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Not peer-reviewed version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: 31 March 2026

doi: 10.20944/preprints202603.2110.v2

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

 doi:10.20944/preprints202603.2110.v2

Article

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

Moriba Kemessia Jah [1,2,3]

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

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 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 unique support point 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. 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 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

 doi:10.20944/preprints202603.2110.v2

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 and Haslett, 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 and Haslett, 2025b).

The zero-temperature limit derived here is also consistent with the Hölder-mean continuum developed in prior work (Jah and Haslett, 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 8.

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 7 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 8 establishes the zero-temperature limit connecting classical and tropical Hamilton–Jacobi, and situates the contact structure within thermodynamic geometry. Section 9 interprets the whitened minimax medioid as the geodesic attractor. Section 10 situates the results within TEAG and identifies open directions. Section 11 concludes.

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

 doi:10.20944/preprints202603.2110.v2

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 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 (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 1 (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 (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 3 (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 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 $\widetilde{\Phi}_\varnothing = \max(-\log\pi, -\log\mathrm{Comp}) = \max(\Phi_\varnothing, \Phi_S)$. □

 doi:10.20944/preprints202603.2110.v2

Remark 3 (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 10.*

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.

Definition 4 (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:*

$$\mathcal{F}_k = \{h \in H : \Phi_\varnothing(h) = \Phi_S(h)\}. \tag{6}$$

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 4 (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 5 (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* $\mathcal{A}_k = \{h : \widetilde{\Phi}_\varnothing(h) \leq c_k^\star\}$*, 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.*

 doi:10.20944/preprints202603.2110.v2

This two-stage structure is the correct formalization of Popper's criterion: $\mathcal{F}_k$ marks where falsification becomes possible; $\mathcal{B}_{\mathrm{adm},k} = \{h : \widetilde{\Phi}_\varnothing(h) = c_k^\star\}$ 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 well's geometric center. 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:

$$\Phi_\varnothing(h) = \frac{(h-h_0)^2}{2\sigma_0^2}, \tag{7}$$

$$\Phi_S(h) = \frac{(h-y)^2}{2\sigma_n^2}, \tag{8}$$

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.,

$$\frac{(h-h_0)^2}{\sigma_0^2} = \frac{(h-y)^2}{\sigma_n^2}. \tag{9}$$

Taking square roots with both signs yields two roots.

Proposition 1 (Scalar active deformation front). *The active deformation front $\mathcal{F} = \{h_-, h_+\}$ consists of:*

$$h_- = \frac{\sigma_n h_0 + \sigma_0 y}{\sigma_n + \sigma_0}, \tag{10}$$

$$h_+ = \frac{\sigma_n h_0 - \sigma_0 y}{\sigma_n - \sigma_0} \qquad (\sigma_n \neq \sigma_0). \tag{11}$$

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 (9), $\sigma_n(h - h_0) = \pm\sigma_0(h - y)$. 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 6 (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 $\sigma_n \to \infty$ (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 $\Phi_\varnothing(h) \leq \Phi_S(h)$.

Proposition 2 (Scalar surviving well width).

$$W = h_+ - h_- = \frac{2\sigma_n\sigma_0\,|h_0 - y|}{\sigma_n^2 - \sigma_0^2}. \tag{12}$$

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 (12): (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 3 (Scalar medioid). *The minimax medioid of the surviving interval $[h_-, h_+]$ is its midpoint:*

$$h_{\text{mid}} = \frac{h_+ + h_-}{2} = \frac{\sigma_n^2 h_0 - \sigma_0^2 y}{\sigma_n^2 - \sigma_0^2}. \tag{13}$$

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

 doi:10.20944/preprints202603.2110.v2

Remark 7 (The medioid 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 medioid is the geometric center of the surviving well — the point of minimum maximum impossibility among all surviving hypotheses — and carries no distributional interpretation. As* $\sigma_0 \to 0$, $h_{\text{mid}} \to h_0$*: the prior well collapses to its center and no observation can displace the commitment point. As* $\sigma_n \to \infty$, $h_{\text{mid}} \to h_0$ *again: an arbitrarily uninformative sensor cannot move the commitment point from the prior well center.*

4.5. The Zero-Temperature Limit in the Scalar Case

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

$$\widetilde{\Phi}_{\varnothing,T}(h) = -T \log\left(\exp\left(-\frac{(h - h_0)^2}{2T\sigma_0^2} \right) + \exp\left(-\frac{(h - y)^2}{2T\sigma_n^2} \right) \right). \tag{14}$$

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:

$$\widetilde{\Phi}_{\varnothing,T}(h) \xrightarrow{T\to 0} \max\big(\Phi_\varnothing(h),\, \Phi_S(h)\big) = \widetilde{\Phi}_\varnothing(h). \tag{15}$$

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 2 (proved in Section 8).

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:

$$\frac{\partial \Phi}{\partial k} \oplus H_{\text{trop}}(h,\, \nabla_h \Phi) = 0, \tag{16}$$

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, Proposition 7).

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

$$T_k \;=\; \Delta\mathcal{E}_{\pi,k} \;=\; \mathcal{E}_{\pi,k-1} - \mathcal{E}_{\pi,k}, \tag{17}$$

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

$$T_k \;\leq\; I_k, \tag{18}$$

where I_k is the Choquet information content of the observation (Proposition 7). 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.

The epistemic Lagrangian and action.

Definition 5 (Epistemic Lagrangian)**.** *The* epistemic Lagrangian *at update step k is*

$$L_k(h,\dot{h}) \;=\; T_k \;-\; V(h) \;=\; \Delta\mathcal{E}_{\pi,k} \;-\; \Phi_\varnothing(h), \tag{19}$$

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*

$$\mathcal{S} \;=\; \sum_{k=1}^{K} L_k. \tag{20}$$

Proposition 4 (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 (Proposition 7) 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, Definition 3). 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 $\mathcal{F}_k$ (Definition 4). □

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

$$p \;=\; \frac{\partial L_k}{\partial \dot{h}} \;=\; \nabla_h \Phi_\varnothing, \tag{21}$$

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

$$H_{\mathrm{trop}}(h,p) \;=\; p \cdot \dot{h} - L_k \;=\; \Phi_{S_k}(h), \tag{22}$$

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 31 March 2026 doi:10.20944/preprints202603.2110.v2

recovering the ESPF tropical Hamiltonian as the surprisal field (Definition 6 below). The momentum-independence of H_{trop} — noted in Remark 8 — 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.

Definition 6 (ESPF tropical Hamiltonian). *The* ESPF tropical Hamiltonian *is the map $H_{\text{trop}} : H \times T^*H \to \mathbb{R}$ defined by*

$$H_{\text{trop}}(h,\ p) \ = \ \Phi_{Sk}(h) \ = \ \tfrac{1}{2}\|z_k(h)\|^2, \tag{23}$$

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 8 (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 1 (ESPF as tropical Hamilton–Jacobi solver / Lax–Oleinik recursion). *With tropical Hamiltonian $H_{\text{trop}}(h, p) = \Phi_{Sk}(h)$ (Definition 6), the ESPF predict–update recursion satisfies the discrete tropical Hamilton–Jacobi equation*

$$\widetilde{\Phi}_{\varnothing,k}(h) = \Phi_{\varnothing k|k-1}(h) \oplus H_{\text{trop}}(h,\ \nabla_h \Phi_{\varnothing k|k-1}) \tag{24}$$

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*

$$\Phi_{\varnothing k|k-1}(h) = \Phi_{\varnothing k-1}(f_k^{-1}(h)), \tag{25}$$

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.

(ii) Update step (wavefront collision): *Substituting Definition 6 into Equation (24), and using the momentum-independence of H_{trop}:*

$$\widetilde{\Phi}_{\varnothing,k}(h) = \Phi_{\varnothing k|k-1}(h) \oplus \Phi_{Sk}(h) = \max\big(\Phi_{\varnothing k|k-1}(h),\ S_k(h)\big), \tag{26}$$

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 31 March 2026 doi:10.20944/preprints202603.2110.v2

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 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 (16) with $H_{\text{trop}} = 0$ (free evolution).

Update step. By Definition 6, $H_{\text{trop}}(h,p) = \Phi_{Sk}(h)$ independently of p. Substituting into Equation (24):

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

This is exactly the possibilistic conjunctive update $\pi'(h) = \min(\pi(h), \mathrm{Comp}(h))$ rewritten in potential coordinates via Definition 3. 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 9 (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 (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 $\Phi^+(h) = \Phi^-(h) \oplus$*

 doi:10.20944/preprints202603.2110.v2

$H_{\text{trop}}(h, \nabla_h \Phi^-)$ 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.

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 6 (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*

$$\frac{\partial \Phi}{\partial t}(h,t) \oplus \Phi_S(h,t) = 0, \tag{27}$$

where $\oplus = \max$ and $\Phi_S(h,t)$ is the surprisal field at time t. The unique viscosity solution of Equation (27) 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 10 (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 (27). 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) forces the Lax–Oleinik reduction. The viscosity solution theory of max-plus Hamilton–Jacobi equations (Fleming–McEneaney 2000; McEneaney 2006) guarantees that Equation (27) has a unique solution in this class, and that solution is the Lax–Oleinik semigroup stated in Proposition 6. 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.

 doi:10.20944/preprints202603.2110.v2

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 and Haslett, 2025b) bounds the minimum volume of the posterior support set:

$$\log\det(\mathrm{MVEE}(\text{posterior})) \geq \mathrm{PCRB}_k, \tag{28}$$

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 7 (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 $\pi_{k|k-1}$:*

$$\bar{S}_k \;=\; (\mathrm{C})\!\int \tfrac{1}{2}q_k^{(i)}\, d\pi_{k|k-1} \;=\; \sup_{i=1,\dots,N} \min\!\Big(\tfrac{1}{2}q_k^{(i)},\ \pi_{k|k-1}^{(i)}\Big), \tag{29}$$

where the second equality is the standard reduction of the Choquet integral under a possibility measure (Jah and Haslett, 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 7 (PCRB as variational minimum action)**.** *The PCRB bounds the possibilistic entropy reduction per update:*

$$E_{\pi,k|k} \;\geq\; E_{\pi,k|k-1} + \tfrac{n}{2}\log(1 - I_k), \tag{30}$$

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

$$\mathrm{PCRB}_k \;=\; \inf_{\pi' :\, \pi'(h) \leq \pi(h)\ \forall h} \bar{S}_k[\pi']. \tag{31}$$

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 (31) 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 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 (30) 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. 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.

7.1. 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.

7.2. The Epistemic Contact Manifold

Definition 8 (Epistemic phase space and contact form)**.** *The* epistemic phase space *is the* $(2n+1)$*-dimensional manifold*

$$\mathcal{M}_{\text{ep}} \;=\; \big\{(h,\, p,\, \Phi) \;:\; h \in H,\; p = \nabla_h \Phi_\varnothing \in T_h^* H,\; \Phi = \Phi_\varnothing(h) \in \mathbb{R}_{\geq 0}\big\}, \tag{32}$$

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 (21)), 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*

$$\eta \;=\; d\Phi \;-\; p \cdot dh \;=\; d\Phi_\varnothing - \nabla_h \Phi_\varnothing \cdot dh. \tag{33}$$

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 11 (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*

 doi:10.20944/preprints202603.2110.v2

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.

7.3. Contact Hamiltonian Flow

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

$$\eta(X_{\mathcal{H}}) = -\mathcal{H}, \qquad \iota_{X_{\mathcal{H}}} d\eta = d\mathcal{H} - (R\mathcal{H})\,\eta, \tag{34}$$

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

$$\dot{h} = \quad \nabla_p \mathcal{H}, \tag{35}$$
$$\dot{p} = -\nabla_h \mathcal{H} - p\,\partial_\Phi \mathcal{H}, \tag{36}$$
$$\dot{\Phi} = p \cdot \nabla_p \mathcal{H} - \mathcal{H}. \tag{37}$$

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

$$\mathcal{H}_{\mathrm{ep}}(h, p, \Phi) \;=\; \Phi_S(h) \;-\; \Phi, \tag{38}$$

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

The contact Hamiltonian equations Equation (35)–Equation (37) yield:

$$\dot{h} = 0, \tag{39}$$
$$\dot{p} = -\nabla_h \Phi_S(h) + p, \tag{40}$$
$$\dot{\Phi} = \Phi_S(h) - \Phi \;\cdot\; 0 - \mathcal{H} = \Phi - \Phi_S(h) + \Phi = \Phi - \Phi_S(h). \tag{41}$$

The fixed point of $\dot{\Phi} = 0$ is $\Phi = \Phi_S(h)$: exactly the active deformation front $\mathcal{B}_{\mathrm{active}}$ (Definition 4). 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 $\mathcal{H}_{\mathrm{ep}} = \Phi_S(h) - \Phi$ into Equation (35)–Equation (37): 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 Remark 8. 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 12 (Contact dissipation and Axiom A3)**.** *The $-p\,\partial_\Phi \mathcal{H}$ term in the contact momentum equation Equation (36) 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.*

 doi:10.20944/preprints202603.2110.v2

7.4. The PCRB as a Contact Energy Bound

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

$$\frac{d\mathcal{H}_{\mathrm{ep}}}{dk} \;=\; -(\partial_\Phi \mathcal{H})\,\mathcal{H} \;=\; \mathcal{H}_{\mathrm{ep}}, \tag{42}$$

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 9 (PCRB as contact energy floor)**.** *The PCRB inequality*

$$\mathcal{E}_{\pi,k|k} \;\geq\; \mathcal{E}_{\pi,k|k-1} + \tfrac{n}{2}\log(1 - I_k) \tag{43}$$

is the contact energy floor of the epistemic Hamiltonian system: it bounds the rate at which the contact Hamiltonian $\mathcal{H}_{\mathrm{ep}}$ 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 $\mathcal{B}_{\mathrm{adm},k}$ — the PCRB equipotential surface — where the contact energy is exactly dissipated and no further impossibility elevation is epistemically justified.

7.5. 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 medoid, resetting the impossibility field to zero over the new support. In contact geometry terms, this is a 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 medoid — 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 $\mathcal{H}_{\mathrm{ep}} = \Phi_S(h) - \Phi$: 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 13 (Symplectic geometry as a limit)**.** *When the PCRB floor is saturated — $\mathcal{E}_{\pi,k|k} - \mathcal{E}_{\pi,k|k-1} = \frac{n}{2}\log(1 - I_k)$ exactly — the contact dissipation term vanishes and the system instantaneously has symplectic character: the contact form η restricted to the $\{\Phi = \mathrm{const}\}$ level sets reduces to the symplectic form $\omega = dh \wedge dp$. In the ESPF this saturation never occurs in the nominal regime (zero saturated steps in the empirical diagnostics), confirming that the filter operates in the genuinely contact regime throughout. The symplectic limit corresponds to the Gaussian collapse limit of TEAG: when epistemic width $W \to 0$ the contact structure collapses to the symplectic structure of classical Hamiltonian mechanics on the Gaussian phase space.*

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 31 March 2026 doi:10.20944/preprints202603.2110.v2

Remark 14 (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 8 — 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.*

8. 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

$$F = -T \log Z = -T \log \int_H \exp\left(-\frac{\Phi(h)}{T}\right) d\mu(h). \tag{44}$$

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:

$$p(h|y) \propto \exp\left(-\frac{\Phi_S(h)}{T}\right) \cdot \exp\left(-\frac{\Phi_\varnothing(h)}{T}\right), \tag{45}$$

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.

$$-T \log\left(e^{-\Phi_\varnothing/T} + e^{-\Phi_S/T}\right) \xrightarrow{T \to 0} \max(\Phi_\varnothing, \Phi_S) = \widetilde{\Phi}_\varnothing. \tag{46}$$

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

Theorem 2 (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 15 (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 and Haslett (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. 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*

 doi:10.20944/preprints202603.2110.v2

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.*

9. The Minimax Medioid as Geodesic Attractor

We first make the metric structure explicit.

Definition 9 (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*

$$\langle h_1 - h_2, h_1 - h_2\rangle_W \;=\; (g(h_1) - g(h_2))^\top \Pi_e^{-1}(g(h_1) - g(h_2)), \tag{47}$$

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:

$$h^\star = \operatorname{argmin}_{h\in H_{\text{surv}}} \max_{\chi^{(j)}\in H_{\text{surv}}} \left\| h - \chi^{(j)} \right\|_W. \tag{48}$$

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 10 (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 10 (Medioid as 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. Specifically, $h^\star$ is the unique point in H_{surv} satisfying*

$$\max_{\chi^{(j)}\in H_{\text{surv}}} d_W(h^\star, \chi^{(j)}) \;\le\; \max_{\chi^{(j)}\in H_{\text{surv}}} d_W(h, \chi^{(j)}) \quad \forall\, h \in H_{\text{surv}}, \tag{49}$$

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 $\{h : \Phi_S(h) = c\}$

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 31 March 2026 doi:10.20944/preprints202603.2110.v2

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 sits at the geometric center of the survivor cloud in the metric that defines both the surprisal field and the geoid boundary — it is the support point most equidistant from that boundary across all directions. □

Remark 16 (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 1(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 (Proposition 7): commit to the geodesic attractor, but assign no further impossibility to the surviving support than the evidence demands.

10. Discussion and Open Directions

10.1. 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 and Haslett, 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 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 1), the PCRB as variational minimum action (Proposition 7), the zero-temperature limit connecting Bayesian and possibilistic inference

 doi:10.20944/preprints202603.2110.v2

(Theorem 2), the whitened metric anchoring of the medioid as geodesic attractor (Proposition 10), and the identification of epistemic phase space as a contact manifold (Section 7) 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 (Definition 6), 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.

10.2. Open directions

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 7, this is the question of whether the contact Hamiltonian flow is Lyapunov stable with respect to the admissible basin volume.

Tropical curvature and the PCRB geometric identity.

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. 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.

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.

Contact quantization and possibilistic field theory.

The contact structure of Section 7 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 1. 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.

11. 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–

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 31 March 2026 doi:10.20944/preprints202603.2110.v2

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 medoid 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.

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 medoid 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.

 doi:10.20944/preprints202603.2110.v2

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

$$\Phi_\varnothing(h) = -\log \pi(h). \tag{A50}$$

Definition A11 (Epistemic support)**.** *The epistemic support is the set*

$$\mathrm{supp}(\pi) = \{h \in H : \pi(h) > 0\}. \tag{A51}$$

Definition A12 (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,

$$\Phi_\varnothing(h^\star) = 0, \qquad \Phi_\varnothing(h \neq h^\star) = \infty, \tag{A52}$$

and the posterior update

$$\Phi^e_\varnothing(h) = \max\big(\Phi_\varnothing(h), \Phi_S(h)\big) \tag{A53}$$

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

Appendix A.2 Active Deformation Front and Epistemic Horizon

Recall that the active deformation front is

$$\mathcal{F}_k = \{h \in H : \Phi_\varnothing(h) = \Phi_S(h)\}. \tag{A54}$$

Definition A13 (Epistemic horizon)**.** *An epistemic horizon is a hypersurface $\mathcal{H} \subset H$ such that for all $h \in \mathcal{H}$,*

$$\|\nabla \Phi_\varnothing(h)\|_W \geq \|\nabla \Phi_S(h)\|_W, \tag{A55}$$

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

$$L_k(h, \dot{h}) = T_k - V(h) = \Delta E_{\pi,k} - \Phi_\varnothing(h), \tag{A56}$$

with epistemic action

$$S[\gamma] = \sum_{k=1}^{K} L_k. \tag{A57}$$

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

$$d\ell_W^2 = dh^\top G_W(h)\, dh, \tag{A58}$$

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 31 March 2026 doi:10.20944/preprints202603.2110.v2

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

$$\dot{E}_\pi = \frac{dE_\pi}{dk}. \tag{A59}$$

The PCRB implies the causal inequality

$$\Delta E_{\pi,k} \le I_k, \tag{A60}$$

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

$$\dot{E}_\pi \le c_{\text{epi}}, \tag{A61}$$

where

$$c_{\text{epi}} := \sup_k I_k \tag{A62}$$

is the maximal epistemic contraction rate admissible under the PCRB.

We now define the quadratic epistemic action density

$$\mathcal{L}(h, \dot{h}, \dot{E}_\pi) = \frac{1}{2}\|\dot{h}\|_W^2 - \frac{1}{2c_{\text{epi}}^2}\dot{E}_\pi^2 - \Phi_\varnothing(h). \tag{A63}$$

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.

The action is therefore

$$S[\gamma] = \int \mathcal{L}\,dk = \frac{1}{2}\int\left(\|\dot{h}\|_W^2 - \frac{1}{c_{\text{epi}}^2}\dot{E}_\pi^2\right)dk - \int \Phi_\varnothing(h)\,dk. \tag{A64}$$

The kinetic part defines the quadratic form

$$ds^2 = d\ell_W^2 - \frac{1}{c_{\text{epi}}^2}(dE_\pi)^2. \tag{A65}$$

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 A11 (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 (A63). Then the kinetic part of the action induces the pseudo-Riemannian line element*

$$ds^2 = d\ell_W^2 - \frac{1}{c_{\text{epi}}^2}(dE_\pi)^2. \tag{A66}$$

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

Proof. The kinetic part of (A63) is the quadratic form

$$\mathcal{T} = \frac{1}{2}\|\dot{h}\|_W^2 - \frac{1}{2c_{\text{epi}}^2}\dot{E}_\pi^2. \tag{A67}$$

Multiplying by dk^2 yields

$$2\mathcal{T}\,dk^2 = dh^\top G_W(h)\,dh - \frac{1}{c_{\text{epi}}^2}(dE_\pi)^2, \tag{A68}$$

which is exactly (A65). The null condition $ds^2 = 0$ gives

$$\|\dot{h}\|_W^2 = \frac{1}{c_{\text{epi}}^2}\dot{E}_\pi^2, \tag{A69}$$

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 A14 (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

$$\Phi^e_{\varnothing,k}(h) = \Phi_{\varnothing,k|k-1}(h) \oplus H_{\text{trop}}(h, \nabla_h \Phi_{\varnothing,k|k-1}), \tag{A70}$$

with

$$H_{\text{trop}}(h, p) = \Phi_{S,k}(h). \tag{A71}$$

In the quadratic action picture above, the canonical momentum is

$$p = \frac{\partial \mathcal{L}}{\partial \dot{h}} = G_W(h)\dot{h}, \tag{A72}$$

and the entropy-conjugate momentum is

$$p_E = \frac{\partial \mathcal{L}}{\partial \dot{E}_\pi} = -\frac{1}{c_{\text{epi}}^2}\dot{E}_\pi. \tag{A73}$$

Hence the characteristic relation induced by the kinetic term is

$$\|p\|_{W^{-1}}^2 - c_{\text{epi}}^2 p_E^2 = 0 \tag{A74}$$

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 A15 (Local epistemic mass). *The local epistemic mass tensor is*

$$M(h) := \nabla_h^2 \Phi_\varnothing(h), \tag{A75}$$

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.

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 31 March 2026 doi:10.20944/preprints202603.2110.v2

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

$$\Phi_\varnothing(h) = \Phi_\varnothing(h^\star) + \frac{1}{2}(h - h^\star)^\top M(h^\star)(h - h^\star) + o(\|h - h^\star\|^2), \tag{A76}$$

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

$$\Phi_\varnothing(h) - \Phi_\varnothing(h^\star) \le \rho_{\mathrm{PCRB}} \qquad \text{for admissible survivors.} \tag{A77}$$

Proposition A12 (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*

$$\lambda_{\min}\big(\nabla^2 \Phi_\varnothing(h^\star)\big)\, \rho_{\mathrm{PCRB}}^2 \;\ge\; 2\, c_{\mathrm{epi}}, \tag{A78}$$

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.

Proof. Under the quadratic expansion,

$$\Phi_\varnothing(h) - \Phi_\varnothing(h^\star) \ge \frac{1}{2}\lambda_{\min}(M(h^\star))\|h - h^\star\|_W^2. \tag{A79}$$

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

$$\Delta\Phi_\varnothing^{\min} \ge \frac{1}{2}\lambda_{\min}(M(h^\star))\rho_{\mathrm{PCRB}}^2. \tag{A80}$$

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

$$\frac{1}{2}\lambda_{\min}(M(h^\star))\rho_{\mathrm{PCRB}}^2 \ge c_{\mathrm{epi}}, \tag{A81}$$

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 A17. *Proposition A12 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 A16 (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*

$$\Phi_\varnothing(\gamma(t)) < \infty \qquad \forall t \in [0,1]. \tag{A82}$$

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 31 March 2026 doi:10.20944/preprints202603.2110.v2

Proposition A13 (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.

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

$$x^\mu = (h^1, \ldots, h^n, \mathcal{E}), \qquad \mathcal{E} := E_\pi, \tag{A83}$$

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

$$ds^2 = g_{\mu\nu}(x)\, dx^\mu dx^\nu, \tag{A84}$$

with block-diagonal metric

$$g_{\mu\nu}(x) = \begin{pmatrix} (G_W)_{ij}(h) & 0 \\ 0 & -c_{\mathrm{epi}}^{-2} \end{pmatrix}. \tag{A85}$$

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.

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

$$S_{\mathrm{geo}}[\gamma] = \frac{1}{2} \int g_{\mu\nu}(x)\, \dot{x}^\mu \dot{x}^\nu \, d\lambda, \tag{A86}$$

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

$$\mathcal{L}_{\mathrm{geo}} = \frac{1}{2} g_{\mu\nu}(x)\, \dot{x}^\mu \dot{x}^\nu. \tag{A87}$$

Applying the Euler–Lagrange equations,

$$\frac{d}{d\lambda}\left(\frac{\partial \mathcal{L}_{\mathrm{geo}}}{\partial \dot{x}^\rho}\right) - \frac{\partial \mathcal{L}_{\mathrm{geo}}}{\partial x^\rho} = 0, \tag{A88}$$

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 31 March 2026 doi:10.20944/preprints202603.2110.v2

gives

$$\frac{d}{d\lambda}(g_{\rho\nu}\dot{x}^{\nu}) - \frac{1}{2}\partial_{\rho}g_{\mu\nu}\dot{x}^{\mu}\dot{x}^{\nu} = 0. \tag{A89}$$

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

$$\ddot{x}^{\alpha} + \Gamma^{\alpha}_{\mu\nu}\dot{x}^{\mu}\dot{x}^{\nu} = 0, \tag{A90}$$

where

$$\Gamma^{\alpha}_{\mu\nu} = \frac{1}{2}g^{\alpha\rho}(\partial_{\mu}g_{\rho\nu} + \partial_{\nu}g_{\rho\mu} - \partial_{\rho}g_{\mu\nu}) \tag{A91}$$

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 (A85), the only nonzero Christoffel symbols are those of the MVEE-whitened metric:

$$\Gamma^{i}_{jk} = \frac{1}{2}(G_W)^{i\ell}\Big(\partial_j(G_W)_{\ell k} + \partial_k(G_W)_{\ell j} - \partial_{\ell}(G_W)_{jk}\Big). \tag{A92}$$

Hence the free geodesic equations reduce to

$$\ddot{h}^{i} + \Gamma^{i}_{jk}(h)\dot{h}^{j}\dot{h}^{k} = 0, \tag{A93}$$

together with

$$\ddot{\mathcal{E}} = 0. \tag{A94}$$

Equation (A93) 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.

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

$$S[\gamma] = \int \left[\frac{1}{2}g_{\mu\nu}(x)\dot{x}^{\mu}\dot{x}^{\nu} - \Phi_{\varnothing}(h)\right]d\lambda. \tag{A95}$$

The Euler–Lagrange equations then become

$$\ddot{x}^{\alpha} + \Gamma^{\alpha}_{\mu\nu}\dot{x}^{\mu}\dot{x}^{\nu} = -g^{\alpha\beta}\partial_{\beta}\Phi_{\varnothing}. \tag{A96}$$

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

$$\ddot{h}^{i} + \Gamma^{i}_{jk}(h)\dot{h}^{j}\dot{h}^{k} = -(G_W)^{ij}\partial_j\Phi_{\varnothing}, \tag{A97}$$

and

$$\ddot{\mathcal{E}} = 0. \tag{A98}$$

Equation (A97) 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.

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.

 doi:10.20944/preprints202603.2110.v2

Jah, M.K. (2026a). The Geometry of Knowing: From possibilistic ignorance to probabilistic certainty. Preprint, arXiv:submit/7362363, 2026 (under review).

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. and Haslett, V. (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.