

Article

Not peer-reviewed version

Theory of Epistemic Abductive Geometry (TEAG): A Unified Theory of Admissibility-Driven Inference Across Dynamical Systems, Measure Theory, and Language

[Moriba Kemessia Jah](#) *

Posted Date: 26 March 2026

doi: 10.20944/preprints202603.2010.v2

Keywords: epistemic abductive geometry; TEAG; possibility theory; admissible inference; non-Bayesian filtering; Possibilistic Cramér–Rao Bound; abductive inference; epistemic width; minimax commitment; language modeling; knowledge graphs



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

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

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

Article

Theory of Epistemic Abductive Geometry (TEAG): A Unified Theory of Admissibility-Driven Inference Across Dynamical Systems, Measure Theory, and Language

Moriba Kemessia Jah

- 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

We introduce the Theory of Epistemic Abductive Geometry (TEAG), a framework for non-Bayesian inference grounded in admissible-support contraction under possibility theory. The central object is the TEAG quintuple $\mathcal{E} = (H, \pi, \{H_\alpha\}_{\alpha \in (0,1]}, C, A)$, where evidence acts by contracting the geometry of admissible hypotheses rather than redistributing probabilistic belief mass. **Falsification has two-stage structure.** Under the log-admissibility transformation $\Phi(h) = -\log \pi(h)$, the canonical TEAG conjunctive update becomes tropical addition in the max-plus semiring: $\Phi^+(h) = \Phi^-(h) \oplus \psi(h) = \max(\Phi^-(h), \psi(h))$, where $\psi(h) = -\log \kappa(y | h)$ is the surprisal of hypothesis h under observation y . The tropical variety of this polynomial, $\mathcal{B}_{\text{active}} = \{h \in H : \Phi^-(h) = \psi(h)\}$, is the *active deformation front*: the exact locus where incoming evidence first matches prior impossibility and begins to deform the posterior field. This is a necessary condition for falsification but not sufficient. Sufficient falsification requires exit from the *PCRB admissible basin* $\mathcal{A}_k = \{h : \Phi_k^+(h) \leq c_k^*\}$, where c_k^* is the equipotential threshold determined by the PCRB at step k . Popper's criterion thus receives a two-stage algebraic formulation: the tropical variety marks where falsification becomes possible; the PCRB basin boundary marks where falsification is complete. Within the class of possibility-theoretic recursive inference systems, this is, to the best of our knowledge, the first exact formulation of this distinction. **Main results.** 1. **Epistemic Contraction Theorem.** Contraction is tropical addition: $\Phi^+ = \Phi^- \oplus \psi$. Posterior α -cuts satisfy $H_\alpha^+ = H_\alpha^- \cap E_\alpha(y)$: geometric intersection, not belief redistribution. The active deformation front is the tropical variety $\mathcal{B}_{\text{active}}$; the falsification boundary is the PCRB admissible basin boundary $\mathcal{B}_{\text{adm}}$. 2. **Possibilistic Cramér–Rao Bound (PCRB).** For any filter in the class $\mathcal{F}$ of epistemically admissible, contraction-based recursive estimators satisfying Axioms 2.1–2.5: $\mathcal{E}_{\pi,k|k} \geq \mathcal{E}_{\pi,k|k-1} + \frac{n}{2} \log(1 - I_k)$, where I_k is the Choquet integral of per-hypothesis surprisal against the prior possibility capacity. Within this class, the ESPF [28] is the unique filter achieving this bound with equality, and is therefore the unique minimax-entropy-optimal set-based recursive estimator under bounded epistemic uncertainty. 3. **Tropical Hamilton–Jacobi structure (summary).** The TEAG update is structurally consistent with a tropical Lagrangian $L = T - V$, Legendre transform to a tropical Hamiltonian equal to the surprisal field, and a Hamilton–Jacobi equation whose solution is the tropical addition rule. The Euler–Lagrange equations on the epistemic manifold yield geodesic motion with explicit Levi–Civita connection and Christoffel symbols. This structure is interpretive and consistent with the axioms; full derivations are in the companion paper [31]. Taken together, this structure admits a precise interpretation: the TEAG update rule is a max-plus dynamical system whose governing equations have the same algebraic form as the Hamilton–Jacobi equations of classical mechanics, instantiated on hypothesis space rather than physical space. 4. **Gaussian collapse.** Probability theory is the collapse limit of TEAG as epistemic width $W \rightarrow 0$: Choquet converges to Lebesgue, the ESPF recovers the Kalman filter, and $\mathcal{E}_\pi \rightarrow \frac{1}{2} \log \det \Sigma + \text{const}(n)$. Probability is earned by evidence, not assumed. **Epistemic neutrality and knowledge-system synthesis.** Because TEAG's axioms require only a hypothesis space, a possibility field, and a contraction operator — not a probability measure,

a likelihood function, or a frequentist grounding — heterogeneous knowledge systems can each instantiate the TEAG quintuple independently. Their joint admissible support intersection is the locus of coherence: the set of hypotheses neither system has falsified. No transformation of one system into the other's representational primitives is required. The composition theory (Section 6) formalizes the coupling architecture. **Four instantiations** provide the unifying structure: the ESPF [28] for recursive state estimation; the Geometry of Knowing [29] for measure-theoretic collapse; the minimax-entropy optimality proof [30]; and the Possibilistic Language Model (PLM, forthcoming [32]).

Keywords: epistemic abductive geometry; TEAG; possibility theory; admissible inference; non-Bayesian filtering; Possibilistic Cramér–Rao Bound; abductive inference; epistemic width; minimax commitment; language modeling; knowledge graphs

1. Introduction

1.1. The Problem with Probabilistic Closure

Most dominant inference frameworks in engineering, statistics, and machine learning share a structural commitment: uncertainty is represented by a normalized probability distribution and evidence acts by updating that distribution through Bayes' rule. This paradigm is extraordinarily powerful when uncertainty arises from repeatable statistical processes whose structure can be modeled probabilistically.

However, many real-world inference problems do not satisfy these assumptions. Observations may be sparse, bounded rather than stochastic, or only partially characterized by statistical models. In such settings the requirement that probability mass sum to one imposes an *epistemic closure condition*: belief must be distributed across the entire hypothesis space even when available evidence is insufficient to justify such commitment. Probability therefore forces a complete ordering of hypotheses even when the underlying evidence supports only partial comparisons of admissibility.

The resulting inference may appear mathematically precise while remaining epistemically fragile. As Kalman himself cautioned, a model may be “an artifact displaying the prejudices of its creator” rather than a faithful representation of the data [34]. Posterior distributions can become sharply concentrated despite weak evidence; likelihood models may assign nonzero mass to states that should be excluded by admissibility constraints; and commitment rules such as posterior means or maximum a posteriori estimates may select hypotheses that surviving evidence does not actually support.

Empirical evidence of this failure mode is concrete. In orbital tracking experiments with the Unscented Kalman Filter (UKF), covariance contraction can continue—communicating apparent confidence—even as the state estimate drifts from truth under unmodeled dynamics or sensor bias. The UKF “fails silently” by concentrating belief around an increasingly incorrect trajectory [28]. An epistemically honest filter should instead *fail informatively*: expose tension, preserve uncertainty, and signal when the model is no longer credible.

These observations motivate a different primitive for inference. Rather than beginning with a normalized probability density and redistributing belief mass under evidence, we begin with a nested admissible support geometry and allow evidence to contract that geometry by eliminating hypotheses incompatible with observation. If $E_\alpha(y)$ denotes the compatibility region induced by observation y , the admissible supports evolve according to

$$H_\alpha^+ = H_\alpha^- \cap E_\alpha(y),$$

so that inference proceeds through progressive intersection of nested admissible hypothesis sets.

1.2. The TEAG Object

Definition 1.1 (TEAG object). A TEAG object is a quintuple

$$\mathcal{E} = (H, \pi, \{H_\alpha\}_{\alpha \in [0,1]}, C, A),$$

where:

- (i) H is a hypothesis space, which may be finite, countable, or measurable.
- (ii) $\pi : H \rightarrow [0, 1]$ is a normalized possibility field satisfying $\sup_{h \in H} \pi(h) = 1$, encoding ordinal admissibility. Possibility values are ordinal: they express relative plausibility without probabilistic interpretation, do not represent frequencies or likelihoods, and are never summed or integrated in the manner of probability mass.
- (iii) For each $\alpha \in [0, 1]$, the α -cut

$$H_\alpha = \{h \in H : \pi(h) \geq \alpha\}$$

defines a nested admissible-support family $H_{\alpha'} \subseteq H_\alpha$ for $\alpha' > \alpha$, with $H_0 = H$ and $H_1 = \arg \max_h \pi(h)$.

- (iv) C is an evidence-driven epistemic contraction operator $C : [0, 1]^H \times \mathcal{Y} \rightarrow [0, 1]^H$ mapping prior possibility field π and observation $y \in \mathcal{Y}$ to posterior possibility field $\pi' = C(\pi, y)$. The operator satisfies the Popperian monotonicity condition

$$\pi'(h) \leq \pi(h) \quad \forall h \in H,$$

ensuring that evidence can only reduce admissibility, never amplify it. The canonical TEAG contraction takes the conjunctive form

$$C(\pi, y)(h) = \min(\pi(h), \kappa(y | h)),$$

where $\kappa(y | h) \in (0, 1]$ denotes the compatibility between hypothesis h and observation y .

- (v) $A : [0, 1]^H \rightarrow H$ is a minimax commitment rule selecting an operational hypothesis from the surviving admissible support. In TEAG instantiations considered here, A is a minimax operator acting on the geometry induced by the α -cut family:

$$A(\pi) = \arg \min_{h \in S_\pi} \max_{h' \in S_\pi} \rho_\pi(h, h'),$$

where ρ_π is a metric or pseudo-metric induced by the admissible-support geometry, and $S_\pi = \{h \in H : \pi(h) > 0\}$ is the surviving admissible support.

Inference in TEAG therefore operates on the geometry of admissible supports rather than on probability mass. The possibility field π serves to index the nested admissible sets $\{H_\alpha\}$ that define this geometry. Probability distributions, likelihood functions, and posterior densities arise only when additional structure—such as additivity, calibration assumptions, or parametric families—is imposed on π .

1.3. Four Instantiations

The TEAG object appears, under different mathematical realizations, in several previously developed frameworks. Table 1 summarizes four domains in which the quintuple $\mathcal{E} = (H, \pi, \{H_\alpha\}, C, A)$ arises naturally.

Table 1. Four realizations of the TEAG object across distinct inference domains.

Instantiation	H	π	C	A	Geometry
ESPF (state est.)	Finite support $\{\chi^{(i)}\}$	Ordinal state admissibility	Residual contraction	Whitened min-max medoid	MVEE ellipsoid in $\mathbb{R}^n$
GOK (measure theory)	Measurable space $\mathcal{X}$	Consonant poss. distribution	Credal contraction & collapse	Singleton prob. limit	Credal set shrinkage to point
ESPF Optimality	α -cut $\{V_\alpha\}$	α -indexed admissibility	Minimax-entropy C^*	PCRB-bounded commit.	α -cut volume manifold
PLM (language)	Vocabulary V	Token admissibility $\pi_k(v)$	Compat.-driven falsification	Minimax medoid token	Embedding-space ellipsoid

1.4. What Naming the Object Enables

Prior to this work, the four instantiations existed as formally separate contributions. Identifying the TEAG object as their shared primitive enables results inaccessible when the domains are treated independently.

- Transfer of results.** Theorems proven at the level of the TEAG object apply to every valid instantiation satisfying the axioms. In particular, the PCRB established operationally in [28] (via asymmetric rate limits $r^+ = 1.15, r^- = 0.97$) and proved formally in [30] for the Euclidean ESPF setting holds for any TEAG instantiation in which the isotropy condition and the epistemically admissible filter class conditions are satisfied. The conditions under which this transfer holds for each instantiation are identified in Section 5.
- Systematic construction of new domains.** New inference frameworks can be built by specifying (H, π, C, A) and verifying the TEAG axioms, rather than rederiving inference principles from scratch. Section 7 illustrates this for biological aging, quantum tomography, and multi-agent planning.
- Cross-domain composition.** Multi-domain reasoning systems such as GaiaVerse require simultaneous inference over dynamical state estimation, language reasoning, and knowledge graph structure. Because each domain shares the TEAG primitive, their contraction and commitment operations become compositionally compatible (Section 6).
- Clarification of intellectual lineage.** The underlying principles—nested admissible supports, contraction without probabilistic redistribution, minimax commitment—originate in the ESPF and GOK work and appear in subsequent developments as domain-specific instantiations of the same mathematical object.

1.5. TEAG as Epistemically Neutral Synthesis Architecture

A persistent challenge in knowledge systems design is the question of how to reason coherently across inference frameworks that operate under incommensurable epistemic commitments. Western scientific inference is grounded in probability, frequentist calibration, and reproducible measurement. Many Indigenous and localized knowledge systems operate under different primitives: relational observation, place-based evidence, ancestral testimony, and admissibility criteria that are not reducible to likelihood functions or statistical priors.

Standard integration attempts resolve this by transformation: Indigenous knowledge is cast as a “prior distribution,” or encoded as soft constraints in a Bayesian model, or treated as qualitative evidence to be assigned numerical weights. Each of these transformations is lossy. They require one framework to serve as the ground representation and the other to be reformatted as input. Epistemic content that is not expressible in the receiving framework’s primitives is discarded.

TEAG offers a different architecture. Its axioms require only:

- a hypothesis space H — the set of what a given knowledge system considers possible;
- a possibility field $\pi : H \rightarrow [0, 1]$ — ordinal admissibility under that system’s evidence criteria;
- a contraction operator C — the system’s rule for eliminating hypotheses when they are incompatible with observation.

No probability measure is required. No likelihood function is required. No shared numerical scale is required. Any knowledge system that can specify what it admits, what it rules out, and how it responds to evidence can be instantiated as a TEAG object.

Two such objects $\mathcal{E}_1 = (H_1, \pi_1, \{H_{\alpha,1}\}, C_1, A_1)$ and $\mathcal{E}_2 = (H_2, \pi_2, \{H_{\alpha,2}\}, C_2, A_2)$ — representing, for instance, an orbital mechanics inference system and a place-based Indigenous observation system over a shared physical domain — need not operate in the same hypothesis space. When their hypothesis spaces intersect, the compositional TEAG framework (Section 6) couples them through a shared admissible support: the joint surviving region is the set of hypotheses neither system has falsified. This is not a transformation of one system into the other. It is a joint falsification boundary — the locus $\mathcal{F}_{12}$ where both impossibility fields and both surprisal fields simultaneously balance.

The result is coherent multi-system inference without requiring either system to abandon its own epistemic primitives. Disagreement between systems is visible as a non-empty falsification gap: the set of hypotheses that one system admits but the other rules out. Agreement is the surviving intersection. The shared PCRB bounds how fast the joint admissible support can contract per observation, preventing either system from asserting certainty faster than the joint evidence warrants.

To the best of our knowledge, TEAG is the first inference framework whose axiomatic structure is epistemically neutral between knowledge systems in this sense: it provides a shared mathematical primitive that does not privilege any particular evidence type, calibration standard, or ontological commitment, while still admitting exact results — the tropical variety, the PCRB, the Gaussian collapse — when the structure of a particular instantiation warrants them.

1.6. Paper Organization

Section 4 establishes general results including the contraction theorem, the universal PCRB, and the Gaussian collapse theorem. Section 5 derives the four domain instantiations explicitly and maps each to material from the companion ESPF paper [28]. Section 6 develops the composition theory for multi-domain reasoning systems. Section 7 outlines several new instantiation directions. Section 8 situates TEAG relative to related work. Section 9 concludes.

1.7. Scope of Claims

To assist the reader in calibrating expectations, we distinguish three levels of claim in this paper.

What is proved within this paper. The TEAG axioms are stated and verified for each instantiation. The tropical contraction geometry (Proposition 3.12) and the active deformation front as tropical variety (Theorem 3.13) are proved here. The three-zone structure (Definition 3.14), the PCRB admissible basin and threshold c_k^* (Definition 3.15), the PCRB-based falsification criterion (Definition 3.17), and the minimax medoid as basin fixed point (Proposition 3.18) are proved here. The Choquet-to-Lebesgue convergence theorem (Theorem 4.7) is proved in [29] and cited. The PCRB (Theorem 4.5) is proved in [30] for the Euclidean ESPF setting and stated here as a universal TEAG result. The epistemic Lagrangian and tropical Hamiltonian (Theorem 3.30, items (i)–(ii)) are proved here as direct consequences of the tropical contraction structure. The compositional convergence theorem (Theorem 6.3) is proved from the TEAG axioms and the Banach fixed-point theorem.

What is established as structural summary. The Euler–Lagrange equations, Levi–Civita connection, and Christoffel symbols (Theorem 3.30, items (iii)–(iv)) are structural summaries consistent with the axioms and proved in full in [31]. The TEAG quintuple is defined, its axioms stated, and its instantiations derived. The impossibility field Φ , the epistemic geoid interpretation of the PCRB, and the Boltzmann justification for $\mathcal{E}_\pi$ are introduced here. These are mathematical contributions but depend on the proved theorems above for their force.

What is interpretive and forward-looking. The “physics of belief” framing — that the TEAG update rule is a max-plus dynamical system whose governing equations have the same algebraic form as the Hamilton–Jacobi equations of classical mechanics — is a mathematical observation about algebraic structure, not a claim that inference is a physical process. It is stated as an interpretation supported by Theorem 3.30, not as a theorem in its own right. The epistemic manifold $\mathcal{M}_{\text{ep}}$, the

constrained geometric flow interpretation, and the optimal transport analogy are geometric interpretations consistent with the established structure. A full differential-geometric theory with smooth global parameterization, a connection, and a Ricci tensor has not been developed; that is reserved for future work. The PLM instantiation is a principled structural mapping rather than a derivation from first principles. The new instantiation directions (Section 7) are research proposals, not completed results.

This paper's primary contribution is the identification and formalization of the TEAG object as the shared primitive underlying several previously distinct frameworks, and the transfer of results that this identification enables.

A consequence of this identification is that the ESPF [28] is not a parallel development to TEAG but an instantiation of it. When the TEAG axioms are applied to the problem of recursive state estimation over a finite support in $\mathbb{R}^n$, and the optimality criterion is possibilistic minimax entropy (minimization of worst-case integrated ignorance per measurement step), the ESPF update rule emerges as the unique solution within the class of epistemically admissible, contraction-based recursive estimators satisfying Axioms 2.1–2.5. Its survivor selection rule, possibility assignment, and commitment rule are not design choices — they are forced by the TEAG structure and the PCRB. This derivational relationship is established in Section 5.1 and 5.5.

2. Axiomatic Scope and Admissible Class

All uniqueness, optimality, and derivational claims in this paper are made with respect to a precisely defined class of inference systems. This section defines that class explicitly. Readers are encouraged to verify, when a claim of uniqueness or optimality appears later in the paper, that the result is scoped to this class.

2.1. The Epistemically Admissible Class $\mathcal{F}$

Definition 2.1 (Epistemically admissible inference system). *An inference system is epistemically admissible and belongs to the class $\mathcal{F}$ if it satisfies all of the following conditions simultaneously:*

- (A1) **Possibilistic representation.** *Uncertainty is represented as a normalized possibility distribution $\pi : H \rightarrow (0, 1]$ over a finite or measurable hypothesis space H , with $\sup_h \pi(h) = 1$. No probability measure, likelihood function, or additive normalization is assumed.*
- (A2) **Popperian contraction.** *Upon receiving evidence $y \in \mathcal{Y}$, the update operator C satisfies $C(\pi, y)(h) \leq \pi(h)$ for all $h \in H$. Evidence can only reduce admissibility, never amplify it.*
- (A3) **Non-resurrection.** *If $\pi(h) = 0$ then $C(\pi, y)(h) = 0$ for all y . A falsified hypothesis cannot be restored by subsequent evidence.*
- (A4) **Geometric non-degeneracy (VFI).** *In Euclidean instantiations, the MVEE shape matrix Π of the surviving support satisfies $\lambda_{\min}(\Pi) \geq \varepsilon > 0$ and $\lambda_{\max}(\Pi) \leq \Lambda < \infty$ at every step. This prevents support collapse to a lower-dimensional subspace.*
- (A5) **Evidence-referencing.** *The survivor selection ranking depends only on the whitened squared innovations $q_k^{(i)} = \|L_e^{-1}(y_k - h(\chi^{(i)}))\|^2$, not on the prior possibility $\pi_{k|k-1}^{(i)}$. The prior may enter the bounding of entropy reduction (via the Choquet information content I_k) but not the selection of survivors.*

2.2. Scope Statement

Remark 2.2 (Scope of uniqueness and optimality claims). *Every claim in this paper of the form “unique,” “optimal,” “only,” or “forced” refers to uniqueness or optimality within the class $\mathcal{F}$ of Definition 2.1, under the TEAG axioms of Section 3.*

These claims do not assert:

- uniqueness among all conceivable inference systems;
- uniqueness among non-possibilistic or probabilistic frameworks;
- uniqueness without the admissibility conditions above.

They do assert:

- that within $\mathcal{F}$, the ESPF is the unique filter achieving the PCRB with equality in the falsification regime;

- that within $\mathcal{F}$, minimum- q survivor selection is the unique rule minimizing possibilistic entropy $\mathcal{E}_\pi$ per measurement step;
- that within $\mathcal{F}$, the TEAG update rule is the unique solution to the tropical Hamilton–Jacobi equation with Hamiltonian equal to the surprisal field.

The class $\mathcal{F}$ is not narrow: it encompasses any inference system that represents uncertainty possibilistically, contracts under evidence without amplification, and selects survivors by evidence geometry rather than prior belief. Within this class, the results are tight.

2.3. Result Hierarchy

Table 2 summarizes the hierarchy of results in this paper, mapping each to its logical type. A reader encountering an unfamiliar claim can locate it here to determine what kind of support it carries.

Table 2. Result hierarchy: logical type of each contribution.

Result	Type	Depends on
TEAG axioms (A1–A5)	Definition	Primitive
Tropical contraction $\Phi^+ = \Phi^- \oplus \psi$	Proved here	Axioms A1–A3
Active deformation front as tropical variety	Proved here	Tropical contraction
Three-zone structure	Proved here	Compatibility assignment + VFI
PCRB admissible basin & threshold c_k^*	Proved here	PCRB + VFI + MVEE scaling
PCRB-based falsification	Proved here	PCRB basin
Minimax medoid as basin fixed point	Proved here	Three-zone + PCRB basin
PCRB	Proved in [30]	A1–A5 + VFI
ESPF uniqueness in $\mathcal{F}$	Proved in [30]	PCRB + A5
Epistemic Lagrangian & tropical HJ equation	Proved here	Tropical contraction
Euler–Lagrange, Christoffel symbols	Summary [31]	HJ structure
Gaussian collapse	Proved in [29]	A1–A3 + consonance
Compositional convergence	Proved here	A1–A5 + Banach
Physics-of-belief interpretation	Interpretation	HJ structure
Knowledge-system neutrality	Structural	A1–A3 only
PLM, aging, quantum instantiations	Research proposals	Axioms A1–A3

3. TEAG Axioms and Formal Structure

3.1. Primitive Definitions

Let H be a hypothesis space and let $\pi : H \rightarrow [0, 1]$ be a normalized possibility field satisfying $\sup_{h \in H} \pi(h) = 1$. For each $\alpha \in [0, 1]$, define the α -cut $H_\alpha = \{h \in H : \pi(h) \geq \alpha\}$. The family $\{H_\alpha\}_{\alpha \in [0, 1]}$ forms a nested admissible-support structure. The surviving admissible support is $S_\pi = \{h \in H : \pi(h) > 0\}$.

The TEAG axioms govern evidence-driven *contraction* of the admissible support. In dynamical instantiations, process uncertainty acts as a complementary *expansion* operator on H or its support, typically via set-valued propagation (e.g., Minkowski sums in the ESPF [28]). Contraction and expansion operate at different phases of the recursion: expansion between observations (Jaynesian maximum entropy), contraction upon observation (Popperian falsification). This reflects a fundamental separation: *state evolution is kinematic* — governed by physical dynamics and process uncertainty acting on H —

while *belief evolution is epistemic* — governed by the impossibility field Φ and its deformation under evidence.

Let $\mathcal{Y}$ denote the observation space. A TEAG contraction operator is a map $C : [0, 1]^H \times \mathcal{Y} \rightarrow [0, 1]^H$, which sends prior possibility field π and observation $y \in \mathcal{Y}$ to posterior possibility field $\pi' = C(\pi, y)$.

Let ρ_π denote a metric or pseudo-metric induced by the geometry of the admissible support. In Euclidean instantiations this may be the Mahalanobis-type metric induced by the minimum-volume enclosing ellipsoid (MVEE) [4,9] of the surviving support.

3.2. Core TEAG Axioms

Axiom 3.1 (Regular admissibility). *The possibility field π induces a nested α -cut family $\{H_\alpha\}_{\alpha \in [0,1]}$ satisfying $H_{\alpha'} \subseteq H_\alpha$ for $\alpha' > \alpha$, and right-continuity $\bigcap_{\alpha < \beta} H_\alpha = H_\beta$ for all $\beta \in (0, 1]$.*

Axiom 3.2 (Popperian contraction). *For every observation $y \in \mathcal{Y}$ and hypothesis $h \in H$: $C(\pi, y)(h) \leq \pi(h)$. Evidence can only reduce admissibility.*

Axiom 3.3 (Non-resurrection). *If $\pi(h) = 0$, then $C(\pi, y)(h) = 0$ for all $y \in \mathcal{Y}$. A falsified hypothesis cannot be restored by subsequent contraction.*

Axiom 3.4 (Commitment admissibility). *The commitment rule selects a hypothesis from the maximal α -cut: $A(\pi) \in H_1$.*

Axiom 3.5 (Minimax commitment). *The commitment rule is minimax with respect to the induced admissible-support geometry: $A(\pi) = \arg \min_{h \in S_\pi} \max_{h' \in S_\pi} \rho_\pi(h, h')$. Thus $A(\pi)$ is a minimax medoid of the surviving support.*

Remark 3.6 (Why minimax rather than expectation). *Expectation is not a primitive commitment rule in TEAG for two reasons. First, possibility values are ordinal and therefore do not support additive expectation operators. Second, an expectation operator may produce a point outside the surviving admissible support. The minimax medoid avoids both failures: it is defined intrinsically on the admissible-support geometry and returns a hypothesis that is admissible by construction. In ESPF implementations this is the whitened minimax medoid—the surviving support point minimizing maximum whitened distance to all other candidates—ensuring the anchor corresponds to an actual surviving hypothesis rather than an interpolated or averaged state [28].*

3.3. Canonical Contraction Rule

The canonical TEAG update takes the conjunctive form

$$\hat{\pi}(h) = \min(\pi(h), \kappa(y | h)), \quad \pi'(h) = \frac{\hat{\pi}(h)}{\sup_{h' \in H} \hat{\pi}(h')}, \quad (1)$$

where $\kappa(y | h) \in (0, 1]$ denotes the compatibility between hypothesis h and observation y . Max-rescaling in (1) preserves ordinal structure and is numerically essential for long-horizon stability: without it, repeated conjunctive updates cause progressive collapse of numerical dynamic range without any epistemic justification [28].

The update in (1) satisfies Axioms 2–3 by construction.

Remark 3.7 (Contraction, posterior possibility, and necessity). *The canonical update (1) establishes a precise logical chain that is worth stating explicitly. Compatibility $\kappa(y | h) \in (0, 1]$ is the falsification mechanism: it measures the degree to which hypothesis h is geometrically consistent with observation y . Posterior possibility $\pi'(h) = \hat{\pi}(h) / \sup_{h'} \hat{\pi}(h')$ is the credibility of hypothesis h after the update: the intersection of prior admissibility and evidentiary compatibility, max-rescaled to preserve ordinal structure.*

Necessity is the dual: $N(A) = 1 - \Pi(A^c) = 1 - \sup_{h \notin A} \pi'(h)$, the degree to which event A is epistemically guaranteed rather than merely plausible. Explicitly:

$$\begin{aligned} \kappa(y | h) &\longrightarrow \text{falsification mechanism,} \\ \pi'(h) = \min(\pi(h), \kappa(y | h)) &(\text{normalized}) \longrightarrow \text{posterior possibility (credibility),} \\ N(A) = 1 - \sup_{h \notin A} \pi'(h) &\longrightarrow \text{necessity (confidence lower bound).} \end{aligned}$$

Necessity is not separately defined — it is derived from posterior possibility via the standard duality of possibility theory. High necessity means few hypotheses outside A remain admissible; it is the TEAG analog of high posterior probability concentrated within A , but without requiring probabilistic additivity.

Proposition 3.8 (α -cut contraction). Let H_α^- denote the prior α -cut and define the compatibility region $E_\alpha(y) = \{h \in H : \kappa(y | h) \geq \alpha\}$. Then the posterior α -cuts satisfy

$$H_\alpha^+ = H_\alpha^- \cap E_\alpha(y).$$

TEAG updating therefore corresponds to geometric intersection of admissible supports under evidence.

Remark 3.9 (Comparison with Bayes' rule). Bayesian updating takes the form $P(h | y) \propto P(y | h)P(h)$, and may increase posterior mass on hypotheses whose likelihood is large relative to the normalization constant. In TEAG, admissibility is never amplified by evidence. Evidence acts asymmetrically: it contracts admissible support by eliminating hypotheses incompatible with observation rather than redistributing belief mass across the hypothesis space.

3.4. Log-Admissibility Coordinates and Tropical Geometry

The TEAG contraction rule admits a natural geometric interpretation in log-admissibility coordinates that connects abductive inference with tropical geometry.

3.4.1. Log-Admissibility Coordinates and the Impossibility Field

Definition 3.10 (Impossibility field). Let $\pi : H \rightarrow [0, 1]$ be a normalized possibility field. The impossibility field induced by π is the map $\Phi : H \rightarrow [0, \infty]$ defined by

$$\Phi(h) = -\log \pi(h).$$

$\Phi(h)$ assigns to each hypothesis a continuous, non-negative degree of impossibility. $\Phi(h) = 0$ when $\pi(h) = 1$ —full admissibility, zero impossibility. $\Phi(h) \rightarrow \infty$ as $\pi(h) \rightarrow 0$ —complete falsification. The α -cut family is exactly the sublevel-set family of Φ :

$$H_\alpha = \{h \in H : \Phi(h) \leq -\log \alpha\},$$

so that the nested admissible-support structure of TEAG is the nested sublevel-set structure of the impossibility field.

Under this transformation, admissibility ordering becomes additive:

$$\pi(h_1) \geq \pi(h_2) \iff \Phi(h_1) \leq \Phi(h_2).$$

Thus higher admissibility corresponds to lower impossibility, and the geometry of admissible supports is equivalently the geometry of sublevel sets of Φ .

3.4.2. Tropicalization of the Contraction Operator

The canonical TEAG update is

$$\pi^+(h) = \min(\pi^-(h), \kappa(y | h)).$$

Define the *surprisal* $\psi(h) = -\log \kappa(y | h)$ of hypothesis h under observation y . Taking negative logarithms of the canonical update yields

$$\Phi^+(h) = \max(\Phi^-(h), \psi(h)).$$

The impossibility of a hypothesis after an observation is the greater of its prior impossibility and its surprisal. Evidence can only increase impossibility, never reduce it—Popperian monotonicity (Axiom 2.2) expressed as the natural dynamics of Φ .

Remark 3.11 (Connection to tropical algebra). *Under the transformation $\Phi = -\log \pi$, the conjunctive min update in possibility space becomes a max update in impossibility space. The resulting algebra $(\oplus, \otimes) = (\max, +)$ corresponds to the max-plus tropical semiring [2,22] widely studied in tropical geometry and optimal control. TEAG contraction is therefore tropical addition of impossibility fields: $\Phi^+ = \Phi^- \oplus \psi$, where $\oplus = \max$. The dual convention $(\min, +)$ yields the same structure under negation; we adopt $(\max, +)$ because it aligns with the direction of impossibility elevation induced by evidence.*

3.4.3. Abductive Inference as Tropical Potential Flow

In impossibility-field coordinates, evidence introduces a potential barrier $\psi(h)$ over the hypothesis space. The update $\Phi^+(h) = \max(\Phi^-(h), \psi(h))$ raises the impossibility of hypotheses that are surprised by the observation while leaving unsurprised hypotheses unchanged. Abductive inference proceeds by progressively elevating the impossibility of incompatible hypotheses until the admissible basin collapses around a surviving explanatory region. The rate at which the basin volume may contract per observation is bounded below by the PCRB (Section 4.3), whose information content is the Choquet integral of per-hypothesis surprisal with respect to the prior possibility capacity.

Proposition 3.12 (Tropical contraction geometry). *Let $\Phi(h) = -\log \pi(h)$ be the impossibility field and $\psi(h) = -\log \kappa(y | h)$ the surprisal of hypothesis h under observation y . Then TEAG contraction corresponds to tropical addition*

$$\Phi^+ = \Phi^- \oplus \psi$$

under the tropical algebra $(\oplus, \otimes) = (\max, +)$.

Theorem 3.13 (Active deformation front as tropical variety). *Let $\Phi^-(h) = -\log \pi^-(h)$ be the prior impossibility field and $\psi(h) = -\log \kappa(y | h)$ the surprisal of hypothesis h under observation y . The posterior impossibility field under the canonical TEAG contraction is the tropical polynomial*

$$\Phi^+(h) = \Phi^-(h) \oplus \psi(h) = \max(\Phi^-(h), \psi(h)).$$

The active deformation front—the locus where incoming evidence first matches prior impossibility and begins to deform the posterior field—is the tropical variety of this polynomial:

$$\mathcal{B}_{\text{active}} = \{h \in H : \Phi^-(h) = \psi(h)\}.$$

This is the zero set of the tropical polynomial $\Phi^- \ominus \psi = \Phi^- - \psi$ in the max-plus semiring, determined entirely by the geometry of the prior impossibility and current surprisal fields. In Euclidean TEAG instantiations, where $\psi^{(i)} = \frac{1}{2}q^{(i)} = \frac{1}{2}\|L_e^{-1}e^{(i)}\|^2$, the tropical variety is the locus of support points $\chi^{(i)}$ satisfying

$$-\log \pi^{(i)} = \frac{1}{2}q^{(i)},$$

the level set on which prior log-admissibility exactly equals whitened squared innovation. $\mathcal{B}_{\text{active}}$ is a necessary condition for falsification: no hypothesis can be falsified without first being reached by the active deformation front. It is not sufficient. Sufficient falsification is defined by basin exit (Definition 3.17).

Definition 3.14 (Three-zone structure). Under the compatibility-based possibility assignment $\pi^{(i)} = e^{-q^{(i)}/2}$ and the VFI condition (Axiom A4), the surviving support S_π admits a canonical three-zone decomposition in whitened coordinates $z^{(i)} = L^{-1}(\chi^{(i)} - \mu)$:

$$\begin{aligned} \text{Zone I (interior): } & \pi^{(i)} > e^{-1/2}, \quad \|z^{(i)}\| < 1, \\ \text{Zone II (contact layer): } & \pi^{(i)} = e^{-1/2}, \quad \|z^{(i)}\| = 1, \\ \text{Zone III (boundary exterior): } & \pi^{(i)} < e^{-1/2}, \quad \|z^{(i)}\| > 1. \end{aligned}$$

Zone II consists of the contact points of the prior MVEE. Under the compatibility assignment, every contact point carries possibility value $e^{-1/2}$ —a structural constant independent of the observation, dimension, and support configuration. The VFI condition guarantees the three-zone structure is non-degenerate at every step: $\lambda_{\min}(\Pi) \geq \varepsilon > 0$ ensures Zone II spans $\mathbb{R}^n$; $\lambda_{\max}(\Pi) \leq \Lambda < \infty$ ensures Zone III is bounded.

Definition 3.15 (PCRB admissible basin). Let I_k be the possibilistic information content of observation y_k (Definition 5) and let r_k^- be the whitened MVEE radius of the prior support. The PCRB-admissible basin radius is

$$r_k^* = r_k^- \cdot \left((1 - I_k) \cdot \frac{M}{M_{\text{surv}}} \right)^{1/n},$$

and the PCRB-admissible threshold is $c_k^* = \frac{1}{2}(r_k^*)^2$. The admissible basin is the sublevel set of the posterior impossibility field at this threshold:

$$\mathcal{A}_k = \{h \in H : \Phi_k^+(h) \leq c_k^*\},$$

with admissibility boundary

$$\mathcal{B}_{\text{adm},k} = \{h \in H : \Phi_k^+(h) = c_k^*\}.$$

In original coordinates $\mathcal{B}_{\text{adm},k}$ is the MVEE ellipsoid of the posterior support at the PCRB-admissible α -cut level $\alpha_k^* = e^{-c_k^*}$. In whitened coordinates it is the ball of radius r_k^* .

Remark 3.16 (Geometric error bound). The closed form for r_k^* is exact in whitened isotropic coordinates. In the anisotropic case the MVEE scaling law $\det P_\alpha \approx f(\alpha)^{2/n} \det P$ introduces a geometric error bounded by

$$\Delta \leq \frac{e^{-1/2}}{n} \cdot \log \frac{M}{n+1},$$

which decays as $\log(n)/n$ and is controlled uniformly by the VFI condition and the Smolyak grid level ℓ .

Definition 3.17 (PCRB-based falsification). A hypothesis $h \in H$ is falsified at step k if and only if

$$\Phi_k^+(h) > c_k^*.$$

It survives at step k if and only if $\Phi_k^+(h) \leq c_k^*$, i.e., it lies inside or on the admissibility boundary $\mathcal{B}_{\text{adm},k}$.

The active deformation front $\mathcal{B}_{\text{active},k}$ and the admissibility boundary $\mathcal{B}_{\text{adm},k}$ are distinct objects operating at different stages. A hypothesis may be touched by $\mathcal{B}_{\text{active},k}$ —its posterior impossibility begins to rise—without being falsified, provided it remains inside $\mathcal{A}_k$. Falsification occurs only upon exit from the PCRB-admissible basin, not upon first contact with the deformation front.

Proposition 3.18 (Minimax medoid as basin fixed point). Under the compatibility-based possibility assignment and the VFI condition, the minimax medoid $A(\pi)$ lies in Zone I at every step and is the unique

support point maximally insulated from the admissibility boundary $\mathcal{B}_{\text{adm},k}$ under the MVEE metric. In whitened coordinates $A(\pi)$ is the support point minimizing $\|z\|$, with geometric insulation

$$d(A(\pi), \mathcal{B}_{\text{adm},k}) = r_k^* - \|z^{(\text{medoid})}\| > 0,$$

bounded below by the VFI eigenvalue condition $\lambda_{\min}(\Pi) \geq \varepsilon > 0$. The minimax medoid is therefore the last support point to be falsified under any sequence of PCRb-admissible contractions. This is not a design property of the commitment rule—it is a theorem about the geometry of the admissible basin.

Remark 3.19 (Correction to Theorem 3.13 and prior statement). Earlier versions of this paper and the companion papers identified the tropical variety $\mathcal{F} = \{h : \Phi^-(h) = \psi(h)\}$ as the falsification boundary—the locus dividing surviving from falsified hypotheses. That identification is here refined. $\mathcal{F}$ is retained and renamed the active deformation front $\mathcal{B}_{\text{active}}$: it correctly identifies where evidence begins to deform the posterior impossibility field, and its tropical variety structure is exact and unchanged. What is corrected is its epistemic role.

The pathology motivating the correction: a hypothesis with very low prior impossibility $\Phi^-(h) \approx 0$ —earned through many observations—is vulnerable to falsification by arbitrarily mild surprisal $\psi(h) > \Phi^-(h)$ under the original definition, even when $\psi(h)$ is small in absolute terms. This violates the principle that epistemic inertia should protect well-established hypotheses. The PCRb-based falsification definition resolves this by requiring basin exit rather than pointwise crossing: a hypothesis survives as long as it lies within the PCRb-admissible equipotential basin, regardless of whether the active front has touched it.

Popper's criterion receives a two-stage formalization: $\mathcal{B}_{\text{active}}$ marks where falsification becomes possible; $\mathcal{B}_{\text{adm}}$ marks where falsification is complete.

Remark 3.20 (Epistemic energy landscape). The impossibility field $\Phi(h)$ plays the role of an epistemic energy landscape. Highly admissible hypotheses occupy low-energy basins; incompatible hypotheses are lifted by evidence-induced surprisal. Inference evolves through progressive deformation of this landscape rather than redistribution of probability mass.

The minimax medoid commitment rule $A(\pi)$ does not select the hypothesis of lowest impossibility—that would be $\arg \min_{h \in S_\pi} \Phi(h)$, the mode of the possibility field, which may lie at the geometric periphery of the surviving support. Instead, $A(\pi)$ selects the hypothesis that is most insulated from the walls of the surviving basin: the support point minimizing worst-case whitened distance to all other surviving hypotheses under the MVEE-induced metric. This is a property of the basin geometry as a whole, not of any individual hypothesis's admissibility.

The distinction matters. A hypothesis at the edge of the surviving support may carry high admissibility yet be maximally exposed to the spread of remaining epistemic uncertainty. The minimax medoid avoids this exposure by selecting the point of greatest geometric centrality among surviving hypotheses. Empirically, this commitment rule yields lower estimation error than both the posterior mean and the Chebyshev center in ESPF implementations [28], confirming that geometric insulation within the admissible support—not individual admissibility extremity—is the correct criterion for epistemic commitment under bounded uncertainty.

Proposition 3.21 (Minimax commitment as basin selection). Let $\Phi(h) = -\log \pi(h)$ be the impossibility field and let ρ_π be the metric induced by the MVEE of the surviving support S_π . Then the minimax medoid commitment rule

$$A(\pi) = \arg \min_{h \in S_\pi} \max_{h' \in S_\pi} \rho_\pi(h, h')$$

selects the hypothesis minimizing the maximum distance within the surviving support under the MVEE-induced metric. In Euclidean instantiations with $H \subset \mathbb{R}^n$, this corresponds to selecting the point of greatest geometric centrality within the admissible basin: the hypothesis that is, in the worst case, closest to all other surviving hypotheses. This is geometric insulation, not expectation, and not impossibility minimization.

Remark 3.22. Proposition 3.21 closes the loop between the three geometric primitives of TEAG: nested admissible supports $\{H_\alpha\}$ define the basin topology as sublevel sets of Φ ; tropical updates $\Phi^+ = \Phi^- \oplus \psi$ deform the basin geometry under evidence; and minimax medoid commitment $A(\pi)$ selects the most geometrically insulated point of the surviving basin. The PCRB (Section 4.3) bounds the admissible rate of contraction of both the basin volume (the MVEE term) and the impossibility gradient (the texture of Φ within the basin), with information content measured by the Choquet integral of surprisal against the prior possibility capacity.

Remark 3.23 (The PCRB as epistemic geoid and $A(\pi)$ as its center). The relationship between the impossibility field Φ , the PCRB, and the commitment rule $A(\pi)$ admits a precise geometric interpretation analogous to the geoid in gravitational theory.

In physical geodesy, the geoid is an equipotential surface of Earth's gravitational field—not a perfect sphere, but the surface of constant potential that reflects the true asymmetric mass distribution of the Earth. Elevation is measured relative to it, and its center serves as the geometric reference for all positional commitment.

The PCRB plays the same role in TEAG. It is an equipotential surface of the impossibility field Φ —a level set $\{\Phi(h) = c^*\}$ at the threshold of admissible contraction, tracing the boundary of what the current observation can epistemically resolve. Just as the geoid is asymmetric when the underlying mass distribution is asymmetric, the PCRB surface is asymmetric when the surviving admissible support has been shaped by asymmetric evidence.

The whitened minimax medoid $A(\pi)$ is the center of this epistemic geoid: the point that is equidistant in the worst case from the PCRB surface in all directions under the MVEE-induced metric ρ_π . It is not the most admissible point $\arg \min_{h \in S_\pi} \Phi(h)$, which may sit close to one wall of the basin. It is the point of maximum geometric symmetry with respect to the equipotential boundary—the commitment that is most sheltered from wherever the next contraction may come.

This picture unifies the three central objects of Section 3.4: Φ defines the impossibility landscape, the PCRB traces its critical equipotential surface, and $A(\pi)$ commits to the center of that surface. Inference in TEAG is the progressive deformation of this epistemic geoid under evidence, with each observation reshaping the equipotential surface and its center. The gravitational analogy reflects a structural equivalence of governing equations — equipotential surfaces, center-of-mass commitment, field deformation under external input — not a physical force model. No claim is made that epistemic inference involves gravitational fields or spacetime curvature.

3.4.4. Connection to Support Geometry

In impossibility-field coordinates, the α -cut family corresponds to sublevel sets of Φ :

$$H_\alpha = \{h : \Phi(h) \leq -\log \alpha\}.$$

Evidence deforms the admissible-support geometry through tropical addition of impossibility fields. The resulting contraction of the α -cut family produces the entropy change measured by the TEAG possibilistic entropy functional introduced in Section 4.

Remark 3.24 (Surprisal as the impossibility increment). The surprisal $\psi(h) = -\log \kappa(y | h)$ is not merely analogous to a potential increment — it is the impossibility increment induced by observation y . In the ESPF, ψ evaluated at support point $\chi^{(i)}$ is exactly the geometric surprisal:

$$\psi(\chi^{(i)}) = \mathcal{S}^{(i)} = \frac{1}{2} q^{(i)} = \frac{1}{2} \|L_e^{-1} e^{(i)}\|^2.$$

This is an identity. The whitened squared innovation norm is the negative log compatibility, which is the surprisal. The tropical contraction update

$$\Phi^+(h) = \max(\Phi^-(h), \psi(h))$$

is therefore surprisal-driven impossibility deformation: each support point's impossibility is elevated by its surprisal under the current observation, and the larger of the prior impossibility and the surprisal is retained. A

hypothesis that is strongly surprised by the observation is pushed toward impossibility; one that is unsurprised remains in the basin.

This identity grounds the tropical geometry of TEAG in the operational mechanics of the ESPF. The epistemic energy landscape is not a metaphor imposed on the filter—it is the structure the filter is already computing. Surprisal is the currency of impossibility deformation, and the α -cut family contracts precisely at the rate permitted by the PCRB (Section 4.3).

3.5. Metric Geometry of Epistemic Contraction

In Euclidean TEAG instantiations, the admissible-support geometry induces a natural local metric structure. Let $S_\pi \subset \mathbb{R}^n$ denote the surviving admissible support and let $P_\pi \succ 0$ be the shape matrix of its minimum-volume enclosing ellipsoid (MVEE). Define the *local metric tensor*

$$g_\pi(u, v) = u^\top P_\pi^{-1} v, \quad (2)$$

with corresponding induced distance

$$\rho_\pi(x, x') = \sqrt{(x - x')^\top P_\pi^{-1} (x - x')}. \quad (3)$$

This is precisely the metric implicit in ESPF whitening: the Cholesky factor L_e with $\Pi_e = L_e L_e^\top$ defines (3) on the predicted measurement support, and the whitened innovation norm $\|L_e^{-1} e^{(i)}\|^2$ is distance-squared under this metric [28].

More generally, each nondegenerate α -cut H_α with MVEE shape matrix $P_\alpha \succ 0$ induces a metric

$$g_\alpha(u, v) = u^\top P_\alpha^{-1} v. \quad (4)$$

A TEAG object therefore carries not only a nested family of admissible supports $\{H_\alpha\}$ but also a corresponding nested family of local metrics $\{g_\alpha\}$. The TEAG possibilistic entropy $\mathcal{E}_\pi = \int_0^1 \log V_\alpha d\alpha$ is the integrated log-volume of this metric family.

Metric deformation under evidence.

Under a single observation, contraction of admissible support deforms the metric family. Let P_α^- and P_α^+ denote the pre- and post-update shape matrices for α -cut H_α . The volume contraction ratio is

$$\frac{V_\alpha^+}{V_\alpha^-} = \left(\frac{\det P_\alpha^+}{\det P_\alpha^-} \right)^{1/2}, \quad (5)$$

so the TEAG entropy change under evidence takes the form

$$\mathcal{E}_\pi^+ - \mathcal{E}_\pi^- = \frac{1}{2} \int_0^1 \log \det(P_\alpha^+ (P_\alpha^-)^{-1}) d\alpha. \quad (6)$$

Evidence does not merely alter a scalar uncertainty value; it deforms the local metric and volume form of the entire admissible-support geometry. In this representation, epistemic contraction is metric deformation: each observation reshapes the inner-product structure that governs distances, basin widths, and minimax commitment within the surviving support.

Remark 3.25 (Geometric interpretation of the PCRB). *In Euclidean TEAG instantiations, the PCRB is a curvature-like bound on admissible epistemic contraction. Equation (6) makes this precise: the PCRB (Section 4.3) bounds $\mathcal{E}_\pi^+ - \mathcal{E}_\pi^-$ from below, constraining both the basin volume term (how rapidly the MVEE boundary contracts) and the impossibility gradient term (how rapidly the texture of Φ steepens within the basin). Just as curvature bounds in differential geometry control how volumes evolve under geodesic flow, the PCRB controls how the full impossibility field geometry evolves under evidence. The asymmetric rate limits $r^+ = 1.15$, $r^- = 0.97$ enforced in the ESPF [28] operationalize this bound: fast metric expansion (embrace of ignorance) is permitted,*

while fast metric contraction (assertion of certainty) is restricted in proportion to the Choquet information content of the observation (Section 4.2).

3.6. Epistemic Manifold Interpretation

The ingredients developed in Sections 3.5 and 3.4—nested α -cut families, MVEE-induced metrics, tropical potential updates, and entropy as integrated log-volume—are collectively consistent with a manifold-like interpretation of TEAG inference. This section makes that interpretive structure explicit. We do *not* claim a full differential-geometric theory with smooth global parameterization, a connection, or a Ricci tensor — those structures require a smooth coordinate chart on the configuration space and a formal proof that $\{P_\alpha\}$ varies smoothly under the TEAG recursion, which remains future work. What is established is the metric-volume foundation that such a theory would require, and the geometric interpretation it supports.

Epistemic states as configuration-space points.

A TEAG state is not a single hypothesis $h \in H$ but an entire admissible-support configuration

$$\Pi = (\pi, \{H_\alpha\}_{\alpha \in [0,1]}, \{P_\alpha\}_{\alpha \in [0,1]}),$$

where π is the possibility field, $H_\alpha = \{h : \pi(h) \geq \alpha\}$ is the α -cut family, and $P_\alpha \succ 0$ is the MVEE shape matrix of H_α . Points in the *epistemic configuration space* $\mathcal{M}_{\text{ep}}$ are states of admissibility: configurations of what has not yet been ruled out, together with the local metric geometry those configurations carry. We use manifold language as geometric shorthand for this configuration space, with the understanding that the full differential-geometric structure has not yet been established.

The triple $(\mathcal{M}_{\text{ep}}, g, \Phi)$.

A Euclidean TEAG object determines a triple

$$(\mathcal{M}_{\text{ep}}, g, \Phi),$$

where $\mathcal{M}_{\text{ep}}$ is the space of admissible-support configurations, g is the MVEE-induced metric family $\{g_\alpha\}$ defined in (4), and Φ is the impossibility field of Definition 3.10. These three objects are not independent: the possibility field π determines $\{H_\alpha\}$, which determines $\{P_\alpha\}$, which determines g ; and $\Phi = -\log \pi$ is the impossibility field whose sublevel sets are exactly the α -cuts. The entire geometric structure is generated by π alone.

Inference as geometric flow.

In this language, a single TEAG update is a map on $\mathcal{M}_{\text{ep}}$:

$$\text{Impossibility deformation: } \Phi^+ = \Phi^- \oplus \psi,$$

$$\text{Metric deformation: } P_\alpha^- \mapsto P_\alpha^+,$$

$$\text{Entropy bound (Section 4.3): } \mathcal{E}_\pi^+ - \mathcal{E}_\pi^- \geq \frac{n}{2} \log(1 - I_k),$$

$$\text{Commitment (basin-center selection): } A(\pi) = \text{minimax medoid under } g_\pi.$$

Recursive TEAG contraction therefore traces a path through $\mathcal{M}_{\text{ep}}$: at each step, evidence deforms the impossibility landscape and the local metric; the PCRB constrains both the rate of basin-volume contraction and the rate of impossibility-gradient steepening; and commitment selects the center of the epistemic geoid.

Proposition 3.26 (Constrained geometric flow interpretation). *Let $\{\Pi_k\}_{k \geq 0}$ be a sequence of Euclidean TEAG states with nondegenerate MVEE family $\{P_{\alpha,k}\}$. Each update induces a deformation of the local metric family $\{g_{\alpha,k}\}$, and the TEAG entropy change satisfies*

$$\mathcal{E}_{\pi_k}^+ - \mathcal{E}_{\pi_k}^- = \frac{1}{2} \int_0^1 \log \det(P_{\alpha,k}^+ (P_{\alpha,k}^-)^{-1}) d\alpha.$$

The PCRB (Section 4.3) bounds this quantity from below via $\mathcal{E}_{\pi_k}^+ - \mathcal{E}_{\pi_k}^- \geq \frac{n}{2} \log(1 - I_k)$, where I_k is the Choquet information content of the observation. The recursion is consistent with a constrained geometric flow selecting the path of least unjustified deformation of the impossibility field geometry — least contraction of the basin volume and least steepening of the impossibility gradient beyond what the evidence warrants. Whether this flow is a geodesic in a fully formalized Riemannian sense remains an open question pending the smooth-parameterization results noted in Section 1.7.

Relation to information geometry.

Classical information geometry [1] assigns to each probability distribution a point on a statistical manifold with Fisher metric. TEAG defines an analogous structure, but with a fundamentally different primitive: points are admissibility configurations rather than distributions, and the metric is induced by MVEE geometry rather than Fisher information. In the Gaussian collapse limit (Section 4), where the possibility field contracts to a Gaussian density, the MVEE shape matrix P_π converges to the covariance Σ , and the epistemic metric (2) converges to the Fisher metric $g_F(u, v) = u^\top \Sigma^{-1} v$. The epistemic manifold therefore contains classical information geometry as a limiting submanifold.

Remark 3.27 (What is and is not claimed). *The present formulation establishes: (i) a local metric structure on admissible-support configurations; (ii) a potential-field dynamics given by tropical addition; (iii) a volume-form contraction law bounded by the PCRB; and (iv) a geodesic-flow interpretation of recursive TEAG contraction. What is not yet claimed is a full differential-geometric theory with smooth global parameterization, a connection, or a Ricci tensor. Those structures require a smooth coordinate chart on $\mathcal{M}_{\text{ep}}$ and a formal proof that $\{P_\alpha\}$ varies smoothly under the TEAG recursion. The present formulation provides the metric-volume foundation that such a theory would require; its development is reserved for future work.*

Remark 3.28 (Connection to optimal transport). *Classical optimal transport equips probability measures with a geometry induced by least-cost mass displacement: given two measures μ_0, μ_1 and a cost $c(x, y)$, it finds the coupling that moves μ_0 to μ_1 with minimum total cost, defining a Wasserstein distance on the space of measures. TEAG admits an analogous but structurally distinct geometry. Rather than transporting probability mass, each observation induces a one-sided deformation of the admissible-support configuration*

$$\Pi^- = (\pi^-, \{H_\alpha^-\}, \{P_\alpha^-\}) \mapsto \Pi^+ = (\pi^+, \{H_\alpha^+\}, \{P_\alpha^+\}),$$

contracting rather than redistributing. The local quadratic cost at level α is $c_\alpha(x, x') = (x - x')^\top P_\alpha^{-1} (x - x')$ —the MVEE-induced cost already implicit in (3)—and the entropy change (6) is the integrated log-volume distortion of this deformation, playing the role that Jacobian distortion plays in Wasserstein geometry. The PCRB (Section 4.3) bounds the admissible rate of this deformation: it constrains both the basin-volume contraction (the MVEE boundary) and the impossibility-gradient steepening (the texture of Φ within the basin), with the bound governed by the Choquet information content I_k of the observation. Three structural distinctions separate TEAG from classical optimal transport: motion is one-sided (evidence only contracts); admissibility rather than mass is conserved until falsified; and the geometry lives on support configurations, not probability laws. For these reasons we call the resulting structure one-sided support-transport geometry: recursive TEAG inference defines a support-transport flow on epistemic state space, dual in spirit to Wasserstein flow on probability space.

Remark 3.29 (Connection to Hamilton–Jacobi dynamics). *The impossibility field update*

$$\Phi^+(h) = \max(\Phi^-(h), \psi(h))$$

has the structure of a discrete max-plus Hamilton–Jacobi update: the impossibility field Φ evolves under an observation-induced barrier field ψ (the surprisal, as established in Remark 3.24), and the α -cuts

$$H_\alpha = \{h : \Phi(h) \leq -\log \alpha\}$$

evolve as sublevel-set fronts under this dynamics. This is a degenerate Hamilton–Jacobi system: the Hamiltonian $H(\Phi, \nabla\Phi) = \psi(h)$ is independent of the gradient of Φ (there is no classical kinetic term), reflecting the fact that TEAG contraction is purely evidence-driven — there is no momentum, no inertia, and no state-space propagation in the epistemic update. The dynamics are entirely determined by the barrier field ψ , not by the geometry of the impossibility field itself. In Euclidean TEAG instantiations, the MVEE-induced metric family $\{g_\alpha\}$ provides the local geometry on which this front propagation occurs. The PCRB (Section 4.3) bounds the admissible rate of contraction of the associated volume form: the total entropy reduction per step cannot exceed $\frac{n}{2} \log(1 - I_k)^{-1}$, where I_k is the Choquet integral of per-hypothesis surprisal against the prior possibility capacity.

Theorem 3.30 (Tropical Hamilton–Jacobi structure: summary). *The following structures are consistent with and follow from the TEAG axioms and the tropical contraction geometry established in Proposition 3.12. Items (i)–(ii) are proved within this paper. Items (iii)–(iv) are structural summaries whose full derivations, including all proofs, are in the companion paper [31]. They are stated here to establish the complete mathematical chain and to make the physics-of-belief interpretation precise.*

(i) Epistemic Lagrangian (proved here).

Define the epistemic Lagrangian

$$L(\Phi, \dot{\Phi}) = T(\dot{\Phi}) - V(\Phi),$$

where $V(\Phi) = \Phi(h)$ is the epistemic potential (the impossibility field) and $T(\dot{\Phi}) = -\frac{n}{2} \log(1 - I_k)$ is the kinetic term: the PCRB-bounded rate of entropy reduction per observation (Theorem 4.5). The action functional over K observations is $\mathcal{A}[\Phi] = \sum_{k=1}^K L(\Phi_k, \dot{\Phi}_k)$. The PCRB is a minimum action principle: no admissible filter reduces the action below the PCRB floor. The ESPF achieves the minimum.

(ii) Tropical Hamiltonian and Hamilton–Jacobi equation (proved here).

The Legendre transform of L with respect to $\dot{\Phi}$ yields the tropical Hamiltonian

$$\mathcal{H}(\Phi, p) = \psi(h),$$

independent of $\nabla\Phi$: the epistemic Hamiltonian is the surprisal field. This reflects the purely evidence-driven character of TEAG contraction — there is no inertia in the epistemic update. The tropical Hamilton–Jacobi equation is

$$\Phi^+(h) = \Phi^-(h) \oplus \mathcal{H}(h) = \max(\Phi^-(h), \psi(h)),$$

which is precisely the TEAG canonical contraction (Proposition 3.12). The TEAG update rule is the solution to the tropical Hamilton–Jacobi equation with Hamiltonian equal to the surprisal field. This is a proved identity within this paper.

(iii) Euler–Lagrange equations and geodesic motion (summary; proved in [31]).

On the epistemic manifold $\mathcal{M}_{\text{ep}}$ with coordinates $\{q^i\}$ induced by the MVEE metric family $\{g_\alpha\}$, the Euler–Lagrange equations of the epistemic action are

$$\ddot{q}^i + \Gamma_{jk}^i \dot{q}^j \dot{q}^k = F^i, \quad F^i = -g^{ij} \partial_j \Phi,$$

where F^i is the epistemic force (the impossibility gradient) and Γ_{jk}^i are the Christoffel symbols of the Levi–Civita connection of g_π . Free inference — no impossibility gradient — follows geodesics of the MVEE metric.

(iv) Christoffel symbols (summary; proved in [31]).

The Christoffel symbols of the MVEE-induced Levi–Civita connection are

$$\Gamma_{jk}^i = \frac{1}{2} g^{il} (\partial_j g_{lk} + \partial_k g_{lj} - \partial_l g_{jk}), \quad g_{ij} = (P_\pi^{-1})_{ij}, \quad g^{ij} = (P_\pi)_{ij}.$$

The impossibility gradient $\nabla\Phi$ deflects inference trajectories from geodesic motion exactly as gravity deflects physical trajectories from inertial paths.

Interpretation.

Items (i)–(iv) together admit a precise interpretation. The TEAG update rule is a max-plus dynamical system. Its governing equations have the same algebraic form as the Hamilton–Jacobi equations of classical mechanics, with hypothesis space playing the role of physical space, the impossibility field playing the role of potential energy, and the surprisal field playing the role of the Hamiltonian. This is a statement about algebraic form, not about physical ontology: inference is not a physical process, and no such claim is made. What is claimed is that the same mathematical structure — max-plus algebra, Legendre duality, geodesic flow, Christoffel symbols — governs both wavefront propagation and belief contraction under evidence. The spaces are different. The governing equations are algebraically the same.

3.7. Choquet Aggregation and Information Functionals

Because TEAG operates under non-additive uncertainty, aggregation of information is naturally expressed through the Choquet integral [6,8] with respect to the possibility capacity induced by π . For any bounded measurable functional $f : H \rightarrow \mathbb{R}$, define

$$\int_H f d\text{Ch}_\pi$$

to be the Choquet integral of f with respect to the possibility measure generated by π . Within the TEAG framework this integral serves three roles: (1) it provides a non-additive aggregation operator for combining multiple sources of evidence; (2) it defines information-content functionals governing admissible contraction of support geometry; and (3) it supplies the measure-theoretic bridge connecting possibilistic inference to probabilistic inference in the collapse limit [29].

In particular, the information content appearing in the PCRB (Section 4.3) is evaluated through a Choquet-type aggregation of compatibility-induced contraction rather than through Shannon or Fisher information defined under additive probability laws.

(Sections 4–8 continue in the full manuscript.)

4. General Results

4.1. Possibilistic Entropy as Ignorance Functional

The central scalar measure of ignorance in TEAG is the possibilistic entropy, defined as the log-geometric mean of the α -cut volume family. It has two geometrically distinct components, each corresponding to a different aspect of the impossibility field Φ .

Definition 4.1 (Possibilistic entropy). Let π be a normalized possibility distribution over a finite support $\{\chi^{(i)}\}_{i=1}^M \subset \mathbb{R}^n$. For each $\alpha \in (0, 1]$, let $V_\alpha = c_n (\det \Pi_\alpha)^{1/2}$ denote the volume of the MVEE of the α -cut $C_\alpha = \{i : \pi^{(i)} \geq \alpha\}$, where $c_n = \pi^{n/2} / \Gamma(n/2 + 1)$ and $\Pi_\alpha \succ 0$ is the MVEE shape matrix. The possibilistic entropy is

$$\mathcal{E}_\pi = \int_0^1 \log V_\alpha d\alpha.$$

Remark 4.2 (Boltzmann, not Shannon). The logarithm in $\mathcal{E}_\pi$ is justified by the Boltzmann tradition, not the Shannon tradition, and the distinction is not merely historical — it reflects a fundamental difference in what the entropy is measuring and what independence structure it respects.

Shannon entropy $-\sum_i p_i \log p_i$ is grounded in the axiomatic theory of communication: it measures the average surprise under a probability distribution, where $-\log p_i$ is the surprise of outcome i . Its additivity across independent random variables follows from the probability sum rule: $p(A \cap B) = p(A) \cdot p(B)$ under independence, so $-\log p(A \cap B) = -\log p(A) - \log p(B)$.

Boltzmann entropy $S = k_B \log W$ measures the logarithm of the volume of the phase space region consistent with a given macrostate. Its additivity is geometric: if two independent state spaces have volumes W_1 and W_2 , the joint volume is $W_1 W_2$, and the logarithm converts the product to a sum. No probability measure is required.

$\mathcal{E}_\pi = \int_0^1 \log V_\alpha d\alpha$ is Boltzmann entropy, and the independence that justifies it is intrinsic to TEAG. At each update step, the impossibility field Φ has two geometrically independent sources: the prior impossibility $\Phi^-(h)$, encoding the accumulated evidence history before the current observation, and the surprisal $\psi(h) = -\log \kappa(y | h)$, determined entirely by the current measurement geometry. These two fields are independent in the possibilistic sense: the prior is fixed before the observation arrives; the surprisal is fixed by the observation alone. Because they are geometrically independent, the volume of the joint admissible region under both constraints is the product of the marginal volumes. The logarithm converts that product to a sum, yielding additivity of $\mathcal{E}_\pi$ across the two sources without requiring a probability measure or a sum rule.

This is directly visible in the tropical update:

$$\Phi^+(h) = \max(\Phi^-(h), \psi(h)).$$

The max selects whichever source of impossibility is larger at each hypothesis, and the resulting α -cut volume V_α^+ factors through the product of the marginal volumes of the two independent geometric constraints. The Boltzmann logarithm converts that product to a sum — exactly what $\mathcal{E}_\pi$ integrates.

Shannon entropy cannot play this role for a structural reason: $-\log p$ requires the probability sum rule, which holds only under additive measures. Possibility theory uses the min rule for conjunction — $\pi(A \cap B) = \min(\pi_A, \pi_B)$ — under which $-\log$ is not additive across independent sources. The independence that matters in TEAG is geometric (volumes multiply) not probabilistic (probabilities multiply), and Boltzmann is the entropy that respects it.

This is not a limitation of possibilistic entropy. It is a precise statement about which entropy concept is appropriate for which uncertainty structure. Shannon entropy is the right tool when uncertainty is probabilistic and additive. Boltzmann entropy is the right tool when uncertainty is geometric and bounded. TEAG operates in the second regime.

The possibilistic entropy decomposes into two geometrically distinct terms [30]:

$$\mathcal{E}_\pi = \underbrace{\log V + \kappa_n}_{\text{basin volume term}} + \underbrace{\int_0^1 \log \frac{V_\alpha}{V} d\alpha}_{\text{impossibility gradient term} \leq 0}, \quad (7)$$

where $V = c_n(\det \Pi)^{1/2}$ is the MVEE volume of the full support and $\kappa_n = \log c_n$.

Basin volume term.

The term $\log V = \frac{1}{2} \log \det \Pi$ is determined entirely by the outermost α -cut — the MVEE of the full surviving support. In the language of the impossibility field Φ (Definition 3.10), this is the volume enclosed by the PCRB equipotential surface (Remark 3.23): how large is the admissible basin at its boundary? This is the Popperian term, measuring the *extent* of remaining ignorance. It is minimized by selecting survivors geometrically closest to the evidence-consistent locus — the minimax entropy selection rule of the ESPF.

Impossibility gradient term.

The term $\int_0^1 \log(V_\alpha/V) d\alpha \leq 0$ is determined by how steeply the impossibility field Φ rises from the basin center toward the equipotential boundary. When π is uniform, Φ is flat within the support and all α -cuts have equal volume: $V_\alpha = V$ for all α , and the gradient term is zero. When π is concentrated — some hypotheses are far more admissible than others — Φ rises steeply within the basin, $V_\alpha \ll V$ for large α , and the gradient term is strictly negative, reducing $\mathcal{E}_\pi$ below $\log V$. This term measures the *texture* of ignorance: not just how large the basin is, but how sharply the impossibility field sculpts it from within.

The gradient term is minimized by the compatibility-based possibility assignment of the ESPF: assigning $\pi^{(i)} = \text{Comp}^{(i)} = e^{-q^{(i)}/2}$ concentrates admissibility on hypotheses closest to the evidence, steepening Φ within the basin and reducing the gradient term simultaneously with the basin volume term.

What the ESPF minimizes.

A filter that prunes the support but reassigns uniform possibility reduces the basin volume term while leaving the gradient term at zero — it reduces the extent of ignorance but discards the texture. The ESPF reduces both simultaneously. This is why $\mathcal{E}_\pi$ is the correct optimality criterion: it is the only scalar that captures the full geometry of the impossibility field, both the size of the admissible basin and the shape of Φ within it.

The possibilistic entropy is the log-geometric mean of $\{V_\alpha\}$ in the Hölder mean hierarchy: $\mathcal{E}_\pi = \log M_0(\{V_\alpha\})$. The basin volume term $\log V$ is $\log M_{+\infty}$ — the coarsest instrument, sensitive only to the outermost boundary. The Kalman MMSE criterion is $\log M_0$ evaluated under Gaussian α -cut geometry. The Hölder ordering $M_0 \leq M_{+\infty}$ establishes $\mathcal{E}_\pi$ as the more stable, more informative criterion: integrated ignorance — combining both basin volume and impossibility gradient — is intrinsically less sensitive to boundary perturbation than the supremum alone [30].

Remark 4.3 (Affine invariance). *All entropy computations are performed in a normalized coordinate system induced by the MVEE metric. Specifically, $V_\alpha = c_n(\det \Pi_\alpha)^{1/2}$ is evaluated after whitening by the MVEE Cholesky factor L_α with $\Pi_\alpha = L_\alpha L_\alpha^\top$, ensuring that $\mathcal{E}_\pi$ is invariant under affine transformations of the state space. This is the same whitening implicit in the ESPF's innovation geometry [28] and is required for geometric comparability of α -cut volumes across iterations.*

4.2. Possibilistic Information Content

The information content of an observation in TEAG is not Shannon entropy evaluated under a probability measure. It is the Choquet integral of the per-hypothesis surprisal with respect to the prior possibility distribution as capacity.

Definition 4.4 (Possibilistic information content [30]). *Let $\{\chi^{(i)}\}$ be the predicted support with prior possibility $\pi_{k|k-1}$ and let $q_k^{(i)} = \|L_e^{-1}(y_k - h(\chi^{(i)}))\|^2$ 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 $\pi_{k|k-1}$:*

$$\bar{S}_k = (\mathcal{C}) \int \frac{1}{2}q_k^{(i)} d\pi_{k|k-1} = \sup_{i=1,\dots,M} \min\left(\frac{1}{2}q_k^{(i)}, \pi_{k|k-1}^{(i)}\right),$$

where the second equality is the standard reduction of the Choquet integral under a possibility measure [36]. The possibilistic information content of observation y_k is

$$I_k = 1 - e^{-\bar{S}_k} \in [0, 1).$$

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 . Crucially, the capacity used in the Choquet integral is the *prior* possibility $\pi_{k|k-1}$, not the current compatibility: the information content is epistemically grounded in accumulated evidence history, not the current measurement alone.

This is the information content that appears in the PCRB — the bound on the rate of admissible impossibility deformation per observation.

4.3. Possibilistic Cramér–Rao Bound

Theorem 4.5 (PCRB [30]). *Let $\mathcal{F} \in F$ be any epistemically admissible filter. At step k , let $\mathcal{E}_{\pi,k|k-1}$ and $\mathcal{E}_{\pi,k|k}$ denote the pre- and post-update possibilistic entropies, and I_k the possibilistic information content (Definition 4.4). Then*

$$\mathcal{E}_{\pi,k|k} \geq \mathcal{E}_{\pi,k|k-1} + \frac{n}{2} \log(1 - I_k). \quad (8)$$

The bound holds for any epistemically admissible filter regardless of its internal implementation.

In the language of the impossibility field, (8) is a constraint on the admissible rate of deformation of the epistemic geoid per observation. The left side is the possibilistic entropy after update — the integrated log-volume of the impossibility field's sublevel-set family. The right side is the pre-update entropy shifted by a term governed entirely by I_k , the Choquet-aggregated surprisal of the observation.

Three structural properties of the PCRB are immediate from its form:

1. **The bound is tight.** The ESPF achieves equality in (8) in the falsification regime, making it the unique optimal filter within the class $\mathcal{F}$ of epistemically admissible estimators (Definition 2.1) [30].
2. **The bound governs integrated ignorance, not pointwise volume.** The PCRB constrains $\int_0^1 \log V_\alpha d\alpha$ simultaneously across all α -levels. A filter may reduce V_α sharply at some levels while barely changing others, but the net reduction in $\mathcal{E}_\pi$ cannot exceed $\frac{n}{2} \log(1 - I_k)^{-1}$.
3. **Selection and bounding are deliberately decoupled.** The prior possibility $\pi_{k|k-1}$ enters only in the Choquet capacity defining $\bar{S}_k$ — a bounding quantity, not a selection criterion. Survivor selection in the ESPF depends solely on $q_k^{(i)}$. This separation preserves the Popperian character of the filter: the rate at which the epistemic geoid may contract is governed by epistemically weighted surprisal, but the direction of contraction is determined by evidence alone.

Remark 4.6 (PCRB as geoid curvature bound). *In the geometric language of Section 3.4, the PCRB is a curvature-like constraint on the deformation of the epistemic geoid. Just as curvature bounds in differential geometry control how volumes evolve under geodesic flow, the PCRB controls how the nested sublevel-set family $\{H_\alpha\}$ of the impossibility field Φ may contract under evidence. The asymmetric rate limits $r^+ = 1.15$, $r^- = 0.97$ enforced in the ESPF [28] operationalize this bound: fast metric expansion (embrace of ignorance) is permitted, while fast metric contraction (assertion of certainty) is restricted in proportion to the Choquet information content of the observation.*

4.4. Gaussian Collapse and Recovery of Probability Theory

The TEAG object has a natural limiting behavior as the impossibility field contracts to zero width. This limit recovers probability theory as the collapse geometry of possibilistic inference [29].

Theorem 4.7 (Choquet-to-Lebesgue convergence [29]). *Let $\{\pi_t\}_{t \geq 0}$ be a sequence of normalized, consonant possibility distributions with associated credal sets $\{\mathcal{P}_{\pi_t}\}$. Suppose: (A1) consonance; (A2) credal contraction: $\sup_A [\Pi_t(A) - N_t(A)] \rightarrow 0$; (A3) L^1 convergence: $\pi_t \rightarrow p^*$ in $L^1(\mu)$; (A4) domination. Then for every bounded measurable f ,*

$$\int f d\text{Ch}_{\pi_t} \longrightarrow \int f(x) p^*(x) d\mu(x) \quad \text{as } t \rightarrow \infty.$$

In TEAG terms, epistemic collapse occurs when the impossibility field Φ flattens to zero outside a single admissible region: the credal set contracts to a singleton, the Choquet integral converges to the Lebesgue integral, and the TEAG object collapses to a classical probability space. Probability

theory is not a competing framework — it is the limiting geometry of TEAG when epistemic width $W = \int (\pi - n) d\mu$ approaches zero.

The ESPF recovers the Kalman filter in this limit: as $\pi_k \rightarrow \mathcal{N}(\hat{x}_k, \Sigma_k)$, the possibilistic entropy satisfies $\mathcal{E}_\pi \rightarrow \frac{1}{2} \log \det \Sigma_k + \text{const}(n)$, so minimizing $\mathcal{E}_\pi$ is asymptotically equivalent to minimizing $\det \Sigma_k$ — the Kalman MMSE criterion [30]. This is convergent optimality, not hierarchical containment: the ESPF and the Kalman filter are optimal solutions to categorically different problems that agree when their domains of applicability coincide.

5. Domain Instantiations

Each instantiation of the TEAG object is obtained by specifying the five components $(H, \pi, \{H_\alpha\}, C, A)$ for a particular inference domain and verifying the TEAG axioms. This section works through the four instantiations of Table 1 explicitly, identifying in each case the hypothesis space, the impossibility field Φ , the contraction operator, the commitment rule, and the geometry of the admissible basin.

5.1. Instantiation I: Recursive State Estimation (ESPF)

The ESPF [28] is the foundational TEAG instantiation and the domain within which the impossibility field, tropical contraction, and the PCRb were first operationalized.

TEAG components.

- $H = \{\chi^{(i)}\}_{i=1}^M \subset \mathbb{R}^n$: a finite set of support points spanning the admissible state space, generated by a Smolyak sparse grid at level ℓ . At level $\ell = 3$, $M = 2n^2 + 1$ points; the minimum survivor count $N_{\min} = 2n + 1$ is a direct consequence of the well-posedness of $\mathcal{E}_\pi$.
- $\pi : H \rightarrow (0, 1]$: ordinal state admissibility, initialized uniform and updated conjunctively. After the first measurement, $\pi^{(i)} = \text{Comp}^{(i)} = e^{-q^{(i)}/2}$ where $q^{(i)} = \|L_e^{-1}e^{(i)}\|^2$ is the whitened squared innovation of hypothesis i .
- $\{H_\alpha\}$: nested α -cuts of the state support cloud, whose MVEE family $\{\Pi_\alpha\}$ defines the local metric $g_\alpha(u, v) = u^\top \Pi_\alpha^{-1}v$ and the impossibility field geometry.
- C : the conjunctive compatibility update $\hat{\pi}^{(i)} = \min(\pi^{(i)}, \text{Comp}^{(i)})$ followed by max-rescaling. In impossibility-field coordinates this is tropical addition $\Phi^+ = \Phi^- \oplus \psi$, where $\psi^{(i)} = \frac{1}{2}q^{(i)}$ is the surprisal — the negative log compatibility, exactly (Remark 3.24).
- A : the whitened minimax medoid — the support point minimizing worst-case whitened distance to all other survivors under the MVEE metric. This is the center of the epistemic geoid (Remark 3.23).

Impossibility field geometry.

The admissible basin is an ellipsoid in $\mathbb{R}^n$ defined by the MVEE of the surviving support. The impossibility field $\Phi(h) = -\log \pi(h)$ rises from zero at the most admissible hypothesis to infinity at falsified hypotheses. The PCRb constrains the rate of basin contraction per observation via $\mathcal{E}_{\pi, k|k} \geq \mathcal{E}_{\pi, k|k-1} + \frac{n}{2} \log(1 - I_k)$, where I_k is the Choquet integral of per-hypothesis surprisal against the prior possibility capacity.

Optimality.

The ESPF is the unique filter in the epistemically admissible class that achieves the PCRb with equality in the falsification regime: minimum- q survivor selection minimizes the basin volume term of $\mathcal{E}_\pi$, and compatibility-based possibility assignment simultaneously minimizes the impossibility gradient term [30].

5.2. Instantiation II: Measure-Theoretic Collapse (GoK)

The Geometry of Knowing [29] instantiates TEAG on a measurable space, proving that possibilistic inference converges to probabilistic inference when the impossibility field contracts to zero width. This is the instantiation that establishes probability theory as the collapse limit of TEAG, not its primitive.

TEAG components.

- $H = \mathcal{X}$: a σ -finite measurable space with reference measure μ .
- $\pi : \mathcal{X} \rightarrow [0, 1]$: a normalized, upper-semicontinuous, consonant possibility distribution. Consonance ensures the necessity kernel $n(x) \equiv \pi(x)$ exists pointwise and the aggregate epistemic width $W = \int_{\mathcal{X}} (\pi - n) d\mu$ is well-defined.
- $\{H_\alpha\} = \{\mathcal{X}_\alpha\}$: the nested α -cut family $\mathcal{X}_\alpha = \{x : \pi(x) \geq \alpha\}$, which under consonance forms a decreasing family of closed sets. These are the sublevel sets of the impossibility field $\Phi(x) = -\log \pi(x)$.
- C : the credal contraction operator. Evidence contracts the credal set $\mathcal{P}_\pi = \{P : N(A) \leq P(A) \leq \Pi(A), \forall A\}$ by tightening the possibility–necessity envelope. The epistemic collapse condition $\Pi(A) = N(A)$ for all A is the limit $\Phi \rightarrow 0$ outside a single admissible region.
- A : in the collapse limit, the unique probability measure P^* to which the credal set contracts — the singleton commitment that possibility theory converges to when epistemic width vanishes.

Impossibility field geometry.

The impossibility field $\Phi(x) = -\log \pi(x)$ defines the credal envelope: the width $W = \int (\pi - n) d\mu$ is the integrated gap between the possibility field and its necessity dual, measuring how far Φ is from the flat zero-width limit. As $W \rightarrow 0$, Φ flattens to zero outside the true state region and the Choquet integral converges to the Lebesgue integral (Theorem 4.7).

The epistemic discipline result.

Probability should be used only when $W < W_{\text{crit}}$: when the impossibility field has contracted sufficiently that the Gaussian approximation error is within the application risk budget. Below this threshold the TEAG object collapses to a probability space and the ESPF recovers the Kalman filter. Above it, possibilistic inference is required [29].

5.3. Instantiation III: Minimax-Entropy Optimality (ESPF Optimality)

The ESPF Optimality paper [30] instantiates TEAG at the level of the α -cut volume family itself, characterizing which contraction operator C is uniquely justified by the minimax-entropy criterion within $\mathcal{F}$ (Definition 2.1). This is the instantiation that proves the PCRB and establishes the ESPF as the unique optimal filter in that class.

TEAG components.

- $H = \{V_\alpha\}_{\alpha \in (0,1]}$: the family of α -cut volumes, treated as the hypothesis space for the optimality analysis. Each $V_\alpha = c_n(\det \Pi_\alpha)^{1/2}$ is a geometric volume.
- π : the α -indexed admissibility, encoding how the possibility distribution concentrates across α -levels.
- $\{H_\alpha\}$: the nested α -cut volume manifold, whose geometry is characterized by the Hölder mean hierarchy $M_p(\{V_\alpha\})$. The possibilistic entropy $\mathcal{E}_\pi = \log M_0(\{V_\alpha\})$ is the $p = 0$ evaluation; the Popperian minimax bound $\log \det(\text{MVEE})$ is $\log M_{+\infty}$.
- C^* : the unique optimal contraction operator — the conjunctive compatibility update with minimum- q survivor selection and compatibility-based possibility assignment. This is the only C in the epistemically admissible class that achieves the PCRB with equality in the falsification regime.
- A : the PCRB-bounded commitment — the whitened minimax medoid, proven here to be the unique commitment minimizing $\mathcal{E}_\pi$ within the evidence-only class.

The optimality result in TEAG language.

The ESPF minimizes the total ignorance functional $\mathcal{E}_\pi$ by simultaneously reducing both components of the entropy decomposition (7): the basin volume term (via minimum- q survivor selection, which minimizes $\log \det(\text{MVEE})$) and the impossibility gradient term (via compatibility-based possibility assignment, which steepens Φ within the basin). No other admissible filter reduces both

simultaneously. The Hölder ordering $M_0 \leq M_{+\infty}$ implies that $\mathcal{E}_\pi$ is a more stable and informative criterion than the minimax bound alone — a direct consequence of the Boltzmann additivity of log-volumes across independent geometric constraints (Remark 4.2).

5.4. Instantiation IV: Possibilistic Language Generation (PLM)

The Possibilistic Language Model [32] instantiates TEAG over vocabulary hypothesis spaces, replacing probabilistic token distributions with possibilistic compatibility fields. It is the instantiation that demonstrates TEAG's reach beyond physical dynamical systems into the domain of language and meaning.

TEAG components.

- $H = \mathcal{V} = \{v_1, \dots, v_{|\mathcal{V}|}\}$: the vocabulary, treated as an epistemic hypothesis space at each generation step k . Tokens in $\mathcal{V} \setminus \mathcal{X}_k$ are categorically excluded — not merely improbable, but falsified.
- $\pi_k : \mathcal{V} \rightarrow [0, 1]$: the vocabulary possibility distribution at step k , encoding ordinal admissibility of each token as the next in sequence. $\pi_k(v) = 0$ means v has been falsified; $\pi_k(v) = 1$ means v is maximally consistent with the context.
- $\{H_\alpha\}$: nested admissible token clouds $\Xi_k^\alpha = \{e_j \in \Xi_k : \pi_k(v_j) \geq \alpha\}$ in embedding space $\mathbb{R}^d$, whose MVEE family defines the vocabulary epistemic volume $V_\alpha = \text{Vol}(\text{MVEE}(\Xi_k^\alpha))$.
- C: Epistemic Possibilistic Attention (EPA) — a falsification-driven attention operator that gates keys by admissible innovation geometry rather than likelihood weighting. The attention innovation residual $r_i = q - k_i$ is whitened by the MVEE of the key cloud, yielding surprisal $S_i = \frac{1}{2} \|L_e^{-1} r_i\|^2$ and compatibility $\kappa_i = e^{-S_i}$. The conjunctive update $\tilde{\pi}^{(i)} = \min(\pi_k^{(i)}, \kappa_i)$ falsifies keys whose geometric surprisal exceeds the admissible innovation bound. This is the tropical update $\Phi^+ = \Phi^- \oplus \psi$ instantiated in embedding space.
- A: the minimax medoid token — the vocabulary token whose embedding is most geometrically central in the surviving admissible cloud $\Xi_{k'}$, committed as the generated token $\hat{v}_k = \arg \min_{v \in \mathcal{X}_k} \max_{v' \in \mathcal{X}_k} \|e_v - e_{v'}\|_{\Sigma_{\Xi_k}^{-1}}$.

Impossibility field geometry.

The impossibility field $\Phi_k(v) = -\log \pi_k(v)$ lives in the token embedding space $\mathbb{R}^d$. Each attention step deforms Φ_k via tropical addition of the attention surprisal field, progressively elevating the impossibility of tokens whose embedding geometry is inconsistent with the context query. The PCRb constrains the rate of vocabulary support contraction per generation step, preventing inadmissible epistemic collapse to a singleton token.

Convergent optimality.

As epistemic width $W_k \rightarrow 0$ (the context becomes unambiguous), $\pi_k \rightarrow \mathcal{N}(e_{\hat{v}}, \Sigma_k)$, the EPA operator converges to scaled dot-product softmax attention, the possibilistic cross-entropy converges to standard cross-entropy, and the PLM recovers a standard transformer LLM [32]. This is convergent optimality: the PLM and the standard LLM are optimal under different epistemic commitments and agree precisely when their domains coincide.

What the PLM instantiation reveals.

By placing language generation within TEAG, the PLM exposes that the standard LLM's softmax is an instance of probabilistic closure applied to a domain where admissibility is bounded and ordinal. The Boltzmann entropy $\mathcal{E}_\pi = \int_0^1 \log V_\alpha d\alpha$ is the correct optimality criterion in both the state estimation and language generation domains because in both cases prior possibility (accumulated context) and surprisal (current token evidence) are geometrically independent sources whose joint admissible volume is the product of the marginal volumes — and the logarithm converts that product to a sum (Remark 4.2).

Remark 5.1 (Asymmetry of instantiations). *The four instantiations are not equally derived. The ESPF is the native TEAG instantiation: the framework was developed in direct correspondence with the ESPF architecture, and the TEAG axioms are satisfied by construction. The GOK instantiation is similarly native, as the measure-theoretic collapse result was developed within the same epistemic program. The PLM instantiation is interpretive: the TEAG structure maps coherently onto the language generation setting, but the mapping is less structurally forced than for state estimation. Whether the embedding geometry of token spaces faithfully tracks semantic admissibility is an open empirical question (Open Problem 6 of [32]). The PLM instantiation should therefore be understood as a principled structural mapping rather than a derivation from first principles of language. Readers should weight these instantiations accordingly.*

5.5. Cross-Instantiation Transfer

Identifying the TEAG object as the shared primitive of the four instantiations enables results that would be inaccessible if the domains were treated independently. Three transfer results follow immediately.

Universal PCRB.

The PCRB (Theorem 4.5) was proved at the level of the TEAG object — for any epistemically admissible contraction operator on any hypothesis space. It therefore holds universally across all four instantiations: it constrains entropy reduction per measurement step in recursive state estimation, per observation in the measure-theoretic collapse, per optimality step in the volume manifold, and per generation step in language modeling. The PCRB is not a filter-specific result; it is a universal geometric bound on the rate of impossibility deformation under evidence.

Boltzmann entropy as universal ignorance functional.

The justification for $\mathcal{E}_\pi = \int_0^1 \log V_\alpha d\alpha$ as the correct entropy in all four domains is the same: prior impossibility and surprisal are geometrically independent in each domain, their joint admissible volume is the product of marginal volumes, and the logarithm converts that product to a sum. This holds for state vectors in $\mathbb{R}^n$, for credal sets on measurable spaces, for α -cut volumes on the optimality manifold, and for token embeddings in $\mathbb{R}^d$.

Convergent optimality.

In each domain, the TEAG object converges to a classical framework in the collapse limit: the ESPF converges to the Kalman filter, the possibilistic credal set converges to a probability measure, the PLM converges to a standard transformer. In every case this convergence is not hierarchical containment but convergent optimality: two frameworks optimal under different epistemic commitments that agree when their domains of applicability coincide. Probability theory is the geometry of TEAG at zero epistemic width.

6. Composition Theory

The four instantiations of Section 5 treat TEAG objects as operating independently over their respective hypothesis spaces. Real-world reasoning systems, however, must simultaneously maintain epistemic integrity across multiple domains: a dynamical state must be estimated, semantic representations must be contracted, contextual adequacy must be audited, and computational attention must be prioritized. This section develops the theory of *compositional* TEAG: how multiple TEAG objects operating over distinct hypothesis spaces can be coupled into a unified epistemic system while preserving the axioms, the PCRB, and the convergent optimality structure of each constituent.

6.1. The Composition Problem

A single TEAG object $\mathcal{E} = (H, \pi, \{H_\alpha\}, C, A)$ governs inference within one domain. When inference must span multiple domains simultaneously — state estimation and language generation, or semantic contraction and contextual adequacy auditing — the question arises: under what conditions can multiple TEAG objects be composed into a system that is itself epistemically admissible?

The challenge is not merely notational. Different domains carry different hypothesis spaces, different impossibility field geometries, and different contraction operators. Naive composition may produce a system that violates Popperian monotonicity across domains, creates self-reinforcing impossibility assignments, or generates commitment rules that are inadmissible in the joint space. The theory of compositional TEAG must specify the coupling conditions that prevent these failures.

Definition 6.1 (Compositional TEAG system). *A compositional TEAG system is a tuple*

$$\mathfrak{E} = (\mathcal{E}_1, \dots, \mathcal{E}_K, \Pi_{\mathfrak{E}}, \mathcal{W}),$$

where each $\mathcal{E}_k = (H_k, \pi_k, \{H_{\alpha,k}\}, C_k, A_k)$ is a TEAG object over domain k , $\Pi_{\mathfrak{E}}$ is a shared epistemic capacity satisfying axioms **A1–A5** of the framework space, and $\mathcal{W}$ is a coupling operator that propagates epistemic width information across domains. The system is admissible if each $\mathcal{E}_k$ satisfies the TEAG axioms and $\mathcal{W}$ preserves Popperian monotonicity: evidence can only increase impossibility, never decrease it, in any domain.

6.2. The ARC/CRA/DAO Instantiation

We develop here a concrete compositional TEAG system operating over three coupled hypothesis spaces: a semantic configuration space (ARC), a representational dimension space (CRA), and a reasoning artifact space (DAO). This architecture is realized in the *Semantic Turning Point Detector* [33], an open-source multi-agent system that wraps arbitrary base LLMs and has been empirically validated across five language models ranging from 1.7B to 30B+ parameters. Together the three TEAG objects form a second-order epistemic controller — a system that not only performs inference but continuously monitors, audits, and steers its own representational adequacy. We derive the TEAG quintuple for each component and establish the coupling structure.

TEAG object 1: ARC (Adaptive Recursive Convergence).

ARC operates over a semantic configuration space (X, ρ) .

- $H_1 = X$: the space of semantic configurations $x = (x_1, \dots, x_m)$, where each x_i is the state of a concurrent reasoning agent.
- π_1 : ordinal admissibility of each configuration under the shared atomic memory tensor $A^{(k)} = \Phi(x_1^{(k)}, \dots, x_m^{(k)})$. The impossibility field $\Phi_1(x) = -\log \pi_1(x)$ is the epistemic divergence $D_{\Pi}(x, x^*)$ from the fixed point.
- $\{H_{\alpha,1}\}$: nested α -cuts of the configuration space, whose contraction under F is bounded by the Banach fixed-point theorem.
- C_1 : the synchronous sweep $F(x) = (f_1(x_1, A), \dots, f_m(x_m, A))$, which is a ρ -contraction. In impossibility-field coordinates, each sweep is a tropical update $\Phi_1^+ = \Phi_1^- \oplus \psi_1$ where ψ_1 is the surprisal of the current configuration against the atomic memory.
- A_1 : the fixed-point commitment x^* — the configuration at which all agents have converged to a stable, internally coherent semantic state.

TEAG object 2: CRA (Cascading Re-Dimensional Attention).

CRA operates over a representational dimension space and serves as the meta-level contraction operator: it monitors whether the current ARC hypothesis space H_1 is epistemically adequate, and escalates to a higher-dimensional space H'_1 when it is not.

- $H_2 = \{0, 1, \dots, n_{\max}\}$: the space of representational dimensions, each corresponding to a distinct semantic scale at which ARC operates.
- π_2 : admissibility of each dimension n , derived from the retro-mask confidence $\text{Conf}(n) \in [0, 1]$. A dimension is admissible if $\text{Conf}(n) \geq \tau_{\text{good}}$; its impossibility $\Phi_2(n) = -\log \pi_2(n)$ rises as confidence falls.
- C_2 : the retroactive coherence assessment. CRA allows unrestricted inference in dimension n , then scores each reasoning fragment via internal coherence C_i and external grounding G_i , yielding

Conf_{in} and Conf_{out} . This is a conjunctive update: dimension n survives only if both internal and external coherence are above threshold.

- A_2 : the escalation commitment — either remain at dimension n (converged) or commit to dimension $n + 1$ (saturated). Saturation occurs when $\text{Conf}_{\text{in}} < \tau_{\text{min}}$ or the gap $\text{Conf}_{\text{out}} - \text{Conf}_{\text{in}} < -\tau_{\Delta}$, signaling that the current hypothesis space H_1 is inadequate and must be expanded.

TEAG object 3: DAO (Differentiating Attentional Orientation).

DAO operates over a space of reasoning artifacts and implements possibilistic salience selection.

- $H_3 = \mathcal{A}_k$: the set of ARC/CRA-validated reasoning artifacts at step k .
- π_3 : the significance field $\phi(a_i) = \lambda\sigma(a_i) + (1 - \lambda)\varphi(a_i)$, where $\sigma(a_i)$ is epistemic significance derived from CRA confidence and $\varphi(a_i)$ is thematic saliency. This is the ordinal admissibility of artifact a_i as a focus of computational attention.
- C_3 : Choquet integration under a consonant possibility measure $\mu(S) = \max_{j \in S} \varphi(a_j)$, computing epistemic necessity $\eta(a_i) = \mathcal{C}_{\mu}(\delta_{\text{pre}}, \delta_{\text{post}})$ as the non-additive aggregate of semantic drift. This is the TEAG contraction in salience space: low-significance artifacts are falsified and excluded from the computational budget.
- A_3 : the budget-constrained minimax selection $\Gamma_B(\mathcal{A}_k) = \arg \max_{S \subseteq \mathcal{A}_k} \sum_{a \in S} w(a)$ subject to $\sum_{a \in S} \kappa(a) \leq B$, where $w(a_i) = \alpha\phi(a_i) + (1 - \alpha)\eta(a_i)$. This is the commitment to the most epistemically significant admissible artifacts within the computational budget.

6.3. The Coupling Operator and Shared Width Monitor

The three TEAG objects are coupled through a shared epistemic width W_f and a confidence-sensitive Hölder dial α_f . This coupling is the mathematical actuator for compositional self-verification.

The aggregate epistemic width

$$W_f = \int_{\mathcal{X}} (\pi_f(x) - n_f(x)) d\mu(x) \quad (9)$$

measures the global gap between possibility and necessity across the shared capacity Π_f . As $W_f \rightarrow 0$, Π_f collapses to a singleton probability measure (the TEAG Gaussian collapse result, Theorem 4.7), and the compositional system reduces to additive probabilistic inference.

The coupling operator $\mathcal{W}$ converts width into a unit-interval confidence capacity and couples it to the Hölder aggregation dial:

$$\nu_{\text{Conf}}(W_f) = \max\left\{0, 1 - \frac{W_f}{W_{\text{crit}}}\right\}, \quad (10)$$

$$\alpha_f = 1 + \frac{1 - \nu_{\text{Conf}}(W_f)}{\varepsilon + \nu_{\text{Conf}}(W_f)}, \quad \varepsilon \in (0, 10^{-2}]. \quad (11)$$

As $W_f \rightarrow 0$, $\nu_{\text{Conf}} \rightarrow 1$ and $\alpha_f \rightarrow 1$ — the additive limit where Choquet integration reduces to Lebesgue integration (Theorem 4.7). As $W_f \rightarrow W_{\text{crit}}$, $\alpha_f \rightarrow 1 + \varepsilon^{-1}$ — the near-maxitive limit where conjunction is governed by the min rule of possibility theory.

This dial is the shared epistemic signal connecting all three TEAG objects:

1. **ARC** uses α_f to determine whether its fixed-point commitment A_1 is epistemically justified. When $\alpha_f \approx 1$ (low width), additive convergence criteria apply. When $\alpha_f \gg 1$ (high width), the system remains in possibilistic caution until the fixed point is verified.
2. **CRA** sets W_{crit} — the risk tolerance that defines the admissible additive window — and monitors W_f via the retro-mask confidence Conf . High Conf implies low W_f , guiding the system toward $\alpha_f \rightarrow 1$. Low Conf implies higher width, forcing conservative aggregation.
3. **DAO** consumes ν_{Conf} to select α_f for its Choquet integration, ensuring that artifact prioritization respects the current epistemic regime: possibilistic when the system is uncertain, additive when collapse is certified.

Remark 6.2 (The coupling as epistemic geoid deformation). *In the language of Section 3.4, the coupling operator $\mathcal{W}$ defines a shared epistemic geoid across all three TEAG domains. The width W_f measures how far the system is from the epistemic collapse surface. The Hölder dial α_f is the shared curvature parameter of that geoid: it is flat ($\alpha_f = 1$) when the system has collapsed to certainty and curved ($\alpha_f \gg 1$) when significant impossibility remains. Each of the three TEAG objects deforms its local geoid through its contraction operator, and the coupling propagates that deformation across domains via W_f .*

6.4. Compositional Convergence

The compositional TEAG system converges in finite time under the following conditions, which correspond precisely to the TEAG axioms applied at the compositional level.

Theorem 6.3 (Compositional convergence). *Let $\mathfrak{E} = (\mathcal{E}_1, \mathcal{E}_2, \mathcal{E}_3, \Pi_{\mathfrak{E}}, \mathcal{W})$ be a compositional TEAG system satisfying Definition 6.1, where: (i) each C_k is a contraction under its domain metric; (ii) $\mathcal{W}$ is monotone in W_f ; (iii) DAO's perturbation is Lipschitz with coupling gain $\lambda_{\text{dao}} \leq \eta < \lambda_{\text{arc}}/L$; (iv) representational depth is bounded by $n_{\text{max}} < \infty$. Then the composite operator $\Psi_{\text{triad}} = \lambda_{\text{arc}}\Phi_{\text{arc}} + \lambda_{\text{cra}}\Psi_{\text{cra}} + \lambda_{\text{dao}}\Omega_{\text{dao}}$ terminates at (x^*, n^*) in at most*

$$T_{\text{max}} = n_{\text{max}} \cdot \left\lceil \frac{\log \varepsilon^{-1}}{\log \kappa^{-1}} \right\rceil$$

iterations, satisfying either high-confidence convergence or budget exhaustion. Furthermore, the shared width satisfies $W_f(T_{\text{max}}) \leq W_f(0)$: evidence can only reduce epistemic width in the compositional system, never increase it.

Proof sketch. Within each domain k , the contraction C_k converges by the Banach fixed-point theorem applied to (H_k, ρ_k) . Popperian monotonicity (Axiom 2.2) applied to each $\mathcal{E}_k$ ensures that $\Phi_k^+ \geq \Phi_k^-$ pointwise, so W_f is non-increasing under each individual contraction. Monotonicity of $\mathcal{W}$ in W_f propagates this property to the coupled system. If CRA detects saturation, escalation to dimension $n+1$ expands H_1 to a strictly larger hypothesis space with lower width $W_{f'} < W_f$, progressing toward collapse. Since escalation consumes one unit of the finite budget n_{max} , the system terminates after at most n_{max} escalations, each requiring at most $\lceil \log \varepsilon^{-1} / \log \kappa^{-1} \rceil$ contraction steps. DAO's Lipschitz perturbation with gain $\lambda_{\text{dao}} < \lambda_{\text{arc}}/L$ is dominated by ARC's contraction rate and does not prevent convergence. $\square$

The final property — W_f is monotone non-increasing — is the compositional analog of Popperian monotonicity (Axiom 2.2): just as evidence can only increase impossibility of individual hypotheses, evidence can only reduce aggregate epistemic width in the compositional system. This is what makes the coupling operator $\mathcal{W}$ epistemically admissible.

6.5. What Compositional TEAG Enables

Beyond ARC/CRA/DAO, the compositional TEAG framework provides a general architecture for any system that must simultaneously maintain epistemic integrity across multiple inference domains. Three properties follow from the framework directly.

Cross-domain PCRB.

Because the PCRB (Theorem 4.5) holds for any epistemically admissible TEAG object, it holds for each constituent of a compositional system. The shared width monitor W_f provides a *joint* constraint: the total entropy reduction across all domains per step is bounded by the sum of per-domain PCRB floors. This prevents any single domain from asserting certainty faster than the evidence warrants, even when other domains have already converged.

Scale equalization.

The compositional convergence theorem (Theorem 6.3) establishes that convergence to canonical turning points is governed by the structure of the TEAG objects and their coupling, not by the parameter

scale of the underlying models. Smaller models equipped with the ARC/CRA/DAO controller converge to the same epistemic fixed points as larger models — empirically confirmed across models ranging from 1.7B to 30B+ parameters in the *Semantic Turning Point Detector* [33] — because self-verification is a property of the contraction structure, not of model size.

Epistemic audit trail.

At every step of a compositional TEAG system, the shared width W_f , the Hölder dial α_f , and the per-domain impossibility fields $\{\Phi_k\}$ provide a continuously available epistemic audit trail. The system cannot assert certainty ($W_f \rightarrow 0$) without that assertion being computable and verifiable. This is the compositional analog of the ESPF’s Epistemic Width Monitor: a multi-domain, multi-scale diagnostic that makes epistemic overconfidence visible rather than silent.

7. New Instantiation Directions

The TEAG object is not limited to the four instantiations developed in Section 5. Any inference domain in which hypotheses are ordered by admissibility, evidence contracts rather than redistributes, and commitment must be epistemically defensible admits a TEAG instantiation. This section sketches three new directions, identifying in each case the hypothesis space, the impossibility field geometry, and the open theoretical obligations.

7.1. Biological Aging as Informational Degradation

Biological aging is conventionally modeled as a stochastic accumulation of damage over time. TEAG suggests a different primitive: aging is the progressive deformation of the biological system’s informational geometry away from a reference state of minimal impossibility.

TEAG components.

Let H be the space of possible molecular and cellular configurations of a biological subsystem (e.g., the epigenome, the proteome, or a tissue-level state). The possibility field $\pi_t : H \rightarrow [0, 1]$ at time t encodes the ordinal admissibility of each configuration relative to the organism’s developmental reference state — the state at which epistemic width $W_t = \int (\pi_t - n_t) d\mu$ is minimized and the impossibility field Φ_t is flattest. The contraction operator C is the biochemical dynamics of the subsystem under environmental and metabolic stress: observations (molecular measurements) falsify configurations incompatible with the current biochemical evidence, contracting the admissible support. The commitment rule A is the dominant phenotypic state — the most geometrically central surviving configuration under the MVEE metric.

The aging hypothesis in TEAG language.

Aging is the monotone increase of the impossibility field’s integrated gradient term — the texture of Φ_t within the admissible basin — rather than a simple reduction in basin volume. Different subsystems age at different rates depending on how rapidly their impossibility fields deform under metabolic and environmental stress. Death occurs when a critical subsystem’s impossibility field contracts to a point — the degenerate collapse $\Phi_t \rightarrow \infty$ outside a singleton — before other subsystems have reached the same limit.

This framing generates a precise empirical prediction: the rate of biological aging in a subsystem is bounded below by a PCRB-like quantity governing how fast informational geometry can degrade per unit time under bounded metabolic perturbation. Subsystems with larger Choquet information content per metabolic cycle (higher I_k) degrade faster; those with smaller I_k are more resilient. Testable proxies include epigenetic clock divergence rates (e.g., Horvath clock [16]), transcriptomic dispersion in longitudinal cohorts such as the Baltimore Longitudinal Study of Aging, and allele-specific expression variance in GTEx.

Open obligations.

The primary open obligation is the formal definition of the compatibility field $\kappa(y|h)$ for biological measurements: what constitutes an admissible molecular configuration relative to a developmental reference, and how measurement noise in high-dimensional omics data maps onto possibilistic surprisal. A formal TEAG instantiation for biological aging would connect the ESPF's possibilistic entropy framework to the information-theoretic aging literature [14].

7.2. Quantum State Tomography

Quantum state tomography recovers an unknown quantum state ρ from measurement statistics. Standard approaches use maximum likelihood estimation or Bayesian inference over density matrices, both of which impose probabilistic closure. TEAG suggests an epistemically conservative alternative.

TEAG components.

Let H be the space of density matrices $\rho \in \mathcal{D}(\mathcal{H})$ on a finite Hilbert space $\mathcal{H}$. The possibility field $\pi : H \rightarrow [0, 1]$ encodes ordinal admissibility of each density matrix relative to the accumulated measurement record: a density matrix is admissible if it is compatible with all observed measurement outcomes, inadmissible if any outcome rules it out. The contraction operator C updates π via compatibility with the Born rule: $\kappa(y|\rho) = \text{Tr}(M_y \rho)$ for measurement operator M_y , falsifying density matrices whose predicted outcome probability is below the admissible threshold. The commitment rule A is the minimax medoid density matrix — the most geometrically central surviving state under the MVEE metric on the space of density matrices.

TEAG vs. standard tomography.

Standard maximum likelihood tomography can produce overconfident estimates when the measurement record is sparse — a failure mode structurally identical to the UKF's silent covariance collapse under unmodeled dynamics. The TEAG instantiation retains the full admissible set of density matrices until the evidence warrants contraction, reporting epistemic width W as a diagnostic of how much quantum state uncertainty remains. The Choquet-to-Lebesgue collapse theorem (Theorem 4.7) provides the formal conditions under which the possibilistic tomographic estimate converges to the maximum likelihood estimate — not as an assumption, but as a provable limit when the measurement record is sufficiently informative.

Open obligations.

The primary obligation is establishing the PCRB in the quantum setting: bounding how rapidly the admissible set of density matrices can contract per measurement, with information content defined via the Choquet integral of Born-rule surprisal against the prior possibility capacity. The connection between possibilistic tomography and the quantum Fisher information bound [15] requires development.

7.3. Multi-Agent Epistemic Planning

Multi-agent planning under uncertainty requires each agent to maintain beliefs about the world and about the beliefs of other agents. Probabilistic approaches (POMDPs, belief-state MDPs) impose probabilistic closure at every level of the belief hierarchy, which is epistemically unjustified when agents have bounded, partial information about each other.

TEAG components.

Let $H^{(i)}$ be the hypothesis space of agent i , encoding possible world states and possible belief states of other agents. The possibility field $\pi^{(i)} : H^{(i)} \rightarrow [0, 1]$ encodes ordinal admissibility of each joint world-belief configuration from agent i 's perspective. The contraction operator $C^{(i)}$ updates $\pi^{(i)}$ via observations of the world and of other agents' actions, falsifying world-belief configurations incompatible with observed behavior. The commitment rule $A^{(i)}$ is the minimax medoid joint configuration — the least inadmissible plan of action given the current admissible set.

The compositional structure.

Multi-agent possibilistic planning is a natural application of compositional TEAG (Section 6): each agent operates as a TEAG object, and the coupling operator $\mathcal{W}$ propagates shared epistemic width through observed actions. When all agents have low width $W^{(i)} \approx 0$, the system collapses to standard Nash equilibrium reasoning. When agents have high width, the compositional system maintains a possibilistic equilibrium: a set of jointly admissible joint plans rather than a single committed strategy. The PCRB bounds how fast each agent can reduce its own uncertainty per observation of other agents' actions.

Open obligations.

Formal definition of the compatibility field $\kappa^{(i)}(y^{(i)}|h^{(i)})$ for multi-agent observations requires a possibilistic theory of belief revision under strategic interaction. The connection to epistemic logic [13] and possibility-based game theory [12] provides a starting point.

8. Related Work

TEAG sits at the intersection of several mature research traditions. This section positions the framework's contributions precisely relative to each, identifying where TEAG extends existing work, where it departs from it, and what it contributes that was previously unavailable.

8.1. Possibility Theory and Imprecise Probability

Possibility theory was introduced by Zadeh [36] as a maxitive, ordinal calculus for reasoning under epistemic uncertainty, and formalized by Dubois and Prade [11] into a comprehensive framework for uncertainty processing. TEAG is grounded in this tradition: the TEAG quintuple's possibility field π is exactly a Zadeh-Dubois-Prade possibility distribution, and the conjunctive update $C(\pi, y)(h) = \min(\pi(h), \kappa(y|h))$ is the standard possibilistic conditioning rule.

TEAG's contribution relative to possibility theory proper is threefold. First, it introduces the impossibility field $\Phi = -\log \pi$ as the natural geometric object for possibility theory, connecting the ordinal structure of π to tropical algebra and differential geometry. Second, it establishes the possibilistic entropy $\mathcal{E}_\pi = \int_0^1 \log V_\alpha d\alpha$ as the correct ignorance functional for geometric hypothesis spaces, with a rigorous Boltzmann justification that distinguishes it from both the U-uncertainty measure of Higashi and Klir [19] and Shannon entropy. Third, it derives the PCRB as a universal lower bound on the rate of admissible support contraction per observation — a result with no prior counterpart in the possibility theory literature.

Walley's theory of imprecise probabilities [35] formalizes the credal set $\mathcal{P}_\pi = \{P : N(A) \leq P(A) \leq \Pi(A)\}$ as the set of all probability measures consistent with current evidence. The broader imprecise probability family [3] encompasses credal sets, p-boxes, and Dempster-Shafer belief functions within a common framework. TEAG is specifically grounded in the *possibilistic* subfamily — the consonant sub-class in which the credal set is generated by a single normalized possibility distribution π and the Choquet integral serves as the natural aggregation operator. This specificity is not a limitation: it provides stronger structural results (the PCRB, the impossibility field geometry) than are available in the general imprecise probability setting. TEAG operates within the credal set framework in the following sense: the epistemic width $W = \int (\pi - n) d\mu$ is a scalar diagnostic on the credal set diameter, and Theorem 4.7 (Choquet-to-Lebesgue convergence) establishes the precise conditions under which the credal set contracts to a singleton. The present work extends Walley's static framework by providing an explicit dynamic contraction mechanism (the TEAG object), a performance bound on that contraction (the PCRB), and a recursive commitment rule (the minimax medoid) that is admissible by construction within the credal set.

Dempster-Shafer theory [26] represents belief through basic probability assignments over subsets of the hypothesis space and combines evidence via Dempster's combination rule. TEAG differs structurally: the TEAG contraction operator satisfies Popperian monotonicity (evidence can only reduce admissibility, never amplify it), while Dempster's rule can amplify belief in subsets. The TEAG

commitment rule is a minimax geometric operation on the surviving support, not a normalized mass assignment. When the focal sets are nested (consonance), Dempster-Shafer theory and possibility theory coincide; TEAG is built specifically on this consonant structure.

8.2. Information Geometry

Classical information geometry [1] assigns to each probability distribution a point on a statistical manifold equipped with the Fisher metric $g_F(u, v) = u^\top \Sigma^{-1} v$, and studies inference as geodesic flow on this manifold. The TEAG epistemic manifold $\mathcal{M}_{\text{ep}}$ is an analogous structure, but with a fundamentally different primitive: points are admissibility configurations $\Pi = (\pi, \{H_\alpha\}, \{P_\alpha\})$ rather than probability distributions, and the metric is induced by the MVEE geometry of α -cuts rather than by Fisher information.

The two structures are related by a limiting containment: in the Gaussian collapse limit where $\pi \rightarrow \mathcal{N}(\hat{x}, \Sigma)$, the MVEE shape matrix P_π converges to Σ and the epistemic metric converges to the Fisher metric. The epistemic manifold therefore contains classical information geometry as a limiting submanifold. Information geometry describes the geometry of certainty; TEAG describes the geometry of the approach to certainty from ignorance.

The PCRB is structurally analogous to the classical Cramér-Rao bound [7,24]: both bound the rate at which uncertainty can be reduced per observation. The classical CRB bounds the variance of any unbiased estimator from below via the Fisher information matrix. The PCRB bounds the rate of possibilistic entropy reduction from below via the Choquet information content I_k . The two bounds are not in a containment relationship — they operate under different uncertainty primitives — but they share the same structural role: constraining how fast evidence can legitimately resolve uncertainty.

8.3. Tropical Geometry

Tropical geometry [22] studies algebraic geometry over the tropical semiring $(\mathbb{R} \cup \{+\infty\}, \min, +)$ or equivalently the max-plus semiring $(\mathbb{R} \cup \{-\infty\}, \max, +)$. It has found applications in optimization, phylogenetics, and neural network expressivity. TEAG's connection to tropical geometry arises through the log-admissibility transformation: under $\Phi = -\log \pi$, the TEAG conjunctive update $\pi^+ = \min(\pi^-, \kappa)$ becomes $\Phi^+ = \max(\Phi^-, \psi)$ — tropical addition of impossibility fields under the $(\max, +)$ semiring.

This connection is not merely notational. It establishes that abductive inference — the progressive elimination of hypotheses incompatible with evidence — is tropical addition of potential fields. The α -cut family $H_\alpha = \{h : \Phi(h) \leq -\log \alpha\}$ evolves as sublevel-set fronts of a tropical potential, with dynamics analogous to Hamilton-Jacobi equations in optimal control [2]. The PCRB is then a curvature-like bound on the admissible rate of contraction of these fronts per observation. To the authors' knowledge, this connection between possibility theory, tropical algebra, and information-theoretic bounds has not been previously established in the literature.

8.4. Optimal Transport

Classical optimal transport [27] equips probability measures with a geometry induced by least-cost mass displacement, defining Wasserstein distances on the space of measures. TEAG defines a structurally distinct geometry on the space of admissibility configurations: rather than transporting probability mass, each observation induces a one-sided deformation of the admissible-support configuration, contracting rather than redistributing. The PCRB constrains the admissible rate of this contraction in the same way that optimal transport geometry bounds the speed of mass displacement under a given cost. Three structural distinctions separate TEAG from optimal transport: motion is one-sided (evidence only contracts); admissibility rather than mass is conserved until falsified; and the geometry lives on support configurations, not probability laws. We call the resulting structure *one-sided support-transport geometry*.

8.5. Bayesian and Particle Filter Optimality

The Kalman filter [21] minimizes mean squared error under Gaussian assumptions and is the optimal linear estimator in its class. Particle filters [17] provide asymptotically exact posterior approximations for non-Gaussian, nonlinear systems. Both are truth-seeking: they assume a generative model, integrate a prior, and produce a posterior whose objective is convergence to the true state in expectation.

TEAG and the ESPF are not truth-seeking. They are impossibility-eliminating: they retain only what the evidence has not ruled out and make no assertion about where truth is. The ESPF is optimal within the class of epistemically admissible evidence-only filters under the possibilistic minimax-entropy criterion — a different optimality criterion applied to a different class of estimators solving a different problem. The convergent optimality result (Section 4.4) establishes that when the world is Gaussian and the model is valid, both the ESPF and the Kalman filter produce the same state estimate by entirely different mechanisms. This is convergent optimality, not containment: the Kalman filter is not a special case of the ESPF, and the ESPF is not a heuristic alternative to the Kalman filter. They are optimal solutions to categorically different problems that agree when their domains of applicability coincide.

Houssineau and Clark [20] and Delande et al. [10] developed subjective random-set formulations connecting Dempster-Shafer belief functions to probability via credal contraction. TEAG complements these domain-specific results by providing a general Choquet-to-Lebesgue convergence theorem on measurable spaces and a universal PCRB that holds across all TEAG instantiations.

8.6. Language Model Uncertainty

Contemporary work on LLM uncertainty quantification — including calibration methods [18], Bayesian neural networks [23], conformal prediction [5], and evidential deep learning [25] — addresses uncertainty within the probabilistic framework. These methods improve distributional calibration but do not question the structural appropriateness of probability as the uncertainty representation for language.

TEAG’s contribution to language modeling, instantiated in the PLM [32], is precisely this structural critique: probabilistic closure is epistemically overcommitted for token uncertainty under bounded admissibility. A token with possibility zero is categorically excluded, not merely improbable. The PLM replaces the softmax output head with a possibilistic compatibility field, replaces attention with a falsification-driven Epistemic Possibilistic Attention operator, and replaces maximum-likelihood training with a possibilistic entropy objective regularized by the PCRB. The PLM is a strict generalization: it recovers a standard transformer in the Gaussian epistemic limit (Section 5.4), providing theoretical coherence with the existing landscape while operating correctly in the broader possibilistic regime.

8.7. Summary of Positioning

Table 3 summarizes the positioning of TEAG relative to the principal related frameworks.

Table 3. Positioning of TEAG relative to related frameworks.

Framework	Uncertainty primitive	Relation to TEAG
Possibility theory [11,36]	Possibility distribution π	Foundation; TEAG adds geometry, PCRB, and impossibility field
Imprecise probability [3,35]	Credal set $\mathcal{P}_\pi$	TEAG is the possibilistic sub-family; adds dynamic contraction and commitment
Dempster-Shafer [26]	Basic probability assignment	Structural departure: TEAG is monotone, consonant, geometric
Information geometry [1]	Fisher metric on distributions	TEAG contains as Gaussian collapse limit
Tropical geometry [22]	Max-plus semiring	TEAG realizes tropical addition as abductive inference
Optimal transport [27]	Wasserstein metric on measures	TEAG is one-sided support-transport geometry
Kalman / particle filters [17,21]	Posterior distribution	Convergent optimality: agree in Gaussian limit, different problems
LLM uncertainty [18,23]	Calibrated probability	TEAG replaces probabilistic closure with possibilistic admissibility

9. Conclusion

This paper has introduced the Theory of Epistemic Abductive Geometry (TEAG), a unified framework for inference under bounded epistemic uncertainty. The central contribution is the identification of a single mathematical object — the TEAG quintuple $\mathcal{E} = (H, \pi, \{H_\alpha\}, C, A)$ — as the shared primitive underlying four previously distinct inference frameworks: the ESPF for recursive state estimation, the Geometry of Knowing for measure-theoretic collapse, the ESPF Optimality theorem for minimax-entropy estimation, and the Possibilistic Language Model for bounded language generation.

What TEAG establishes.

The theoretical contributions of this paper are five.

First, TEAG introduces the *impossibility field* $\Phi(h) = -\log \pi(h)$ as the natural geometric object for representing epistemic uncertainty. Where possibility theory encodes what remains admissible, the impossibility field encodes its dual: the continuous, non-negative degree to which each hypothesis has been eliminated by evidence. Inference in TEAG is the progressive deformation of this field under evidence — tropical addition $\Phi^+ = \Phi^- \oplus \psi$, where $\psi(h) = -\log \kappa(y|h)$ is the surprisal of hypothesis h under observation y . Surprisal and the impossibility increment are the same object: an identity, not an analogy.

Second, TEAG establishes the *epistemic geoid* as the geometric interpretation of the Possibilistic Cramér–Rao Bound. The PCRB is an equipotential surface of the impossibility field — the boundary of admissible contraction per observation, analogous to a geoid in gravitational theory. The whitened minimax medoid $A(\pi)$ is the center of this geoid: the commitment point most geometrically insulated from its boundary, not the mode of the possibility field. This distinction — between geometric centrality and possibilistic extremality — is the correct characterization of the commitment rule, and is confirmed empirically by the ESPF’s lower estimation error relative to both the posterior mean and the Chebyshev center.

Third, TEAG establishes that the possibilistic entropy $\mathcal{E}_\pi = \int_0^1 \log V_\alpha \, d\alpha$ is Boltzmann entropy, not Shannon entropy. Its logarithm is justified by geometric independence: prior impossibility and surprisal are independent sources whose joint admissible volume is the product of marginal volumes, and the logarithm converts that product to a sum. Shannon entropy cannot play this role because possibility theory uses the min rule for conjunction, under which $-\log$ is not additive across independent possibilistic sources. This is a precise statement about which entropy concept is appropriate for which uncertainty structure.

Fourth, TEAG proves that probability theory is the *collapse limit* of possibilistic inference, not its foundation. As epistemic width $W = \int (\pi - n) \, d\mu$ approaches zero, the Choquet integral converges to the Lebesgue integral, the ESPF converges to the Kalman filter, and the PLM converges to a standard

transformer. In every domain, this convergence is convergent optimality: two frameworks optimal under different epistemic commitments that agree precisely when their domains of applicability coincide. Probability is earned by evidence, not assumed.

Fifth, TEAG gives Popper's falsification criterion a *two-stage algebraic formulation*. The tropical variety $\mathcal{B}_{\text{active}} = \{h : \Phi^-(h) = \psi(h)\}$ is the active deformation front: the exact locus where evidence begins to deform the posterior impossibility field, and a necessary condition for falsification. Sufficient falsification requires exit from the PCRB-admissible basin $\mathcal{A}_k = \{h : \Phi_k^+(h) \leq c_k^*\}$, whose threshold c_k^* is determined by the PCRB and the MVEE geometry. This resolves a structural pathology in the pointwise crossing definition: a hypothesis with very low prior impossibility — earned through many observations — is protected from mild surprisal by the basin, not exposed to it. The minimax medoid is the fixed point of this basin: the unique support point that survives as long as any hypothesis survives.

What TEAG unifies.

By naming the shared primitive, TEAG enables three classes of results that were previously inaccessible.

The PCRB — proved in the ESPF Optimality paper for state estimation — now holds universally across all TEAG instantiations: it constrains entropy reduction per measurement step in state estimation, per observation in measure-theoretic collapse, per generation step in language modeling, and per contraction step in compositional systems. It is a universal geometric bound on the rate of impossibility deformation under evidence, not a filter-specific result.

The Boltzmann justification for $\mathcal{E}_\pi$ carries across all four domains because the independence structure is the same in each: prior impossibility and surprisal are geometrically independent wherever the TEAG axioms hold. The entropy is additive across independent domains of state space by the geometry of volumes, not by the axioms of probability.

Compositional TEAG — the coupling of multiple TEAG objects through a shared epistemic width monitor — provides a principled architecture for systems that must simultaneously maintain epistemic integrity across multiple inference domains. The convergent optimality structure and the PCRB transfer automatically to the compositional setting, and scale equalization — smaller models reaching the same epistemic fixed points as larger ones — emerges as a consequence of the contraction structure rather than parameter volume.

What TEAG does not claim.

TEAG is a framework for inference under bounded epistemic uncertainty in information-fusion contexts. It makes no claim about the ontological status of probability in quantum mechanics, statistical mechanics, or other domains where probability has well-established interpretations orthogonal to the present work. The claim is operational: in contexts where epistemic and aleatory uncertainty must be distinguished, where evidence eliminates rather than merely reweights, and where the commitment must be epistemically defensible rather than merely optimal in expectation, the TEAG object is the appropriate primitive.

Open directions.

Four directions remain open and are taken up in companion and future work. The formal convergence proof for the epistemic manifold — smooth parameterization of $\mathcal{M}_{\text{ep}}$, a connection, and a Ricci tensor — requires proving that $\{P_\alpha\}$ varies smoothly under the TEAG recursion. The exact calibration of c_k^* in the fully anisotropic non-quadratic regime beyond the MVEE scaling law approximation requires deriving closed-form correction terms from the contact point geometry. The ESPF vs. UKF head-to-head comparison under the convergent optimality framing remains as the primary empirical obligation. And the formal PAC-possibilistic learnability theory for the PLM training objective has not yet been established.

The epistemic posture of TEAG.

TEAG is built on a single asymmetry: evidence can only increase impossibility, never decrease it. A hypothesis that survives is not confirmed — it is merely not yet falsified. The commitment $A(\pi)$ is not the truth — it is the least inadmissible hypothesis, the center of the surviving basin, the point most sheltered from wherever the next contraction may come. This posture — quick to embrace ignorance, slow to assert certainty — is not a design philosophy. In the falsification regime, it is a theorem.

TEAG as the physics of belief.

The TEAG update rule is a max-plus dynamical system. Its governing equations have the same algebraic form as the Hamilton–Jacobi equations of classical mechanics. The PCRB is a minimum action principle. Free inference follows geodesics of the MVEE metric. The impossibility gradient deflects inference from those geodesics with explicit Christoffel symbols.

This is a statement about algebraic structure. The TEAG governing equations and the classical Hamilton–Jacobi governing equations are the same equations, applied to different spaces — hypothesis space rather than physical space. Inference is not a physical process. What is claimed is the algebraic identity of the governing structure. That identity is what makes TEAG the physics of belief: not a metaphor, but a precise statement about which mathematical object governs the dynamics of belief contraction under evidence.

TEAG as epistemically neutral synthesis.

Because TEAG’s axioms require only a hypothesis space, a possibility field, and a contraction operator — not a probability measure, a likelihood function, or a frequentist grounding — the framework is structurally neutral between knowledge systems. Western scientific inference and Indigenous localized knowledge can each instantiate the TEAG quintuple independently, without transformation of one into the other’s representational system. Their joint admissible support intersection is the locus of coherence. The composition theory provides the coupling architecture.

This is, to the best of our knowledge, the first inference framework that enables coherent synthesis of heterogeneous knowledge systems without requiring lossless transformation between them. It is a consequence of the architecture, not an add-on.

The final word belongs to the structure itself. Possibility theory, tropical algebra, Hamiltonian mechanics, and the geometry of convex bodies are not four separate contributions assembled for convenience. They are four views of the same object: the epistemically minimal description of what it means to reason under bounded ignorance. TEAG names that object.

References

1. S. Amari, *Information Geometry and Its Applications*. Springer, 2016.
2. M. Akian, R. Bapat, and S. Gaubert, “Max-plus algebra,” in *Handbook of Linear Algebra*, L. Hogben, Ed. Chapman & Hall/CRC, 2007, ch. 25.
3. T. Augustin, F. P. A. Coolen, G. de Cooman, and M. C. M. Troffaes, Eds., *Introduction to Imprecise Probabilities*. Wiley, 2014.
4. F. John, “Extremum problems with inequalities as subsidiary conditions,” in *Studies and Essays Presented to R. Courant on his 60th Birthday*, K. O. Friedrichs, O. E. Neugebauer, and J. J. Stoker, Eds. Interscience, New York, pp. 187–204, 1948.
5. A. N. Angelopoulos and S. Bates, “A gentle introduction to conformal prediction and distribution-free uncertainty quantification,” *arXiv:2107.07511*, 2022.
6. G. Choquet, “Theory of capacities,” *Annales de l’Institut Fourier*, vol. 5, pp. 131–295, 1953–1954.
7. H. Cramér, *Mathematical Methods of Statistics*. Princeton University Press, 1946.
8. M. Grabisch, *Set Functions, Games and Capacities in Decision Making*. Springer, 2016.
9. M. J. Todd, *Minimum-Volume Ellipsoids: Theory and Algorithms*. SIAM, Philadelphia, 2016.
10. E. Delande, J. Houssineau, and D. Clark, “Propagation of epistemic uncertainty in dynamical systems,” *IEEE Trans. Aerospace and Electronic Systems*, vol. 56, no. 6, pp. 4754–4768, 2020.

11. D. Dubois and H. Prade, *Possibility Theory: An Approach to Computerized Processing of Uncertainty*. Plenum Press, New York, 1988.
12. D. Dubois and H. Prade, "Possibility theory in constraint satisfaction problems: handling priority, preference and uncertainty," *Applied Intelligence*, vol. 6, no. 4, pp. 287–309, 1996.
13. R. Fagin, J. Y. Halpern, Y. Moses, and M. Y. Vardi, *Reasoning About Knowledge*. MIT Press, 1995.
14. V. N. Gladyshev, "The information theory of aging," *Aging and Disease*, vol. 12, no. 2, pp. 579–586, 2021.
15. C. W. Helstrom, *Quantum Detection and Estimation Theory*. Academic Press, 1976.
16. S. Horvath, "DNA methylation age of human tissues and cell types," *Genome Biology*, vol. 14, no. 10, p. R115, 2013.
17. N. J. Gordon, D. J. Salmond, and A. F. M. Smith, "Novel approach to nonlinear/non-Gaussian Bayesian state estimation," *IEE Proceedings F*, vol. 140, no. 2, pp. 107–113, 1993.
18. C. Guo, G. Pleiss, Y. Sun, and K. Q. Weinberger, "On calibration of modern neural networks," *Proc. ICML*, pp. 1321–1330, 2017.
19. M. Higashi and G. J. Klir, "Measures of uncertainty and information based on possibility distributions," *Int. J. General Systems*, vol. 9, no. 1, pp. 43–58, 1983.
20. J. Houssineau and D. Clark, "A subjective approach to random sets," *Proc. Royal Society A*, vol. 472, no. 2190, 2016.
21. R. E. Kalman, "A new approach to linear filtering and prediction problems," *J. Basic Engineering*, vol. 82, no. 1, pp. 35–45, 1960.
22. D. Maclagan and B. Sturmfels, *Introduction to Tropical Geometry*. American Mathematical Society, 2015.
23. W. J. Maddox, P. Izmailov, T. Garipov, D. P. Vetrov, and A. G. Wilson, "A simple baseline for Bayesian uncertainty in deep learning," *Advances in Neural Information Processing Systems*, vol. 32, 2019.
24. C. R. Rao, "Information and the accuracy attainable in the estimation of statistical parameters," *Bull. Calcutta Math. Soc.*, vol. 37, pp. 81–91, 1945.
25. M. Senséoy, L. Kaplan, and M. Kandemir, "Evidential deep learning to quantify classification uncertainty," *Advances in Neural Information Processing Systems*, vol. 31, 2018.
26. G. Shafer, *A Mathematical Theory of Evidence*. Princeton University Press, 1976.
27. C. Villani, *Optimal Transport: Old and New*. Springer, 2009.
28. M. K. Jah and V. Haslett, "The Epistemic Support-Point Filter (ESPF): A bounded possibilistic framework for ordinal state estimation," *arXiv:2508.20806*, 2025.
29. M. K. Jah, "The Geometry of Knowing: From possibilistic ignorance to probabilistic certainty," preprint, *arXiv:submit/7362363*, 2026 (under review).
30. M. K. Jah, "The Epistemic Support-Point Filter: Jaynesian maximum entropy meets Popperian falsification," *arXiv:2603.10065v2*, 2026.
31. M. K. Jah, "Epistemic Wavefronts: A Tropical Hamilton–Jacobi Theory of Possibilistic Inference," preprint, GaiaVerse, Ltd., 2026 (submitted).
32. M. K. Jah, "Toward a Possibilistic Language Model," manuscript in preparation, GaiaVerse, Ltd., 2026 (forthcoming).
33. Z. Liu and M. K. Jah, *Semantic Turning Point Detector* (v4.0.0). GaiaVerse Ltd., Zenodo, 2025. <https://doi.org/10.5281/zenodo.16734909>
34. R. E. Kalman, "System identification from noisy data," in *Dynamical Systems II*, A. Bednarek, Ed., vol. 210, pp. 135–164. Academic Press, New York, 1982. Presented at the International Symposium on Dynamical Systems, Gainesville, FL, February 1981.
35. P. Walley, *Statistical Reasoning with Imprecise Probabilities*. Chapman and Hall, 1991.
36. L. A. Zadeh, "Fuzzy sets as a basis for a theory of possibility," *Fuzzy Sets and Systems*, vol. 1, no. 1, pp. 3–28, 1978.

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