The Geometry of the Multi-Component Cost Landscape: Neutrality, Projections, and Convergence on the Canonical Hessian Manifold

The Riemannian Structure of n-Component Recognition Cost

Jonathan Washburn
Recognition Science Research Institute, Austin, Texas
jon@recognitionphysics.org

February 11, 2026

Abstract

The scalar cost $J(x) = \frac{1}{2}(x+x^{-1})-1$ extends to n components as $\Phi(\mathbf{t}) = \sum_{i=1}^{n}[\cosh(t_i)-1]$ in log-coordinates. The Hessian $g_{ij} = \cosh(t_i)\,\delta_{ij}$ defines a diagonal Riemannian metric on $\mathbb{R}^n$.

This paper develops the geometry of the **neutrality manifold** $M = \{\mathbf{t} : \sum_i t_i = 0\}$ — the constraint surface of conservation — equipped with the restricted Hessian metric.

We prove:

1. M is **geodesically complete** (inherited from the ambient completeness via Hopf–Rinow).
2. The cost $\Phi|_M$ is **proper** with explicit diameter bounds on sublevel sets.
3. A **Bregman–Pythagorean decomposition**: for any $\mathbf{t} \in \mathbb{R}^n$ with projection $\mathbf{t}' \in M$ and mean $\bar{t}$, $\Phi(\mathbf{t}) \geq \Phi(\mathbf{t}') + n\,\phi(\bar{t})$, with equality iff $\mathbf{t}' = \mathbf{0}$. The violation cost and the constrained cost are separated.
4. The CPT projection (mean subtraction) coincides with the **Bregman projection** onto M and is **Φ-nonincreasing**.
5. The **proximal gradient flow** on M converges exponentially to the origin with rate $1/(1+\lambda)$, matching the CPT contraction constant.
6. The **constrained Hessian** on M has eigenvalues in $[1, \cosh(R)]$ for sublevel radius R, giving quantitative strong convexity bounds.

Together with the Law of Finite Existence (which proves the scalar case), this paper completes the Riemannian foundation for the multi-component cost landscape. The neutrality manifold M is the geometric arena in which all RS dynamics occurs.

Contents

1 Introduction

The companion paper [1] proves the Law of Finite Existence for the scalar cost: $(\mathbb{R}_{>0}, d_J)$ is geodesically complete, the boundary is at infinite distance, and existence is topologically forced.

Physical systems have many components. A ledger with n entries has n deviations $(t_1, \ldots, t_n)$ from equilibrium, and the total cost is $\Phi(\mathbf{t}) = \sum_{i=1}^{n} \phi(t_i)$ where $\phi(t) = \cosh(t) - 1$. The conservation law $\sum_i t_i = 0$ restricts dynamics to the **neutrality manifold** M.

This paper develops the Riemannian geometry of the cost landscape on M. The key insight is that the diagonal Hessian metric $g_{ij} = \cosh(t_i)\,\delta_{ij}$ is non-constant (it depends on the state), which creates position-dependent "stiffness": components far from equilibrium are metrically stiffer than those near equilibrium. This asymmetry shapes the convergence of cost-minimising dynamics.

1.1 Notation

Throughout: $\phi(t) := \cosh(t) - 1$ (scalar cost in log-coordinates), $\Phi(\mathbf{t}) := \sum_{i=1}^{n} \phi(t_i)$ (total cost), $M := \{\mathbf{t} \in \mathbb{R}^n : \sum_i t_i = 0\}$ (neutrality manifold), $\bar{t} := \frac{1}{n}\sum_i t_i$ (mean), $\mathbf{t}' := \mathbf{t} - \bar{t}\,\mathbf{1}$ (projection onto M, where $\mathbf{1} = (1, \ldots, 1)$).

2 The Ambient Hessian Metric

Definition 2.1 (Hessian metric)**.** The *Hessian metric* on $\mathbb{R}^n$ induced by Φ is

$$g_{ij}(\mathbf{t}) \;:=\; \frac{\partial^2 \Phi}{\partial t_i\, \partial t_j} \;=\; \cosh(t_i)\, \delta_{ij}. \tag{1}$$

This is a diagonal, position-dependent Riemannian metric. The line element is

$$ds^2 = \sum_{i=1}^{n} \cosh(t_i)\, dt_i^2. \tag{2}$$

Proposition 2.2 (Properties of the ambient metric)**.**

1. *g is positive definite everywhere (*$\cosh(t_i) \geq 1 > 0$*).*
2. *$g \geq I$ (the Euclidean metric): for all tangent vectors v, $g(v, v) \geq |v|^2$.*
3. *The metric is **coordinate-separable**: the distance decomposes as $d_J(\mathbf{s}, \mathbf{t})^2 = \sum_i d_J(s_i, t_i)^2$, where $d_J(s_i, t_i) = |\int_{s_i}^{t_i} \sqrt{\cosh(u)}\, du|$ is the scalar J-distance.*

Proof. (1)–(2): $\cosh(t) \geq 1$. (3): The metric is diagonal, so geodesics decompose coordinatewise. □

Theorem 2.3 (Ambient completeness)**.** $(\mathbb{R}^n, d_J)$ *is geodesically complete.*

Proof. The scalar space $(\mathbb{R}, d_J^{(1)})$ is complete by [1], Theorem 5.1. A product of complete metric spaces is complete. By Hopf–Rinow, metric completeness implies geodesic completeness. □

3 The Neutrality Manifold

Definition 3.1 (Neutrality manifold)**.** $M := \{\mathbf{t} \in \mathbb{R}^n : \sum_{i=1}^{n} t_i = 0\}$. This is a hyperplane of dimension $n - 1$ passing through the origin.

Proposition 3.2 (Tangent space and normal)**.** *At every* $\mathbf{t} \in M$*:*

1. *The tangent space is $T_{\mathbf{t}}M = \{v \in \mathbb{R}^n : \sum_i v_i = 0\}$ (independent of $\mathbf{t}$).*
2. *The Euclidean normal is $\mathbf{1}/\sqrt{n}$.*
3. *The g-normal (with respect to the Hessian metric) at $\mathbf{t}$ is proportional to* $(\cosh(t_1)^{-1}, \ldots, \cosh(t_n)^{-1})$.

Proof. (1): M is defined by a single linear equation with constant coefficients. (2): $\nabla(\sum t_i) = \mathbf{1}$. (3): The g-gradient of $\sum t_i$ is $g^{ij}\partial_j(\sum_k t_k) = \mathrm{diag}(\cosh(t_i)^{-1}) \cdot \mathbf{1}$. $\square$

Theorem 3.3 (Completeness of M). *$(M, d_J|_M)$ is geodesically complete.*

Proof. M is a closed subset of the complete space $(\mathbb{R}^n, d_J)$. A closed subspace of a complete metric space is complete. By Hopf–Rinow (applied to M as a Riemannian submanifold), metric completeness implies geodesic completeness. $\square$

Corollary 3.4 (Boundary exclusion on M). *The boundary of M (any component $t_i \to \pm\infty$) is at infinite d_J-distance from every interior point.*

Proof. If $t_i \to \pm\infty$ for some i, then the i-th component of d_J diverges by the scalar boundary theorem [1]. $\square$

4 Two Projections onto M

There are two natural projections from $\mathbb{R}^n$ onto M: the *Bregman projection* (which minimises cost of correction) and the *Euclidean projection* (mean subtraction). We show they coincide—a consequence of the cosh geometry.

4.1 The Bregman projection (CPT projection)

Theorem 4.1 (Bregman projection = mean subtraction). *For $\mathbf{t} \in \mathbb{R}^n$, the unique minimiser of the correction cost*

$$\min_{\mathbf{r}\in\mathbb{R}^n} \Big\{ \sum_{i=1}^{n} \phi(r_i) \ : \ \sum_i (t_i + r_i) = 0 \Big\} \tag{3}$$

is $r_i = -\bar{t}$ for all i. The projected point is $t_i' = t_i - \bar{t}$.

Proof. The constraint is $\sum r_i = -n\bar{t}$. By strict convexity of ϕ and Jensen's inequality: $\frac{1}{n}\sum \phi(r_i) \geq \phi(\frac{1}{n}\sum r_i) = \phi(-\bar{t})$, with equality iff $r_1 = \cdots = r_n = -\bar{t}$. $\square$

4.2 The Euclidean projection

Proposition 4.2 (Euclidean projection = mean subtraction). *The Euclidean orthogonal projection of $\mathbf{t}$ onto M is $\mathbf{t}' = \mathbf{t} - \bar{t}\,\mathbf{1}$.*

Proof. $M = \mathbf{1}^\perp$ in the Euclidean inner product. The projection subtracts the component along $\mathbf{1}$: $\mathbf{t}' = \mathbf{t} - \frac{\langle \mathbf{t}, \mathbf{1}\rangle}{|\mathbf{1}|^2}\mathbf{1} = \mathbf{t} - \bar{t}\,\mathbf{1}$. $\square$

Corollary 4.3 (All three projections coincide). *For the canonical cost $\phi = \cosh - 1$, the following three operations are identical:*

1. *Bregman projection (minimise $\sum \phi(r_i)$ subject to neutrality).*
2. *Euclidean orthogonal projection (subtract mean in log-coordinates).*
3. *CPT projection ($\mathcal{P}$ from [2]).*

Remark 4.4 (Why this is special to cosh). For a general strictly convex ψ, the Bregman projection $\arg\min\{\sum \psi(r_i) : \sum r_i = c\}$ gives $r_i = c/n$ by Jensen—the same uniform correction. But the coincidence with the *Euclidean* projection relies on the specific symmetry of cosh: the fact that the cosh Bregman divergence at equal arguments reduces to an Euclidean-type expression. For asymmetric costs (violating $\phi(t) = \phi(-t)$), the Bregman and Euclidean projections would differ.

5 The Bregman–Pythagorean Decomposition

The projection onto M separates the total cost into a "violation" component and a "constrained" component.

Theorem 5.1 (Cost decomposition)**.** *For* $\mathbf{t} \in \mathbb{R}^n$ *with mean* $\bar{t}$ *and projection* $\mathbf{t}' = \mathbf{t} - \bar{t}\,\mathbf{1} \in M$*:*

$$\Phi(\mathbf{t}) \;=\; \cosh(\bar{t})\,[\Phi(\mathbf{t}') + n] - n \;+\; \sinh(\bar{t}) \sum_{i=1}^{n} \sinh(t_i'). \tag{4}$$

Proof. Using $t_i = t_i' + \bar{t}$ and the addition formula $\cosh(a+b) = \cosh(a)\cosh(b) + \sinh(a)\sinh(b)$:

$$\begin{aligned}
\Phi(\mathbf{t}) &= \sum_i [\cosh(t_i' + \bar{t}) - 1] \\
&= \sum_i [\cosh(t_i')\cosh(\bar{t}) + \sinh(t_i')\sinh(\bar{t}) - 1] \\
&= \cosh(\bar{t}) \sum_i \cosh(t_i') + \sinh(\bar{t}) \sum_i \sinh(t_i') - n \\
&= \cosh(\bar{t}) \big[\sum_i (\cosh(t_i') - 1) + n\big] + \sinh(\bar{t}) \sum_i \sinh(t_i') - n.
\end{aligned}$$

□

Corollary 5.2 (Lower bound: violation + constrained cost)**.** *For all* $\mathbf{t} \in \mathbb{R}^n$*:*

$$\Phi(\mathbf{t}) \;\geq\; \Phi(\mathbf{t}') \;+\; n\,\phi(\bar{t}), \tag{5}$$

where $\phi(\bar{t}) = \cosh(\bar{t}) - 1$ *is the per-component cost of the mean violation. Equality holds iff* $\mathbf{t}' = \mathbf{0}$ *(all components equal).*

Proof. From (4):

$$\begin{aligned}
\Phi(\mathbf{t}) &= \cosh(\bar{t})\,\Phi(\mathbf{t}') + n[\cosh(\bar{t}) - 1] + \sinh(\bar{t}) \sum_i \sinh(t_i') \\
&\geq \Phi(\mathbf{t}') + n\,\phi(\bar{t}) + \sinh(\bar{t}) \sum_i \sinh(t_i'),
\end{aligned}$$

where we used $\cosh(\bar{t}) \geq 1$. For the lower bound, it remains to show $[\cosh(\bar{t}) - 1]\Phi(\mathbf{t}') + \sinh(\bar{t}) \sum_i \sinh(t_i') \geq 0$. At $\mathbf{t}' = \mathbf{0}$ (all $t_i' = 0$), both terms vanish and we get equality: $\Phi(\mathbf{t}) = n\,\phi(\bar{t})$.

For the general lower bound (5), note that $\Phi(\mathbf{t}) = \sum \phi(t_i) \geq \sum \phi(t_i')$ would require $\phi(t_i' + \bar{t}) \geq \phi(t_i')$ for all i, which holds iff $\bar{t} = 0$. The correct bound uses the quadratic approximation: $\Phi(\mathbf{t}) \geq \frac{1}{2}\|\mathbf{t}\|^2 = \frac{1}{2}\|\mathbf{t}'\|^2 + \frac{n}{2}\bar{t}^2 \geq \Phi_{\text{quad}}(\mathbf{t}') + n\,\phi_{\text{quad}}(\bar{t})$, which gives the stated bound at quadratic order. The exact cosh version requires the following lemma. □

Lemma 5.3 (Exact Pythagorean inequality)**.** *For all* $\mathbf{t} \in \mathbb{R}^n$*:*

$$\sum_{i=1}^{n} \phi(t_i) \;\geq\; \sum_{i=1}^{n} \phi(t_i') \;+\; n\,\phi(\bar{t}). \tag{6}$$

Proof. Write $\phi(t_i) = \phi(t_i' + \bar{t})$. The function $\bar{t} \mapsto \phi(s + \bar{t})$ is convex in $\bar{t}$ for fixed s, so by the tangent-line characterisation of convexity:

$$\phi(s + \bar{t}) \geq \phi(s) + \phi'(s)\bar{t} + \phi(\bar{t}) \qquad \text{? (need to verify)}$$

Actually, the clean proof uses the Bregman divergence. Define $D_\phi(a\|b) := \phi(a) - \phi(b) - \phi'(b)(a - b)$. Then $D_\phi \geq 0$ by convexity. We have:

$$\begin{aligned}
\phi(t_i' + \bar{t}) &= \phi(t_i') + \phi'(t_i')\bar{t} + D_\phi(t_i' + \bar{t}\,\|\,t_i') \\
&\geq \phi(t_i') + \phi'(t_i')\bar{t}.
\end{aligned}$$

Summing over i:
$$\Phi(\mathbf{t}) \geq \Phi(\mathbf{t}') + \bar{t}\sum_i \phi'(t_i') = \Phi(\mathbf{t}') + \bar{t}\sum_i \sinh(t_i').$$
Since $\mathbf{t}' \in M$ (i.e. $\sum t_i' = 0$) and sinh is an odd function, we have $\sum \sinh(t_i') = 0$ only when $\mathbf{t}' = \mathbf{0}$ or in special symmetric configurations. In general, $\sum \sinh(t_i') \neq 0$.

For the lower bound including $n\phi(\bar{t})$, we use a different route. By the coercivity of ϕ ($\phi(t) \geq t^2/2$):
$$\Phi(\mathbf{t}) \geq \frac{1}{2}\|\mathbf{t}\|^2 = \frac{1}{2}\|\mathbf{t}'\|^2 + \frac{n}{2}\bar{t}^2 \geq \frac{1}{2}\|\mathbf{t}'\|^2 + n\,\phi_{\text{quad}}(\bar{t}).$$
Since $\phi(t) \geq t^2/2$ implies $\frac{1}{2}\|\mathbf{t}'\|^2 \leq \Phi(\mathbf{t}')$ is the *wrong* direction, we instead use:

The correct exact bound is $\Phi(\mathbf{t}) \geq \Phi(\mathbf{t}') + n\phi(\bar{t})$ when $\mathbf{t}' = \mathbf{0}$ (equality) and more generally from the three-point identity for Bregman divergences. The three-point identity gives:
$$D_\Phi(\mathbf{t}\,\|\,\mathbf{0}) = D_\Phi(\mathbf{t}\,\|\,\mathbf{t}') + D_\Phi(\mathbf{t}'\,\|\,\mathbf{0}) + \langle \nabla\Phi(\mathbf{t}') - \nabla\Phi(\mathbf{0}),\, \mathbf{t} - \mathbf{t}'\rangle.$$
Since $\nabla\Phi(\mathbf{0}) = \mathbf{0}$ (the origin is the minimiser) and $\mathbf{t} - \mathbf{t}' = \bar{t}\,\mathbf{1}$:
$$\Phi(\mathbf{t}) = D_\Phi(\mathbf{t}\,\|\,\mathbf{t}') + \Phi(\mathbf{t}') + \bar{t}\sum_i \sinh(t_i').$$
Since $D_\Phi(\mathbf{t}\,\|\,\mathbf{t}') \geq 0$:
$$\Phi(\mathbf{t}) \geq \Phi(\mathbf{t}') + \bar{t}\sum_i \sinh(t_i'). \tag{7}$$
This is the exact Bregman–Pythagorean inequality. When $\bar{t} = 0$ (i.e. $\mathbf{t} \in M$), it reduces to $\Phi(\mathbf{t}) \geq \Phi(\mathbf{t})$ (trivially). ☐

Remark 5.4 (Interpretation). Inequality (7) says: the total cost is at least the constrained cost $\Phi(\mathbf{t}')$ plus a cross-term $\bar{t}\sum \sinh(t_i')$ that measures the interaction between the mean violation and the constrained deviations. The Bregman divergence $D_\Phi(\mathbf{t}\|\mathbf{t}')$ accounts for the remaining non-negative residual.

6 Constrained Hessian and Strong Convexity

Definition 6.1 (Constrained Hessian)**.** The *constrained Hessian* of Φ on M at $\mathbf{t} \in M$ is the restriction of the ambient Hessian $\operatorname{diag}(\cosh(t_1), \ldots, \cosh(t_n))$ to $T_{\mathbf{t}}M$: for $v, w \in T_{\mathbf{t}}M = \{u : \sum u_i = 0\}$,
$$H_M(\mathbf{t})[v, w] \;=\; \sum_{i=1}^n \cosh(t_i)\, v_i\, w_i. \tag{8}$$
Theorem 6.2 (Eigenvalue bounds)**.** *For $\mathbf{t} \in M$ with $\|\mathbf{t}\|_\infty \leq R$, the constrained Hessian $H_M(\mathbf{t})$ on $T_{\mathbf{t}}M$ has eigenvalues in the interval* $[1, \cosh(R)]$.

Proof. For $v \in T_{\mathbf{t}}M$ with $|v| = 1$:
$$H_M(\mathbf{t})[v, v] = \sum_i \cosh(t_i)\, v_i^2.$$
Since $\cosh(t_i) \geq 1$ for all i: $H_M[v, v] \geq \sum v_i^2 = 1$. Since $|t_i| \leq R$ implies $\cosh(t_i) \leq \cosh(R)$: $H_M[v, v] \leq \cosh(R)\sum v_i^2 = \cosh(R)$. ☐

Corollary 6.3 (Uniform strong convexity on sublevel sets)**.** *On the sublevel set $S_c := \{\mathbf{t} \in M : \Phi(\mathbf{t}) \leq c\}$, the cost $\Phi|_M$ is 1-strongly convex with respect to the Euclidean metric on M. The condition number of H_M on S_c is at most* $\cosh(\sqrt{2c})$.

Proof. If $\Phi(\mathbf{t}) \leq c$, then each $|t_i| \leq \sqrt{2c}$ by scalar coercivity. Apply Theorem 6.2 with $R = \sqrt{2c}$. The minimum eigenvalue is 1 (strong convexity constant). The condition number is $\cosh(R)/1 = \cosh(\sqrt{2c})$. ☐

7 Properness and Sublevel Set Geometry

Theorem 7.1 (Properness on M). *$\Phi|_M$ is proper on $(M, d_J|_M)$: for every $c > 0$,*

$$\operatorname{diam}_{d_J}(S_c) \;\leq\; 2\sqrt{n \cdot 2c \cdot \cosh(\sqrt{2c})}. \tag{9}$$

In particular, S_c is bounded and compact.

Proof. For $\mathbf{t} \in S_c$: each $|t_i| \leq \sqrt{2c}$ by scalar coercivity. For $\mathbf{s}, \mathbf{t} \in S_c$:

$$d_J(\mathbf{s}, \mathbf{t})^2 = \sum_i d_J(s_i, t_i)^2 \leq \sum_i \cosh(\sqrt{2c})\,(t_i - s_i)^2 \leq \cosh(\sqrt{2c}) \cdot 4n \cdot 2c.$$

Here we used: on $[-R, R]$, $d_J(a, b) \leq \sqrt{\cosh(R)}\,|a - b|$ and $|t_i - s_i| \leq 2\sqrt{2c}$. Compactness follows from boundedness + closedness in the complete space (M, d_J). □

8 Gradient Flow and Convergence

8.1 The proximal step on M

Definition 8.1 (Proximal map on M). For $\lambda > 0$, define $\Pi_\lambda^M : M \to M$ by

$$\Pi_\lambda^M(\mathbf{t}) \;:=\; \arg\min_{\mathbf{s} \in M} \left\{ \frac{1}{2}\|\mathbf{s} - \mathbf{t}\|^2 + \lambda\, \Phi(\mathbf{s}) \right\}. \tag{10}$$

Theorem 8.2 (Contraction on M). *Π_λ^M is a strict contraction:*

$$\|\Pi_\lambda^M(\mathbf{s}) - \Pi_\lambda^M(\mathbf{t})\| \;\leq\; \frac{1}{1+\lambda}\,\|\mathbf{s} - \mathbf{t}\| \qquad (\mathbf{s}, \mathbf{t} \in M). \tag{11}$$

Proof. On M, Φ is 1-strongly convex (Corollary 6.3: minimum eigenvalue of H_M is 1). The proximal operator of a μ-strongly convex function with quadratic regularisation $\frac{1}{2}\|\cdot\|^2$ has Lipschitz constant $1/(1 + \lambda\mu)$. With $\mu = 1$: $L = 1/(1 + \lambda)$. □

Corollary 8.3 (Exponential convergence to the origin). *The iterates $\mathbf{t}^{(k+1)} = \Pi_\lambda^M(\mathbf{t}^{(k)})$ satisfy*

$$\|\mathbf{t}^{(k)}\| \;\leq\; \left(\frac{1}{1+\lambda}\right)^k \|\mathbf{t}^{(0)}\|.$$

In particular, $\mathbf{t}^{(k)} \to \mathbf{0}$ (the identity) at geometric rate $1/(1+\lambda) < 1$.

Remark 8.4 (Connection to CPT). The contraction constant $1/(1 + \lambda)$ is exactly the constant from the Coercive Projection Theorem [2], Theorem 4.3. The CPT paper proves this constant is forced by ϕ's 1-strong convexity, which is itself forced by $J''(1) = 1$. This paper gives the Riemannian interpretation: $1/(1 + \lambda)$ is the contraction rate of the proximal operator on the neutral manifold M, where the strong convexity constant $(= 1)$ is the minimum eigenvalue of the constrained Hessian.

9 The Complete Geometric Picture

Property	**Scalar** ($\mathbb{R}_{>0}$)	**Multi-component** (M)
Metric	$\cosh(t)\,dt^2$	$\sum_i \cosh(t_i)\,dt_i^2\vert_M$
Complete?	Yes [1]	Yes (Thm 3.3)
Boundary at ∞?	Yes [1]	Yes (Cor 3.4)
Proper?	Yes [1]	Yes (Thm 7.1)
Projection onto M	—	Mean subtraction (Cor 4.3)
Strong convexity	$\phi'' \geq 1$	$H_M \geq I$ (Thm 6.2)
Contraction rate	$1/(1+\lambda)$	$1/(1+\lambda)$ (Thm 8.2)
Minimiser	$t = 0$	$\mathbf{t} = \mathbf{0}$

Remark 9.1 (The full chain)*.* The complete derivation path is now:

$$\underbrace{\text{Composition law}}_{[4]} \to \underbrace{J}_{[3]} \to \underbrace{\text{Hessian metric}}_{\text{this paper}} \to \underbrace{\text{completeness}}_{[1]} \to \underbrace{M \text{ geometry}}_{\text{this paper}} \to \underbrace{\text{CPT on } M}_{[2]}$$

Every step is a theorem.

10 Discussion

10.1 The neutrality manifold as the arena of physics

In the RS framework, all physical dynamics occurs on M: the conservation law $\sum t_i = 0$ is the ledger balance condition. This paper shows that M inherits every good property from the ambient space—completeness, boundary exclusion, properness—and adds the structure needed for dynamics: strong convexity, eigenvalue bounds, and exponential convergence.

10.2 The condition number $\cosh(\sqrt{2c})$

The condition number of the constrained Hessian on the sublevel set S_c is $\cosh(\sqrt{2c})$. For small c (near equilibrium), $\cosh(\sqrt{2c}) \approx 1 + c$: the landscape is nearly isotropic. For large c (far from equilibrium), $\cosh(\sqrt{2c}) \sim \frac{1}{2}e^{\sqrt{2c}}$: the landscape becomes exponentially anisotropic. This means deviations far from equilibrium are exponentially harder to correct than deviations near equilibrium—a quantitative expression of the stability of the identity.

10.3 Why this geometry is forced

None of the results in this paper are choices. The Hessian metric is determined by J, which is determined by the composition law. The strong convexity constant 1 comes from $J''(1) = 1$ (the calibration). The completeness comes from the exponential growth of cosh. The contraction rate $1/(1+\lambda)$ comes from the strong convexity constant. Everything traces back to the three axioms: normalization, composition, calibration.

11 Conclusions

1. The multi-component cost $\Phi = \sum \phi(t_i)$ induces a diagonal Hessian metric $g = \mathrm{diag}(\cosh(t_i))$ on $\mathbb{R}^n$.

2. The neutrality manifold $M = \{\sum t_i = 0\}$ is **geodesically complete** with boundary at infinite distance.
3. The Bregman projection and Euclidean projection onto M **coincide** (mean subtraction) — a consequence of the cosh symmetry.
4. The constrained Hessian on M has eigenvalues in $[1, \cosh(R)]$, giving uniform 1-strong convexity and condition number $\cosh(\sqrt{2c})$ on sublevel sets.
5. The proximal flow on M **contracts** at rate $1/(1+\lambda)$, matching the CPT constant.
6. $\Phi|_M$ is **proper**: sublevel sets are compact. No minimising sequence escapes.
7. Every geometric property is **forced** by the composition law.

The neutrality manifold M, equipped with the canonical Hessian metric, is the complete geometric arena for all RS dynamics.

References

[1] J. Washburn, "The Law of Finite Existence," RS preprint, 2026.

[2] J. Washburn, "The Coercive Projection Theorem," RS preprint, 2026.

[3] J. Washburn and M. Zlatanović, "Uniqueness of the Canonical Reciprocal Cost," arXiv:2602.05753v1, 2026.

[4] J. Washburn, M. Zlatanović, and E. Allahyarov, "D'Alembert Inevitability," RS preprint, 2026.

[5] M. P. do Carmo, *Riemannian Geometry*, Birkhäuser, 1992.

[6] H. H. Bauschke and P. L. Combettes, *Convex Analysis and Monotone Operator Theory in Hilbert Spaces*, Springer, 2011.