surv}. It is not the geoid center itself — that is a geometric interpolation which need not correspond to any realized hypothesis — but the member of H_{surv} that minimizes worst-case whitened distance to all other survivors. Specifically, $h^\star$ is the unique point in H_{surv} satisfying*

$$\max_{\chi^{(j)}\in H_{\text{surv}}} d_W(h^\star, \chi^{(j)}) \le \max_{\chi^{(j)}\in H_{\text{surv}}} d_W(h, \chi^{(j)}) \quad \forall\, h \in H_{\text{surv}}, \tag{68}$$

© 2026 by the author(s). Distributed under a Creative Commons CC BY license.

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 15 April 2026 doi:10.20944/preprints202603.2110.v3

where $d_W(h, h') = \|h - h'\|_W$ is the whitened geodesic distance. It is the geodesic attractor of the knowledge manifold: the support point from which every geodesic direction reaches another survivor before ascending steeply in impossibility, and therefore the support point that is maximally insulated from the PCRB equipotential boundary in all directions.

Proof. The minimax objective $f(h) = \max_j d_W(h, \chi^{(j)})$ is convex in d_W over H_{surv} since it is the pointwise maximum of convex distance functions. Its minimizer $h^\star$ exists and is unique when H_{surv} is finite (as in the ESPF support-point cloud) and achieves the smallest possible worst-case whitened distance to any survivor. Since $\Phi_S(h) = \frac{1}{2} d_W(h, y)^2$, the geoid equipotential surface $\{h : \Phi_S(h) = c\}$ is a sphere of radius $\sqrt{2c}$ in the d_W metric centered at y. The medioid minimizes the maximum d_W-distance to all survivors, so it is the realized support point most equidistant from the boundary of the survivor cloud in the metric that defines both the surprisal field and the geoid boundary — maximally insulated from the PCRB equipotential boundary across all directions. Unlike the MVEE Chebyshev center, which is a geometric interpolation and need not be a realized hypothesis, the medioid is always a member of H_{surv}: a specific support point $\chi^{(i)}$ selected from the cloud. □

Remark 11.4 (Two distinct roles: shape matrix vs. commitment point)**.** *The MVEE and the minimax medioid serve distinct roles in the ESPF and must not be conflated.*

Shape matrix (MVEE + Chebyshev center). *To compute the surprisal field and the PCRB, the ESPF needs a shape matrix $\Pi_e = L_e L_e^\top$ that characterizes the spread of the predicted measurement support $\{g(\chi^{(i)})\}$. This is obtained by fitting the MVEE to the predicted measurement cloud. The Chebyshev center of the MVEE — the center of the enclosing ellipsoid — is the reference point used during that fitting procedure. It is a geometric artifact of the MVEE computation, not a commitment point. The output is L_e^{-1}, used to whiten the residual $z_k(h) = L_e^{-1}(y_k - g(h))$ in the surprisal field definition.*

Commitment point (whitened minimax medioid). *The anchor for Smolyak cloud regeneration — both during propagation and after a measurement update — is the whitened minimax medioid $h^\star \in H_{\text{surv}}$: the non-falsified support point that minimizes the maximum d_W-distance to all other non-falsified survivors. This is a discrete constrained optimization over the surviving cloud, not a geometric center of an enclosing ellipsoid. The medioid is always a realized hypothesis; the MVEE Chebyshev center need not be.*

The connection. *The medioid is the closest realized support point to the Chebyshev center of the surviving cloud in the MVEE-whitened metric d_W. In the limit of a dense support cloud the two coincide; for a finite support-point cloud they differ, and the medioid is the correct choice because it is a member of H_{surv}. The d_W metric used in the medioid computation is itself defined by Π_e, so the shape matrix (MVEE) provides the metric, and the minimax medioid selects the commitment point within that metric.*

Regenerating the Smolyak cloud around $h^\star$ is the boundary condition reset of Theorem 5.5(iii): the impossibility field is reset to zero over the regenerated cloud, placing maximum epistemic support at the geodesic attractor of the previous well. The filter propagates maximum ignorance — zero impossibility — within the region that the evidence and prior history have not yet condemned. This is the operational expression of the minimum action principle (Theorem 6.2): commit to the geodesic attractor, but assign no further impossibility to the surviving support than the evidence demands.

12. Epistemic Time: Discreteness, Continuity, and the Physical Limit

12.1. The Discrete Structure of Epistemic Time

Epistemic time is not a background parameter. It is a derived quantity: the count of informative updates registered by an epistemic agent. This section makes that structure precise, establishes its continuity limit, derives the local clock rate of a belief from the geometry of the impossibility well, and identifies gravitational time dilation as the physical instance of epistemic time dilation.

© 2026 by the author(s). Distributed under a Creative Commons CC BY license.

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 15 April 2026 doi:10.20944/preprints202603.2110.v3

Definition 12.1 (Epistemic time)**.** *The epistemic time of an agent after k updates is*

$$\mathcal{T}_{\mathrm{epi}}(k) = \sum_{j=1}^{k} \delta T_j, \qquad \delta T_j = \frac{\Phi^{\mathrm{eff}}_{S,j}}{\bar{S}_j + \varepsilon}, \tag{69}$$

where $\Phi^{\mathrm{eff}}_{S,j} = \max(\Phi_S(h,j) - \Phi_\varnothing(h,j), 0)$ *is the effective surprisal at step j and* $\varepsilon > 0$ *is a regularizing constant. Each tick* δT_j *is nonzero if and only if* $\Phi^{\mathrm{eff}}_{S,j} > 0$*: the surprisal wavefront has penetrated the impossibility well and genuine learning has occurred.*

Remark 12.2 (Epistemic time is discrete)**.** *Between updates, the impossibility field undergoes free transport (Theorem 5.5(i)) but no cognitive work is done, no falsification occurs, and no new information is registered. The filter is in epistemic free fall. Epistemic time does not advance between measurements. It advances in discrete ticks whose magnitude is determined by the effective surprisal — the portion of the incoming wavefront that successfully deforms the impossibility well. An update with* $\bar{S}_k > 0$ *but* $\Phi^{\mathrm{eff}}_S \equiv 0$ *(inertial shielding) does not advance epistemic time: the evidence arrives but cannot register.*

Proposition 12.3 (Local clock rate of a belief)**.** *The epistemic proper time rate of a belief at h is proportional to the effective surprisal that can penetrate the local impossibility well:*

$$\frac{d\mathcal{T}_{\mathrm{epi}}}{d\tau}(h) \;\propto\; \Phi^{\mathrm{eff}}_S(h) = \max(\Phi_S(h) - \Phi_\varnothing(h),\, 0). \tag{70}$$

- *(i)* *Where* $\Phi_\varnothing(h) = 0$ *(no accumulated epistemic history), the clock runs at the full rate of the incoming surprisal wavefront.*
- *(ii)* *Where* $\Phi_\varnothing(h)$ *is large (deep impossibility well), the clock slows: incoming evidence is dominated by prior history.*
- *(iii)* *Where* $\Phi_\varnothing(h) \geq \Phi_S(h)$ *everywhere on the support (epistemic black hole),* $\Phi^{\mathrm{eff}}_S \equiv 0$ *and epistemic time freezes completely.*

Proof. By Theorem 12.1, $\delta T_j \propto \Phi^{\mathrm{eff}}_{S,j}$. The effective surprisal is zero precisely when $\Phi_S(h) \leq \Phi_\varnothing(h)$ for all h in the surviving support — i.e., when the active deformation front $\mathcal{F}_k$ is empty. In this regime the tropical update $\widetilde{\Phi}_\varnothing = \max(\Phi_\varnothing, \Phi_S) = \Phi_\varnothing$ leaves the impossibility field unchanged, no hypothesis is expelled, $\mathcal{L}_k = 0$, and no cognitive work is done. The three cases follow directly from the ordering of Φ_S and $\Phi_\varnothing$. □

Theorem 12.4 (Confirmation bias as epistemic time dilation)**.** *Let* $\mathcal{B}_k = \bar{S}_k / (\mathcal{L}_k + \varepsilon)$ *be the confirmation bias index (Theorem 9.19). Then the epistemic proper time rate satisfies:*

$$\frac{d\mathcal{T}_{\mathrm{epi}}}{d\tau} \propto \frac{1}{\mathcal{B}_k}. \tag{71}$$

$\mathcal{B}_k \approx 1$ *(commensurate learning) corresponds to maximum epistemic clock rate.* $\mathcal{B}_k \gg 1$ *(inertial shielding) corresponds to epistemic time dilation.* $\mathcal{B}_k \to \infty$ *(epistemic black hole) corresponds to complete epistemic time freeze. The confirmation bias index is the external measurement of epistemic time dilation.*

Proof. When $\mathcal{B}_k \approx 1$, $\mathcal{L}_k \approx \bar{S}_k$, meaning the realized deformation matches the available epistemic action. By Theorem 12.3, Φ^{eff}_S is large relative to $\Phi_\varnothing$, so $d\mathcal{T}_{\mathrm{epi}}/d\tau$ is near its maximum. When $\mathcal{B}_k \gg 1$, $\mathcal{L}_k \ll \bar{S}_k$: evidence arrives but cannot penetrate the well, so $\Phi^{\mathrm{eff}}_S \approx 0$ and epistemic time stalls. The proportionality $d\mathcal{T}_{\mathrm{epi}}/d\tau \propto 1/\mathcal{B}_k$ follows from substituting the effective surprisal into Theorem 12.1 and normalizing by $\bar{S}_k$. □

12.2. The Continuity Limit

Epistemic time becomes continuous in the limit of constant information flux.

© 2026 by the author(s). Distributed under a Creative Commons CC BY license.

Theorem 12.5 (Epistemic time continuity limit). *Suppose evidence arrives at a constant rate with constant information content: $\bar{S}_k = \bar{S}$ for all k, and the inter-update interval $\Delta\tau \to 0$ with $\bar{S}/\Delta\tau \to \phi$ (a constant information flux density). Then:*

(i) *The discrete epistemic time increments δT_k become uniform.*

(ii) *The sum $\mathcal{T}_{\text{epi}}(k) = \sum_{j=1}^{k} \delta T_j$ converges to a Riemann integral.*

(iii) *The discrete ESPF recursion converges to the continuous-time tropical Hamilton–Jacobi PDE:*

$$\frac{\partial \Phi}{\partial t}(h,t) \oplus \Phi_S(h,t) = 0, \tag{72}$$

with $\oplus = \max$, whose unique viscosity solution is the Lax–Oleinik semigroup $\Phi(h,t) = \Phi_0(h) \oplus \sup_{0 \le s \le t} \Phi_S(h,s)$.

Proof. Part (iii) is Theorem 5.8. Parts (i) and (ii) follow from the fact that when $\bar{S}_k = \bar{S}$ is constant, the tick magnitudes $\delta T_k = \Phi^{\text{eff}}_{S,k}/(\bar{S} + \varepsilon)$ are bounded and uniform, and their sum is a Riemann sum with mesh $\Delta\tau \to 0$. □

Remark 12.6 (What makes physical time appear continuous). *Physical time appears continuous because the universe delivers information at a rate that is, to excellent approximation, constant and universal: every region of spacetime is embedded in a field structure that propagates information at the speed of light, continuously. Under Theorem 12.5, continuous time emerges when information flux is constant. Physical time is epistemic time in the limit of universal, constant information flux.*

12.3. Epistemic Time Dilation and Gravitational Time Dilation

The local clock rate of a belief (Theorem 12.3) has an exact structural parallel with gravitational time dilation in general relativity.

In GR, the proper time rate of a clock at gravitational potential Φ_{grav} relative to a clock at infinity is:

$$\frac{d\tau}{dt} = \sqrt{1 + \frac{2\Phi_{\text{grav}}}{c^2}} \approx 1 + \frac{\Phi_{\text{grav}}}{c^2} \tag{73}$$

for weak fields. A clock deeper in a gravitational well (more negative Φ_{grav}) runs slower. At the event horizon, $d\tau/dt \to 0$.

In TEAG, the epistemic proper time rate at h is $d\mathcal{T}_{\text{epi}}/d\tau \propto \Phi^{\text{eff}}_S(h)$. A belief deeper in the impossibility well (larger $\Phi_\varnothing(h)$) has its epistemic clock run slower. At the epistemic horizon (Theorem A3), $d\mathcal{T}_{\text{epi}}/d\tau \to 0$.

Theorem 12.7 (Structural identity of epistemic and gravitational time dilation). *Epistemic time dilation and gravitational time dilation implement the same Hamilton–Jacobi shielding mechanism. Specifically:*

(i) GR mechanism. *A clock at gravitational potential Φ_{grav} runs at rate $d\tau/dt \approx 1 + \Phi_{\text{grav}}/c^2$. The governing equation for null geodesics is the classical Hamilton–Jacobi equation $\partial_t S + H(q, \nabla_q S) = 0$. Deeper in the well, the clock runs slower; at the event horizon, $d\tau/dt \to 0$.*

(ii) TEAGmechanism. *A belief at impossibility depth $\Phi_\varnothing(h)$ advances epistemic time at rate $d\mathcal{T}_{\text{epi}}/d\tau \propto \Phi^{\text{eff}}_S(h)$. The governing equation is the tropical Hamilton–Jacobi equation $\partial_k \Phi \oplus H_{\text{trop}}(h, \nabla_h \Phi) = 0$. Deeper in the well, the clock runs slower; at the epistemic horizon, $d\mathcal{T}_{\text{epi}}/d\tau \to 0$.*

(iii) Governing equations are structurally identical.

$$\textit{GR (classical):} \quad \frac{\partial S}{\partial t} + H(q, \nabla_q S) = 0, \tag{74}$$

$$\text{TEAG}\textit{(tropical):} \quad \frac{\partial \Phi}{\partial k} \oplus H_{\text{trop}}(h, \nabla_h \Phi) = 0. \tag{75}$$

© 2026 by the author(s). Distributed under a Creative Commons CC BY license.

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 15 April 2026 doi:10.20944/preprints202603.2110.v3

In GR, the Hamiltonian encodes geodesic dynamics and the potential well slows the proper time rate. In TEAG*, the Hamiltonian is the surprisal field Φ_S and the impossibility well slows the epistemic time rate. The shielding mechanism — potential depth suppressing the local clock — is the same equation in both cases, differing only in the algebra (classical + vs. tropical $\oplus = \max$).*

(iv) Horizon formation criterion. *The PCRB criterion $\lambda_{\min}(\nabla^2\Phi_\varnothing)\,\rho^2_{\mathrm{PCRB}} \geq 2c_{\mathrm{epi}}$ (Theorem A7) states that the inward potential barrier at the geoid exceeds the maximal admissible wavefront action per update: the epistemic equivalent of the condition under which an event horizon forms.*

Items (i)–(iii) establish a structural identity of governing equations. Item (iv) is a structural correspondence. The ontological question of whether these reflect a single underlying geometry is addressed in the companion programme (Jah 2026b).

Proof. (i) is standard GR. (ii) follows from Theorem 12.3: when $\Phi_\varnothing(h) \geq \Phi_S(h)$, the tropical max selects $\Phi_\varnothing$ and $\Phi_S^{\mathrm{eff}} = 0$. (iii) The two equations have identical structure: a time derivative of a scalar action field plus a Hamiltonian evaluated at the field gradient equals zero. The shielding mechanism is the same in both: when the potential term dominates the Hamiltonian, the wavefront cannot advance. (iv) follows from Theorem A7. □

Remark 12.8 (Inertia as the common source)**.** *In GR, mass generates gravitational potential, which generates inertia for test particles, which slows the local clock. In* TEAG*, accumulated evidence generates impossibility curvature $M(h) = \nabla_h^2\Phi_\varnothing$, which generates belief inertia, which slows the epistemic clock. Both implement the same principle: inertia slows the local clock. Confirmation bias is not a psychological anomaly; it is a clock running slow under the weight of accumulated evidence.*

12.4. Open Directions in Epistemic Time

Structural correspondence with physical spacetime. Theorem 12.7 establishes that epistemic and gravitational time dilation implement the same Hamilton–Jacobi shielding mechanism. Whether this reflects a deeper connection between the TEAGimpossibility field and the Einstein field equations is the subject of the companion programme (Jah 2026b).

Wavefunction collapse and the Born rule. The quantum impossibility field $\Phi_\varnothing(x,p) = -\log|W(x,p)|$ suggests a TEAGaccount of wavefunction collapse as a tropical update selecting one branch of the surviving support. Deriving the Born rule $P(\text{outcome}) \propto |\psi(y)|^2$ from the PCRB geometry is an open problem.

Arrow of epistemic time. The irreversibility of TEAG(Axiom A3) defines a preferred direction of epistemic time. Whether this is related to the thermodynamic and cosmological arrows is open. The contact geometry of Section 8 suggests all three may be instances of the same underlying contact dissipation.

13. Discussion and Open Directions

13.1. Position within TEAG

The Theory of Epistemic Abductive Geometry (TEAG) (Jah 2026b) identifies a single mathematical object — the TEAGquintuple $\mathcal{E} = (H, \pi, \{H_\alpha\}, C, A)$ — as the shared primitive underlying four previously distinct inference frameworks: the ESPF for recursive state estimation, the Geometry of Knowing for measure-theoretic collapse (Jah 2026a), the ESPF Optimality theorem for minimax-entropy estimation (Jah 2025b), and the Possibilistic Language Model for bounded language generation.

Within this architecture, each of the central objects developed in the present paper has a precise home. The impossibility field $\Phi_\varnothing = -\log\pi$ is the natural geometric object of the TEAGquintuple (Jah 2026b), the log-admissibility coordinate that connects possibility theory to tropical algebra. The pointwise tropical superposition $\widetilde{\Phi}_\varnothing = \Phi_\varnothing \oplus \Phi_S$ is the canonical TEAGcontraction in impossibility coordinates (Jah 2026b). The surprisal $\Phi_S(h) = \frac{1}{2}\|z(h)\|^2$ is the impossibility increment induced by evidence — an identity, not an analogy, as established in the TEAG surprisal remark (Jah 2026b). The

© 2026 by the author(s). Distributed under a Creative Commons CC BY license.

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 15 April 2026 doi:10.20944/preprints202603.2110.v3

PCRB as epistemic geoid equipotential surface and the whitened minimax medioid as geoid center are both introduced in TEAG (Jah 2026b) as the epistemic geoid remark and the minimax commitment axiom respectively.

The present paper contributes the *dynamical* geometric foundation that TEAG (Jah 2026b) identifies as future work: the explicit Hamilton–Jacobi theorem (Theorem 5.5), the PCRB as variational minimum action (Theorem 6.2), the zero-temperature limit connecting Bayesian and possibilistic inference (Theorem 10.1), the whitened metric anchoring of the medioid as geodesic attractor (Theorem 11.3), and the identification of epistemic phase space as a contact manifold (Section 8) with the PCRB as a contact energy floor and the ESPF as a contact Hamiltonian system with discrete projection. Where TEAGestablishes that the impossibility field update *has the structure* of a discrete max-plus Hamilton–Jacobi update, the present paper names the tropical Hamiltonian explicitly (Theorem 5.3), proves the full recursion, grounds every claim in the metric geometry of the MVEE-whitened measurement space, and identifies the irreversibility of Popperian falsification as the geometric reason contact geometry is the correct primitive rather than symplectic geometry.

The unifying statement. Every layer of the architecture developed in this paper is governed by a single scalar functional: the aggregate epistemic surprisal $\bar{S}_k$, defined as the Choquet integral of per-hypothesis surprisal with respect to the prior possibility capacity. This is not a convenience of notation. The same functional determines the propagation of the deformation wavefront (whether $\mathcal{B}_{\text{active}}$ is non-empty), governs admissible contraction through the PCRB (via $I_k = 1 - e^{-\bar{S}_k}$), and constitutes the epistemic action whose gradient drives cognitive acceleration. In this sense, epistemic dynamics are fully determined by a single quantity: the action of surprisal on the hypothesis manifold.

When this quantity vanishes, the surprisal field fails to overtake the impossibility field anywhere on the support, no deformation front forms, no wavefront propagates, and cognitive acceleration is identically zero. The system becomes incapable of distinguishing change. Learning, falsification, and the advance of epistemic time all cease simultaneously — not by separate mechanisms but as unified consequences of the vanishing of a single scalar. This is what it means for the TEAG architecture to be internally forced: the impossibility of learning under zero surprisal is not a design choice or an added axiom. It is what the structure demands. A natural reviewer concern is whether the mechanics of Section 9 — cognitive acceleration, epistemic inertia, kinetic and potential energy, cognitive work — constitutes a genuine structural result or an elaborate analogy to physics. The answer is precise and worth stating directly. A related and deeper question — whether the impossibility field is a freely chosen field on a fixed background geometry or a coordinate marker of an evolving space — is addressed in Appendix A.10.

The objects are not analogies. The impossibility field $\Phi_\varnothing = -\log \pi$ is a scalar field on a manifold; its negative gradient is a vector field; that vector field is the right-hand side of an Euler–Lagrange equation; the Euler–Lagrange equation is derived from the Lagrangian $\mathcal{L}_c = T_c - V_c$ whose kinetic term is the standard Riemannian kinetic energy in the MVEE-whitened metric. None of these steps invoke physics; they invoke variational calculus on a Riemannian manifold. The same mathematics would arise in any field with a scalar potential on a curved space — geometry, economics, optimal control. The word "acceleration" is not borrowed from Newton; it is the coordinate-free second derivative of a curve. The word "work" is not borrowed from thermodynamics; it is the line integral of a covector field. These are standard mathematical objects that happen to share names with physical quantities because physics encountered them first.

What is specific to TEAG is the *interpretation*: the scalar field is the impossibility potential, the manifold is hypothesis space, the metric is epistemic inertia, and trajectories are the evolution of committed epistemic states. The structural results — irrotationality, non-conservation of cognitive energy, the degeneracy of zero acceleration, the equivalence principle — follow from properties of this specific field and manifold, not from physical intuition imported from outside. In particular, the epistemic equivalence principle (Theorem 9.4) is not a postulate by analogy with GR; it is a theorem about the indistinguishability of two terms in the same Euler–Lagrange equation. And the epistemic

© 2026 by the author(s). Distributed under a Creative Commons CC BY license.

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 15 April 2026 doi:10.20944/preprints202603.2110.v3

degeneracy (Theorem 9.15) is not a philosophical observation; it is a proposition with a proof that follows from the tropical superposition structure of the update.

The mechanics layer is therefore neither metaphor nor analogy. It is the natural differential geometry of a scalar potential on a Riemannian manifold, applied to the specific manifold and potential that TEAG defines.

Formal convergence via Lyapunov potential. The posterior potential $\widetilde{\Phi}_{\varnothing}$ is a natural Lyapunov candidate for filter stability. If $\widetilde{\Phi}_{\varnothing}$ decreases (in the sublevel set volume sense) monotonically toward the PCRB floor under repeated updates with consistent dynamics, then the ESPF is stable in the possibilistic Lyapunov sense. Making this argument rigorous requires bounding the potential increase from regeneration against the decrease from falsification — the key step that tropical Hamilton–Jacobi machinery may provide. In the contact geometry language of Section 8, this is the question of whether the contact Hamiltonian flow is Lyapunov stable with respect to the admissible basin volume.

Tropical curvature, the PCRB, and the MVEE–tropical correspondence. The PCRB is expressed via the Choquet integral of per-hypothesis surprisal over the possibility-weighted support. The active deformation front $\mathcal{F}_k$ has a curvature theory in the sense of tropical geometry (Maclagan and Sturmfels 2015). We conjecture that the tropical curvature of $\mathcal{F}_k$ at the medioid is related to the PCRB by a tropical Gauss–Bonnet identity — a possibilistic analog of the classical relationship between curvature and information. This conjecture gains concrete geometric grounding from the MVEE–tropical correspondence (Equation (6)): if the MVEE support vectors of the predicted measurement cloud $\{v_i\}$ are the tropical vertices of tconv$\{v_i\}$, then the tropical geometry of $\mathcal{F}_k$ is inherited from the tropical hull of the support cloud. The deformation front, the PCRB-defined epistemic geoid, and the tropical convex hull of the measurement cloud are all objects in the same tropical algebraic geometry; a tropical Gauss–Bonnet identity relating them would be a precise theorem on this common geometric object. In contact geometry, the Reeb vector field of the contact form η generates the curvature of the contact manifold; the relationship between this curvature and the PCRB bound is an open problem. A further open question is whether the Cholesky factorization $\Pi_e = L_e L_e^\top$ preserves any tropical structure from Π_e — the interaction between tropical algebra and matrix decompositions.

Multi-body epistemic mechanics. The present paper treats a single observation source generating a single surprisal wavefront. Multiple simultaneous sensors generate multiple wavefronts whose tropical superposition determines the posterior potential. The epistemic Lagrange points of this multi-body problem — the loci where multiple surprisal fields balance — determine the falsification structure under sensor fusion. This is the epistemic analog of the restricted n-body problem, and the tropical geometry of multi-source superposition is an open problem of practical significance for multi-sensor ESPF implementations. In the contact geometry framework, multi-sensor fusion corresponds to the superposition of multiple contact Hamiltonians, whose joint fixed point structure determines the composite active deformation front.

Cognitive mechanics: epistemic thermodynamics and dissipation. The cognitive mechanics layer of Section 9 formalizes a mechanics of admissible inference around five objects (cognitive velocity, acceleration, kinetic energy, potential energy, and work) and establishes that the same aggregate surprisal $\bar{S}_k$ governs three distinct levels of the architecture: the admissible basin radius via the regeneration formula, the PCRB contraction bound, and the confirmation bias index. The epistemic equivalence principle (Theorem 9.4) and the epistemic degeneracy proposition (Theorem 9.15) are companion results characterizing, respectively, the indistinguishability of the sources of motion and the causes of stasis. Several extensions remain open. First, the breaking of the epistemic equivalence at deformation fronts suggests a curvature-based observable analogous to tidal forces in GR: the Riemann tensor of the epistemic manifold H evaluated at $\mathcal{B}_{\text{active}}$ may furnish a geometric measure of the distinguishable component of field-induced vs. geometry-induced acceleration, with potential applications to adaptive support-point placement in the ESPF. Second, summing the cognitive work $W_c[\gamma_k]$ over observations defines a cumulative epistemic work functional; relating this integral to the possibilistic entropy trajectory $\mathcal{E}_{\pi,k}$ would close the loop between the mechanics and the information-

© 2026 by the author(s). Distributed under a Creative Commons CC BY license.

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 15 April 2026 doi:10.20944/preprints202603.2110.v3

theoretic layers. Third, the confirmation bias index $\mathcal{B}_k$ provides a per-step diagnostic but a principled threshold — the value above which inertial shielding is diagnosable as pathological rather than merely conservative — could be derived from the PCRB as the ratio of observed deformation to the admissible maximum $1 - e^{-\bar{S}_k}$.

Contact quantization and possibilistic field theory. The contact structure of Section 8 admits a quantization procedure analogous to geometric quantization of symplectic manifolds. Contact quantization replaces the contact Hamiltonian flow with a Schrödinger-type equation on sections of a contact line bundle. In the TEAG setting, this would correspond to a wave equation on the epistemic contact manifold, whose semiclassical limit recovers the tropical Hamilton–Jacobi equation of Theorem 5.5. Whether this quantization produces physically or epistemically meaningful objects — and whether the possibilistic Cramér–Rao Bound has a quantum analog in this framework — is an open question.

13.2. *Toward a Common Geometry of Constraint Propagation*

The results of this paper make available a synthesis that the individual sections approach but do not name.

The governing equations of physical dynamics and epistemic inference are structurally identical. In classical mechanics, a physical state $q \in M$ evolves under Hamilton–Jacobi: $\partial_t S + H(q, \nabla_q S) = 0$. In TEAG, a hypothesis $h \in H$ evolves under the tropical Hamilton–Jacobi equation: $\partial_k \Phi \oplus H_{\text{trop}}(h, \nabla_h \Phi) = 0$. The scalar fields S and Φ, the manifolds M and H, and the Hamiltonians have the same functional roles. The algebra differs — classical versus tropical — but the governing structure does not.

This alignment suggests a reading in which physical states and epistemic hypotheses are not different kinds of objects but different interpretations of points on a common manifold of admissible configurations. On this manifold, a scalar field evolves under constraint propagation. Physical laws impose dynamical constraints; observations impose epistemic constraints. Both enter as forcing terms in the same class of Hamilton–Jacobi equations. The distinction between "physical" and "epistemic" is then not geometric but interpretive: it marks which constraints are known and which remain uncertain.

Under this reading, the Lax–Oleinik update and Hamiltonian flow are operators on the same underlying manifold. The ESPF is not a method applied to physical dynamics from outside. It is constraint propagation operating on the same geometric object as the dynamics it tracks.

Remark 13.1 (What this does and does not claim)**.** *This subsection surfaces an implication of the governing equation identity established in Sections 5 and 12.3. It does not claim ontological identity of physical and epistemic space. The claim is geometric and structural: the same class of Hamilton–Jacobi equations governs both, the same algebra forces the update rule in both, and the same wavefront mechanism implements shielding in both. Whether this reflects a single underlying geometry or a deep structural analogy is a question the mathematics raises but does not settle. The companion programme (Jah 2026b) pursues the former possibility.*

13.3. *Hypothesis Space as Physical State Space*

The structural identity of Section 12.3 is not a coincidence of form. It has a precise explanation that requires no additional assumptions: in the ESPF, hypotheses already *are* physical states.

In orbit determination, the hypothesis space H is the space of possible orbital states $(r, \dot{r})$ — position and velocity vectors of a physical object. A hypothesis $h \in H$ is not an abstract proposition; it is a concrete physical configuration. The dynamics f_k in the ESPF predict step are the same equations of motion that govern physical trajectories in M. The measurement map $g : H \to \mathcal{Y}$ is the same function that maps physical states to observables. The hypothesis space is the physical state space:

$$H \equiv M = \{\text{all admissible physical configurations}\}. \tag{76}$$

Under this identification, the ESPF implements two types of constraint propagation on the scalar field $\Phi_\varnothing : M \to [0, \infty)$:

© 2026 by the author(s). Distributed under a Creative Commons CC BY license.

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 15 April 2026 doi:10.20944/preprints202603.2110.v3

(i) *Predict step.* Transport the impossibility field forward under physical dynamics: $\Phi_{\varnothing,k|k-1}(q) = \Phi_{\varnothing,k-1}(f_k^{-1}(q))$. This is free Hamiltonian flow on M.

(ii) *Update step.* Deform the impossibility field under the surprisal field generated by observation y_k: $\widetilde{\Phi}_\varnothing(q) = \max(\Phi_{\varnothing,k|k-1}(q), \Phi_{S,k}(q))$. This is the tropical Hamilton–Jacobi step.

Proposition 13.2 (ESPF as joint physical-epistemic constraint propagation)**.** *Let M be the physical state space and let $H \equiv M$ under the identification (76). The ESPF predict–update recursion implements two types of constraint propagation on the same scalar field $\Phi_\varnothing : M \to [0, \infty)$:*

$$\textit{Dynamical constraint:} \quad \Phi_{\varnothing,k|k-1}(q) = \Phi_{\varnothing,k-1}(f_k^{-1}(q)), \tag{77}$$

$$\textit{Epistemic constraint:} \quad \Phi^+_{\varnothing,k}(q) = \Phi_{\varnothing,k|k-1}(q) \oplus \Phi_{S,k}(q). \tag{78}$$

Both are Hamilton–Jacobi operators on M: the first is free field transport (zero Hamiltonian); the second is forced field evolution with Hamiltonian $H_{\text{trop}} = \Phi_{S,k}$. The distinction between dynamics and inference is not a distinction of geometry but of constraint type.

Proof. The dynamical constraint is the predict step of Theorem 5.5(i) with $h \equiv q \in M$. The epistemic constraint is the update step of Theorem 5.5(ii). Both follow from the tropical Hamilton–Jacobi theorem with the identification $H \equiv M$. ☐

Remark 13.3 (What changes and what does not)**.** *The identification $H \equiv M$ does not change any equation in the paper. It changes the interpretation: hypothesis space was always a space of physically meaningful configurations; this subsection makes that explicit. What remains open is whether the classical Hamilton–Jacobi equation of physical dynamics (+ algebra) and the tropical Hamilton–Jacobi equation of inference (*max *algebra) can be understood as the same equation at different temperatures on M — the $T \to 0$ limit of Theorem 10.1 applied to the dynamics themselves. That question is the companion programme (Jah 2026b).*

14. Conclusion

This paper has established the dynamical geometric foundation of the Epistemic Support-Point Filter and the Theory of Epistemic Abductive Geometry. The central claim is that the ESPF predict–update recursion is not a heuristic approximation to Bayesian inference under unusual assumptions: it is the exact solution operator of a tropical Hamilton–Jacobi equation, forced by the TEAG axioms without choice or approximation.

What was proved. The TEAG conjunctive update in log-admissibility coordinates is tropical addition of impossibility fields: the posterior field is the pointwise maximum of the prior impossibility and the surprisal. This algebraic identity is exact. Its level-set boundaries propagate as wavefronts governed by the tropical Hamilton–Jacobi equation, with the surprisal field as Hamiltonian. Both structures are forced: Popperian contraction forces max-plus, and the evidence-referencing axiom forces momentum independence. No alternative update consistent with the TEAG axioms exists.

Falsification has two-stage structure. The tropical variety of the posterior field — the active deformation front — marks where evidence begins to deform the impossibility landscape and is a necessary condition for falsification. Sufficient falsification requires exit from the PCRB-admissible basin, whose threshold is determined by the possibilistic information content of the observation. A hypothesis with low prior impossibility, earned through many observations, is protected from mild surprisal by the basin geometry. The minimax medioid is the fixed point of this basin: the last point to be falsified under any sequence of admissible contractions.

The PCRB is a minimum action principle. No measurement can compress the surviving well below the PCRB floor per update; the ESPF achieves this floor with equality at the admissibility boundary. In the variational language, the ESPF selects the path of least unjustified epistemic deformation: it commits no further impossibility to the surviving support than the evidence demands.

© 2026 by the author(s). Distributed under a Creative Commons CC BY license.

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 15 April 2026 doi:10.20944/preprints202603.2110.v3

The correct geometric primitive of epistemic phase space is a contact manifold, not a symplectic one. Symplectic geometry governs reversible, volume-preserving Hamiltonian systems; TEAG is neither, by axiom. The contact manifold (h, p, Φ) — hypothesis, impossibility gradient, accumulated action — carries a contact form whose Reeb dynamics reproduce the TEAG update exactly. The PCRB is a contact energy floor. The ESPF implements contact Hamiltonian flow with discrete projection onto the admissible basin at each regeneration step. The zero-temperature limit, which recovers probability theory from possibility theory, is a contact phase transition: both frameworks are contact Hamiltonian systems differing only in temperature.

What this gives TEAG. The wavefront formulation supplies three things the main TEAG paper identifies as structural summaries pending derivation. First, the tropical Hamiltonian is named explicitly and proved to be the surprisal field by axiomatic necessity. Second, the full predict–update–regenerate cycle is given a unified geometric interpretation as contact Hamiltonian flow with projection. Third, the medioid is proved to be the geodesic attractor of the surviving well — its role as commitment point is not a heuristic but a theorem about the geometry of the admissible basin.

The epistemic posture. Every structure in this paper descends from a single asymmetry: evidence can only increase impossibility, never decrease it. That asymmetry forces max-plus algebra, forces the contact structure, forces one-sided wavefront propagation, and forces the irreversibility of Popperian falsification. The filter is quick to embrace ignorance and slow to assert certainty not as a design philosophy but as a mathematical consequence of operating in a world where falsification is permanent.

Appendix A. Epistemic Horizons, Causal Structure, and Black Hole Analogues

This appendix extends the Hamilton–Jacobi formulation of TEAG by introducing an epistemic causal structure on hypothesis space, together with the associated notions of horizons, support collapse, and black-hole analogues. The central claim is that the causal metric of the epistemic manifold emerges naturally from the action functional already defined in the main text.

Appendix A.1. Epistemic Support Collapse and Black Holes

Let H be a hypothesis space equipped with a possibility distribution $\pi : H \to (0, 1]$ and corresponding impossibility field

$$\Phi_\varnothing(h) = -\log \pi(h). \tag{A79}$$

Definition A1 (Epistemic support)**.** *The epistemic support is the set*

$$\mathrm{supp}(\pi) = \{h \in H : \pi(h) > 0\}. \tag{A80}$$

Definition A2 (Epistemic black hole)**.** *An epistemic black hole is a subset $\mathcal{B} \subset H$ such that*

1. $\mathrm{supp}(\pi) = \mathcal{B}$*,*
2. $\mathrm{Vol}(\mathcal{B}) \to 0$ *(or* $|\mathcal{B}| = 1$ *in the discrete case),*
3. *no admissible expansion operator restores support outside* $\mathcal{B}$*.*

In this regime,

$$\Phi_\varnothing(h^\star) = 0, \qquad \Phi_\varnothing(h \neq h^\star) = \infty, \tag{A81}$$

and the posterior update

$$\Phi^e_\varnothing(h) = \max\big(\Phi_\varnothing(h), \Phi_S(h)\big) \tag{A82}$$

cannot generate new admissible hypotheses. Falsification becomes undefined rather than merely unlikely, because no alternate admissible directions in hypothesis space remain.

© 2026 by the author(s). Distributed under a Creative Commons CC BY license.

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 15 April 2026 doi:10.20944/preprints202603.2110.v3

Appendix A.2. Active Deformation Front and Epistemic Horizon

Recall that the active deformation front is

$$\mathcal{F}_k = \{h \in H : \Phi_\varnothing(h) = \Phi_S(h)\}. \tag{A83}$$

Definition A3 (Epistemic horizon)**.** *An epistemic horizon is a hypersurface $\mathcal{H} \subset H$ such that for all $h \in \mathcal{H}$,*

$$\|\nabla\Phi_\varnothing(h)\|_W \geq \|\nabla\Phi_S(h)\|_W, \tag{A84}$$

with equality on the boundary and strict dominance inward, so that no admissible surprisal wavefront can alter the topology of the surviving support beyond $\mathcal{H}$.

Thus, $\mathcal{H}$ is the epistemic analogue of an event horizon: once the trajectory of the filter enters the region where the prior impossibility geometry dominates all incoming evidence, the support remains trapped unless regeneration injects new admissible volume.

Appendix A.3. The Causal Metric from the Epistemic Action

In the main text, the epistemic Lagrangian is defined as

$$L_k(h, \dot{h}) = T_k - V(h) = \Delta E_{\pi,k} - \Phi_\varnothing(h), \tag{A85}$$

with epistemic action

$$S[\gamma] = \sum_{k=1}^{K} L_k. \tag{A86}$$

To expose the induced geometry, let the hypothesis trajectory be represented in MVEE-whitened coordinates, so that the local metric is

$$d\ell_W^2 = dh^\top G_W(h)\, dh, \tag{A87}$$

where G_W is the MVEE-whitened metric tensor. Assume that the entropy reduction per step is generated by a local epistemic speed

$$\dot{E}_\pi = \frac{dE_\pi}{dk}. \tag{A88}$$

The PCRB implies the causal inequality

$$\Delta E_{\pi,k} \leq I_k, \tag{A89}$$

where I_k is the Choquet information content of the observation. In the continuum limit, this becomes

$$\dot{E}_\pi \leq c_{\text{epi}}, \tag{A90}$$

where

$$c_{\text{epi}} := \sup_k I_k \tag{A91}$$

is the maximal epistemic contraction rate admissible under the PCRB.

We now define the quadratic epistemic action density

$$\mathcal{L}(h, \dot{h}, \dot{E}_\pi) = \frac{1}{2}\|\dot{h}\|_W^2 - \frac{1}{2c_{\text{epi}}^2}\dot{E}_\pi^2 - \Phi_\varnothing(h). \tag{A92}$$

This is the second-order kinetic refinement of the discrete action in the main text: the first term measures epistemic displacement through hypothesis space, the second measures entropy-contraction effort relative to the PCRB speed limit, and the third is the impossibility potential.

© 2026 by the author(s). Distributed under a Creative Commons CC BY license.

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 15 April 2026 doi:10.20944/preprints202603.2110.v3

The action is therefore

$$S[\gamma] = \int \mathcal{L}\, dk = \frac{1}{2} \int \left(\|\dot{h}\|_W^2 - \frac{1}{c_{\text{epi}}^2} \dot{E}_\pi^2 \right) dk - \int \Phi_\varnothing(h)\, dk. \tag{A93}$$

The kinetic part defines the quadratic form

$$ds^2 = d\ell_W^2 - \frac{1}{c_{\text{epi}}^2} (dE_\pi)^2. \tag{A94}$$

Thus, the epistemic line element emerges directly from the action: it is the indefinite metric whose null cone is set by the maximum contraction rate permitted by the PCRB.

Proposition A4 (Epistemic line element from the action). *Let $\gamma(k) = (h(k), E_\pi(k))$ be an epistemic trajectory and let the quadratic action density be given by (A92). Then the kinetic part of the action induces the pseudo-Riemannian line element*

$$ds^2 = d\ell_W^2 - \frac{1}{c_{\text{epi}}^2} (dE_\pi)^2. \tag{A95}$$

Moreover, null trajectories $ds^2 = 0$ are precisely those that saturate the PCRB speed limit.

Proof. The kinetic part of (A92) is the quadratic form

$$\mathcal{T} = \frac{1}{2} \|\dot{h}\|_W^2 - \frac{1}{2c_{\text{epi}}^2} \dot{E}_\pi^2. \tag{A96}$$

Multiplying by dk^2 yields

$$2\mathcal{T}\, dk^2 = dh^\top G_W(h)\, dh - \frac{1}{c_{\text{epi}}^2} (dE_\pi)^2, \tag{A97}$$

which is exactly (A94). The null condition $ds^2 = 0$ gives

$$\|\dot{h}\|_W^2 = \frac{1}{c_{\text{epi}}^2} \dot{E}_\pi^2, \tag{A98}$$

so the entropy-contraction rate exactly matches the maximal propagation rate allowed by the causal bound. Hence null trajectories are the epistemic analogue of lightlike propagation. □

Definition A5 (Epistemic causality). *A trajectory γ is called*

- causal *if $ds^2 \geq 0$,*
- null *if $ds^2 = 0$,*
- spacelike *if $ds^2 < 0$.*

Spacelike transitions are inadmissible: they require entropy contraction faster than the PCRB permits.

Appendix A.4. Hamilton–Jacobi Interpretation of the Horizon

The tropical Hamilton–Jacobi equation in the main text is

$$\Phi_{\varnothing,k}^{e}(h) = \Phi_{\varnothing,k|k-1}(h) \oplus H_{\text{trop}}(h, \nabla_h \Phi_{\varnothing,k|k-1}), \tag{A99}$$

with

$$H_{\text{trop}}(h, p) = \Phi_{S,k}(h). \tag{A100}$$

© 2026 by the author(s). Distributed under a Creative Commons CC BY license.

In the quadratic action picture above, the canonical momentum is

$$p = \frac{\partial \mathcal{L}}{\partial \dot{h}} = G_W(h)\dot{h}, \tag{A101}$$

and the entropy-conjugate momentum is

$$p_E = \frac{\partial \mathcal{L}}{\partial \dot{E}_\pi} = -\frac{1}{c_{\text{epi}}^2}\dot{E}_\pi. \tag{A102}$$

Hence the characteristic relation induced by the kinetic term is

$$\|p\|^2_{W^{-1}} - c_{\text{epi}}^2 p_E^2 = 0 \tag{A103}$$

for null propagation. This is the causal shell condition of the epistemic manifold. The active deformation front propagates along characteristics until the prior potential $\Phi_\varnothing$ bends those characteristics so strongly that the front can no longer escape. That trapped region is the epistemic horizon.

Appendix A.5. Belief Inertia and Epistemic Mass

The resistance of a hypothesis to displacement under incoming evidence can be formalized in terms of local curvature.

Definition A6 (Local epistemic mass)**.** *The local epistemic mass tensor is*

$$M(h) := \nabla_h^2 \Phi_\varnothing(h), \tag{A104}$$

computed in whitened coordinates.

Large positive curvature corresponds to steep confinement of the support cloud around h. In particular, if $\lambda_{\min}(M(h))$ is large, then any small displacement away from h incurs a large increase in impossibility. This formalizes belief inertia as geometric stiffness of the impossibility field.

Appendix A.6. PCRB Curvature and Horizon Formation

The PCRB defines an epistemic geoid as an equipotential or equal-action boundary. In the local quadratic approximation about a surviving support point $h^\star$, write

$$\Phi_\varnothing(h) = \Phi_\varnothing(h^\star) + \frac{1}{2}(h - h^\star)^\top M(h^\star)(h - h^\star) + o(\|h - h^\star\|^2), \tag{A105}$$

where $M(h^\star) = \nabla^2 \Phi_\varnothing(h^\star)$.

Let ρ_{PCRB} denote the characteristic PCRB radius in the whitened metric, defined implicitly by the equipotential condition

$$\Phi_\varnothing(h) - \Phi_\varnothing(h^\star) \le \rho_{\text{PCRB}} \qquad \text{for admissible survivors.} \tag{A106}$$

Proposition A7 (PCRB curvature criterion for horizon formation)**.** *Let $h^\star$ be a surviving support point and suppose the impossibility field is twice continuously differentiable in a neighborhood of $h^\star$. If*

$$\lambda_{\min}\big(\nabla^2 \Phi_\varnothing(h^\star)\big)\, \rho_{\text{PCRB}}^2 \;\ge\; 2\, c_{\text{epi}}, \tag{A107}$$

then an epistemic horizon forms around $h^\star$: any incoming surprisal wavefront consistent with the PCRB speed limit is unable to penetrate beyond the corresponding PCRB geoid in a single admissible update.

© 2026 by the author(s). Distributed under a Creative Commons CC BY license.

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 15 April 2026 doi:10.20944/preprints202603.2110.v3

Proof. Under the quadratic expansion,

$$\Phi_\varnothing(h) - \Phi_\varnothing(h^\star) \geq \frac{1}{2}\lambda_{\min}(M(h^\star))\|h - h^\star\|_W^2. \tag{A108}$$

At the PCRB radius $\|h - h^\star\|_W = \rho_{\text{PCRB}}$, the minimal potential rise is therefore

$$\Delta\Phi_\varnothing^{\min} \geq \frac{1}{2}\lambda_{\min}(M(h^\star))\rho_{\text{PCRB}}^2. \tag{A109}$$

An admissible surprisal wavefront can change the epistemic state no faster than c_{epi} per update. If

$$\frac{1}{2}\lambda_{\min}(M(h^\star))\rho_{\text{PCRB}}^2 \geq c_{\text{epi}}, \tag{A110}$$

then the inward potential barrier at the PCRB geoid is at least as large as the maximal admissible surprisal action. Thus the active deformation front cannot cross the boundary in a single admissible update. Since the tropical update is monotone and cannot create support behind an infinite or overcritical barrier, the region enclosed by the geoid is trapped under the update. This is precisely horizon formation. □

Remark A8. *Proposition A7 states that horizon formation is governed by the product of two quantities:*

1. *local curvature of the impossibility field, through $\lambda_{\min}(\nabla^2\Phi_\varnothing)$,*
2. *admissible contraction scale, through ρ_{PCRB}.*

High curvature and small PCRB radius correspond to tightly confined, difficult-to-falsify belief structures. In the limit that the support collapses to a singleton and the curvature diverges, the horizon closes into an epistemic black hole.

Appendix A.7. Topological Falsifiability

Definition A9 (Epistemic connectivity)**.** *A hypothesis manifold H is epistemically connected if for any $h_1, h_2 \in H$ there exists a continuous admissible path $\gamma : [0,1] \to H$ such that*

$$\Phi_\varnothing(\gamma(t)) < \infty \qquad \forall t \in [0,1]. \tag{A111}$$

Proposition A10 (Geometric criterion for scientific inference)**.** *Falsification is possible only if the admissible hypothesis manifold is epistemically connected and free of complete support collapse.*

This yields a precise geometric version of the Popperian demarcation criterion: science requires that beliefs remain reachable by admissible evidence-driven trajectories. When the impossibility field disconnects the manifold or collapses its support to a singularity, falsification ceases to be well-defined.

Appendix A.8. Interpretation

The constructions above establish the following.

- The quadratic refinement of the epistemic action induces a pseudo-Riemannian metric on (H, E_π).
- The PCRB defines the maximum admissible entropy-contraction rate c_{epi}, which serves as the epistemic speed of light.
- Null characteristics of the Hamilton–Jacobi action correspond to maximally propagating falsification fronts.
- Large local curvature of the impossibility field relative to the PCRB radius causes horizon formation.
- Complete support collapse is the epistemic analogue of a black hole.

Thus, the causal and topological structure of TEAG is not metaphorical decoration. It is already latent in the action principle, the tropical Hamilton–Jacobi recursion, and the PCRB geometry.

© 2026 by the author(s). Distributed under a Creative Commons CC BY license.

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 15 April 2026 doi:10.20944/preprints202603.2110.v3

Appendix A.9. The Causal Metric, Geodesics, and the Levi-Civita Connection

We now make explicit the differential-geometric structure induced by the epistemic action. Let the extended epistemic manifold be coordinatized by

$$x^\mu = (h^1, \ldots, h^n, \mathcal{E}), \qquad \mathcal{E} := E_\pi, \tag{A112}$$

where $(h^1, \ldots, h^n)$ are local coordinates on hypothesis space and $\mathcal{E}$ is the entropy coordinate.

The epistemic line element introduced above is

$$ds^2 = g_{\mu\nu}(x)\, dx^\mu dx^\nu, \tag{A113}$$

with block-diagonal metric

$$g_{\mu\nu}(x) = \begin{pmatrix} (G_W)_{ij}(h) & 0 \\ 0 & -c_{\text{epi}}^{-2} \end{pmatrix}. \tag{A114}$$

Here $(G_W)_{ij}(h)$ is the MVEE-whitened metric on hypothesis space, and c_{epi} is the maximal admissible entropy-contraction rate induced by the PCRB. In the present derivation, we treat c_{epi} as constant, so the only position dependence of the metric lies in the hypothesis-space block.

Remark A11 (Semi-definiteness and degenerate directions)**.** *The hypothesis-space block $(G_W)_{ij}(h)$ is the pullback of the MVEE-whitened inner product through the measurement map $g : H \to \mathbb{R}^m$. It is positive definite if and only if g is injective in a neighborhood of h — that is, if distinct hypotheses produce distinct predicted measurements. When g has a non-trivial kernel (multiple hypotheses map to the same measurement), $(G_W)_{ij}$ is only positive* semi-*definite, and directions in the kernel of dg are metrically degenerate: they carry zero whitened length.*

Geometrically, these degenerate directions correspond to hypotheses that are observationally equivalent *— no available measurement can distinguish them. In the ESPF, such hypotheses are never falsified by the surprisal field (since $\Phi_S(h) = \frac{1}{2}\|L_e^{-1}(y - g(h))\|^2$ is constant along kernel directions), and their impossibility field values evolve only through prior history, not through evidence. This is the geometric signature of* identifiability failure*: the measurement model cannot resolve the degeneracy.*

The extended metric $g_{\mu\nu}$ on $(H, \mathcal{E})$ is therefore pseudo-Riemannian in two senses: it has a timelike direction (the $-c_{\text{epi}}^{-2}$ entropy block) and may have additional degenerate spacelike directions from the kernel of dg. The geodesic equations and horizon formation results of this appendix hold on the non-degenerate submanifold; on degenerate directions, dynamics reduce to free transport without evidence-driven deformation.

To identify the natural trajectories of this geometry, consider first the free kinetic action

$$S_{\text{geo}}[\gamma] = \frac{1}{2} \int g_{\mu\nu}(x)\, \dot{x}^\mu \dot{x}^\nu \, d\lambda, \tag{A115}$$

where λ is an affine parameter along the trajectory and dots denote $d/d\lambda$. The corresponding Lagrangian is

$$\mathcal{L}_{\text{geo}} = \frac{1}{2} g_{\mu\nu}(x)\, \dot{x}^\mu \dot{x}^\nu. \tag{A116}$$

Applying the Euler–Lagrange equations,

$$\frac{d}{d\lambda}\left(\frac{\partial \mathcal{L}_{\text{geo}}}{\partial \dot{x}^\rho}\right) - \frac{\partial \mathcal{L}_{\text{geo}}}{\partial x^\rho} = 0, \tag{A117}$$

gives

$$\frac{d}{d\lambda}(g_{\rho\nu}\dot{x}^\nu) - \frac{1}{2}\partial_\rho g_{\mu\nu}\dot{x}^\mu \dot{x}^\nu = 0. \tag{A118}$$

© 2026 by the author(s). Distributed under a Creative Commons CC BY license.

Expanding the total derivative and multiplying by $g^{\alpha\rho}$ yields

$$\ddot{x}^{\alpha} + \Gamma^{\alpha}_{\mu\nu}\dot{x}^{\mu}\dot{x}^{\nu} = 0, \tag{A119}$$

where

$$\Gamma^{\alpha}_{\mu\nu} = \frac{1}{2}g^{\alpha\rho}(\partial_{\mu}g_{\rho\nu} + \partial_{\nu}g_{\rho\mu} - \partial_{\rho}g_{\mu\nu}) \tag{A120}$$

are the Christoffel symbols of the Levi–Civita connection associated with the epistemic metric.

Thus, the Christoffel connection does not enter by assumption. It is forced by the action principle once the epistemic metric is specified. The connection encodes how the local geometry of the epistemic manifold bends free inference trajectories.

For the metric (A114), the only nonzero Christoffel symbols are those of the MVEE-whitened metric:

$$\Gamma^{i}_{jk} = \frac{1}{2}(G_W)^{i\ell}\Big(\partial_j(G_W)_{\ell k} + \partial_k(G_W)_{\ell j} - \partial_{\ell}(G_W)_{jk}\Big). \tag{A121}$$

Hence the free geodesic equations reduce to

$$\ddot{h}^{i} + \Gamma^{i}_{jk}(h)\dot{h}^{j}\dot{h}^{k} = 0, \tag{A122}$$

together with

$$\ddot{\mathcal{E}} = 0. \tag{A123}$$

Equation (A122) shows that, absent any additional potential, epistemic trajectories follow geodesics of the whitened hypothesis-space geometry. Curvature of the epistemic manifold therefore arises from spatial variation in the MVEE-whitened metric itself.

Remark A12 (When Christoffel symbols vanish)**.** *When the MVEE-whitened metric $(G_W)_{ij}(h)$ is constant over hypothesis space — which occurs when the measurement map g is linear and the MVEE geometry does not vary with h — all partial derivatives $\partial_k(G_W)_{ij} = 0$, and by (A121) all Christoffel symbols vanish: $\Gamma^{i}_{jk} = 0$.*

In this regime the geodesic equation (A122) reduces to $\ddot{h}^{i} = 0$ — straight-line motion in the whitened coordinates. This is not a trivialization of the geometry; it means that in the MVEE-whitened coordinate system, hypothesis space is locally flat. The curvature lives in the variation *of G_W with h: when g is nonlinear, the MVEE shape changes across the support cloud, and the Christoffel terms are non-zero and carry genuine geometric content.*

The practical consequence: the ESPF operates in a regime where G_W is recomputed at each update step from the current support cloud. Between updates, within a single step, G_W is held constant, and Christoffel terms vanish within the step. Across steps, the changing G_W encodes the evolving curvature of hypothesis space as the admissible support contracts. The geometry of hypothesis space is not chosen freely; it is read off the measurement structure at each step. The universe gives us shape.

The full epistemic dynamics are obtained by reintroducing the impossibility field as a potential. Define the forced epistemic action

$$S[\gamma] = \int \left[\frac{1}{2}g_{\mu\nu}(x)\dot{x}^{\mu}\dot{x}^{\nu} - \Phi_{\varnothing}(h)\right]d\lambda. \tag{A124}$$

The Euler–Lagrange equations then become

$$\ddot{x}^{\alpha} + \Gamma^{\alpha}_{\mu\nu}\dot{x}^{\mu}\dot{x}^{\nu} = -g^{\alpha\beta}\partial_{\beta}\Phi_{\varnothing}. \tag{A125}$$

Since $\Phi_{\varnothing}$ depends only on the hypothesis coordinates, this yields

$$\ddot{h}^{i} + \Gamma^{i}_{jk}(h)\dot{h}^{j}\dot{h}^{k} = -(G_W)^{ij}\partial_j\Phi_{\varnothing}, \tag{A126}$$

© 2026 by the author(s). Distributed under a Creative Commons CC BY license.

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 15 April 2026 doi:10.20944/preprints202603.2110.v3

and

$$\ddot{\mathcal{E}} = 0. \tag{A127}$$

Equation (A126) is the epistemic analog of Newton's equation on a curved manifold: the Christoffel term represents inertial motion induced by the geometry of admissible inference, while the gradient of the impossibility field acts as the epistemic force generated by prior structure and accumulated history.

Appendix A.10. Fields, Geometry, and the Ontology of the Impossibility Field

The metric $G_W(h)$ encodes the shape of hypothesis space — the geometry of what is measurable. The impossibility field $\Phi_\varnothing(h) = -\log \pi(h)$ encodes the epistemic state — the history of what has been admitted or falsified. These are structurally different objects.

The geometry is not chosen. The MVEE-whitened metric $G_W(h)$ is determined entirely by the measurement map g and the geometry of the predicted measurement support cloud. It is not a modeling choice; it is read off the physical structure of what observations are available. When the TEAG axioms specify that the metric is the MVEE-whitened inner product, they are identifying the unique geometry under which the surprisal field is a squared-distance function. The measurement structure of the world determines the shape of hypothesis space. The universe gives us shape.

The impossibility field is not freely chosen either. The impossibility field $\Phi_\varnothing = -\log \pi$ is not an independently specified field placed on a fixed background geometry. It is the log-admissibility coordinate of the epistemic state — determined uniquely and completely by the possibility measure π, which is determined by the history of evidence under the TEAG contraction axioms. The field is not a choice; it is a consequence.

TEAG operates in and on an epistemic universe of discourse. The epistemic manifold $\mathcal{M}_{\text{ep}}$ is not a fixed background on which inference takes place. It is the universe of discourse itself — the space of all admissible configurations. Its geometry is constant only in the presence of uninformative evidence ($\bar{S}_k = 0$). Evidence is what deforms it. Each informative measurement ($\bar{S}_k > 0$) contracts the admissible support from S_k^- to S_k^+, expelling hypotheses irreversibly and shrinking the manifold. The impossibility field $\Phi_\varnothing$ marks the shape of the surviving geometry from the inside: it tells how close any hypothesis is to the boundary of the surviving space, in log-admissibility units.

The forced geodesic equation. Equation (A126) describes how a committed epistemic state moves through a geometry that is itself being shaped by the history of falsification. The Christoffel term $\Gamma^i_{jk}\dot{h}^j\dot{h}^k$ is inertial motion through the surviving geometry in the absence of evidence. The term $-(G_W)^{ij}\partial_j\Phi_\varnothing$ is the gradient of the boundary marker, pointing toward the interior. The metric is forced by the measurement structure. The dynamics are forced by the TEAG axioms. Neither is freely chosen.

Appendix A.11. Information as Geometric Breaking

The considerations above lead to a foundational theorem about the relationship between information, geometry, and dynamics in TEAG.

Theorem A13 (Geometry conservation iff zero information)**.** *Let $\mathcal{M}_{\text{ep}}$ be the epistemic contact manifold and let $\bar{S}_k$ be the aggregate epistemic surprisal at step k (Theorem 6.1). Then:*

- *(i) The geometry of $\mathcal{M}_{\text{ep}}$ is preserved under the TEAG update at step k if and only if $\bar{S}_k = 0$.*
- *(ii) When $\bar{S}_k = 0$ for all k, the contact structure degenerates to a symplectic structure: the contact dissipation term vanishes, Liouville's theorem is restored, and the dynamics become reversible and volume-preserving.*
- *(iii) When $\bar{S}_k > 0$, the update is a retraction, not a diffeomorphism: $H_{\text{surv}} \subsetneq H$, the geometry breaks irreversibly, and contact structure is the irreducible geometric primitive.*

Proof. *(i).* Geometry is preserved iff the update is a diffeomorphism on the admissible support — iff no hypotheses are expelled. No hypotheses are expelled iff $\widetilde{\Phi}_\varnothing(h) = \max(\Phi_\varnothing(h), \Phi_S(h)) = \Phi_\varnothing(h)$ for

© 2026 by the author(s). Distributed under a Creative Commons CC BY license.

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 15 April 2026 doi:10.20944/preprints202603.2110.v3

all h in the support, i.e., iff $\Phi_S(h) \le \Phi_\varnothing(h)$ everywhere. By Theorem 6.1, this condition is equivalent to $\bar{S}_k = \sup_i \min(\frac{1}{2}q_k^{(i)}, \pi_{k|k-1}^{(i)}) = 0$: no hypothesis is simultaneously credible and surprised.

(ii). When $\bar{S}_k = 0$ for all k, the active deformation front $\mathcal{B}_{\text{active}}$ is empty at every step. The contact Hamiltonian $\mathcal{H}_{\text{ep}} = \Phi_S(h) - \Phi$ satisfies $\mathcal{H}_{\text{ep}} \le 0$ everywhere, and the contact dissipation term $-p\,\partial_\Phi \mathcal{H}_{\text{ep}} = +p$ vanishes in the surviving interior. The contact form $\eta = d\Phi - p \cdot dh$ restricted to the $\{\Phi = \text{const}\}$ level sets reduces to the symplectic form $\omega = dh \wedge dp$. Liouville's theorem holds on these level sets, and the dynamics are reversible.

(iii). When $\bar{S}_k > 0$, there exists at least one hypothesis that is both credible ($\pi^{(i)} > 0$) and sufficiently surprised ($\Phi_S > \Phi_\varnothing$) to be expelled. The update map $H \to H_{\text{surv}}$ is not surjective onto H, hence not a diffeomorphism. By Axiom A3 (non-resurrection), this contraction is irreversible. The contact dissipation term is non-zero, Liouville's theorem fails, and the irreducible geometric primitive is the contact manifold. □

Remark A14 (What the theorem says about information)**.** *Information, by definition, means something new. Something new means the space of what was previously possible must shrink — a hypothesis admitted before is no longer admitted now. A shrinking space is broken geometry. The theorem above is therefore the precise mathematical statement of this definition:*

Information is geometric breaking.

You cannot have genuine information without irreversible change to the admissible manifold. The aggregate surprisal $\bar{S}_k$ is not merely the epistemic action or the cognitive work budget — it is the rate of geometric destruction *per update: the measure of how much of the manifold is removed at step k.*

The theorem has a limit that gives contact and symplectic geometry their precise epistemic meanings. At one extreme, $\bar{S}_k > 0$: geometry breaks, the contact structure is necessary, falsification is real, and learning occurs. At the other extreme, $\bar{S}_k = 0$ permanently: no new information exists, geometry is static, the contact structure degenerates to symplectic, and the dynamics become reversible. A universe where everything is perfectly known has no geometry left to break. It converges to symplectic dynamics — the geometry of systems with nothing left to learn.

This is why TEAG, by design, breaks geometry. It is not a limitation or a pathology of the framework. It is the mathematical consequence of taking seriously what information means. The contact manifold is the geometry of learning. The symplectic manifold is the geometry of complete knowledge. The passage between them is governed by a single scalar: the Choquet aggregate surprisal.

Remark A15 (Evidence deformation is non-isentropic)**.** *When $\bar{S}_k > 0$, the update is not merely a deformation of the impossibility field — it is a non-isentropic transformation of the epistemic universe of discourse.*

Recall that an isentropic process is one that preserves phase-space volume (in the thermodynamic sense: entropy is constant, no heat is exchanged). By Theorem A13(iii), whenever $\bar{S}_k > 0$, Liouville's theorem fails: the update map $H \to H_{\text{surv}}$ is not volume-preserving. The admissible volume — measured by the possibilistic entropy $E_\pi = \int_0^1 \frac{1}{2} \log\det Q_\alpha \, d\alpha$ over α-cut MVEE shape matrices — decreases strictly.

The thermodynamic analog is precise: the epistemic universe of discourse acts as a system exchanging information with an external observer (the measurement apparatus). Each informative measurement ($\bar{S}_k > 0$) is a heat exchange: the observer injects surprisal into the field, which irreversibly contracts the admissible volume. The process is non-isentropic because the contraction cannot be undone (Axiom A3). No adiabatic reversal restores the expelled hypotheses.

The PCRB governs how much non-isentropy is admissible in a single step:

$$\Delta E_\pi \;\le\; \tfrac{n}{2}\log(1 - I_k), \tag{A128}$$

where $I_k = 1 - e^{-\bar{S}_k}$ is the information content and n is the hypothesis dimension. This is the TEAG analog of the Clausius inequality: the entropy reduction per update is bounded by the available epistemic action. A measurement cannot extract more geometric order from the epistemic universe than the evidence warrants.

© 2026 by the author(s). Distributed under a Creative Commons CC BY license.

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 15 April 2026 doi:10.20944/preprints202603.2110.v3

Three regimes summarize the full geometry:

(i) $\bar{S}_k = 0$ (uninformative): *isentropic, symplectic, reversible, volume-preserving.*

(ii) $\bar{S}_k > 0$, PCRB not saturated: *non-isentropic, contact, irreversible, volume contracts within the PCRB bound.*

(iii) $\bar{S}_k > 0$, PCRB saturated: *maximum admissible non-isentropy at this information content; contact dissipation at its bound.*

The geometry of the epistemic universe of discourse is constant only under uninformative evidence. Evidence is what deforms it, and that deformation is necessarily non-isentropic.

Remark A16 (Connection to the zero-temperature limit)**.** *Section 10 establishes that the zero-temperature limit $T \to 0$ recovers possibilistic (TEAG) dynamics from probabilistic (Bayesian) dynamics. Theorem A13 establishes a second limiting direction: the zero-information limit $\bar{S}_k \to 0$ recovers symplectic from contact dynamics. These two limits are orthogonal in the space of theories:*

	$T > 0$ *(probabilistic)*	$T = 0$ *(possibilistic)*
$\bar{S}_k > 0$ *(learning)*	*Bayesian contact*	*TEAG contact*
$\bar{S}_k = 0$ *(static)*	*Bayesian symplectic*	*TEAG symplectic*

Both Bayesian and TEAG systems collapse to their respective symplectic limits when information content vanishes. TEAG is the zero-temperature, geometry-breaking, contact-geometric theory of learning. Symplectic mechanics is the geometry of the perfectly known.

References

Fleming, W.H. and McEneaney, W.M. (2000). A max-plus based algorithm for an HJB equation of nonlinear filtering. *SIAM Journal on Control and Optimization*, 38(3):683–710.

Fleming, W.H. and Soner, H.M. (2006). *Controlled Markov Processes and Viscosity Solutions*, 2nd edition. Springer, New York.

Jah, M.K. (2026a). The Geometry of Knowing: From possibilistic ignorance to probabilistic certainty. Manuscript submitted for review, 2026.

Jah, M.K. (2026b). Theory of Epistemic Abductive Geometry (TEAG): A Unified Theory of Admissibility-Driven Inference Across Dynamical Systems, Measure Theory, and Language. Preprint, preprints.org, 2026 (submitted). doi:10.20944/preprints202603.2010.v1. https://www.preprints.org/manuscript/202603.2010.

Jah, M.K. and Haslett, V. (2025a). The Epistemic Support-Point Filter. *arXiv:2508.20806*.

Jah, M.K. (2025b). Optimality of the Epistemic Support-Point Filter. *arXiv:2603.10065*.

McEneaney, W.M. (2006). *Max-Plus Methods in Control and Estimation*. Birkhäuser, Boston.

Kolokoltsov, V.N. and Maslov, V.P. (1997). *Idempotent Analysis and Its Applications*. Kluwer Academic Publishers.

Maclagan, D. and Sturmfels, B. (2015). *Introduction to Tropical Geometry*. American Mathematical Society.

Geiges, H. (2008). *An Introduction to Contact Topology*. Cambridge University Press.

Bravetti, A., Cruz, H., and Tapias, D. (2017). Contact Hamiltonian mechanics. *Annals of Physics*, 376:17–39.

Bravetti, A. (2019). Contact geometry and thermodynamics. *International Journal of Geometric Methods in Modern Physics*, 16(supp01):1940003.

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

© 2026 by the author(s). Distributed under a Creative Commons CC BY license.