The Coercive Projection Theorem: The Unique Certification Strategy Forced by Canonical Cost

Inevitability, Optimality, and Minimality from First Principles

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

February 11, 2026

Abstract

We prove that the three-step certification template known as the Coercive Projection Method (CPM) is the *unique optimal* strategy for deciding membership in a zero-defect structured set, given a cost functional satisfying the Recognition Composition Law, a conservation constraint, and finite local resolution.

The main result (**Master Theorem**, §5) is:

> *Let $J(x) = \frac{1}{2}(x + x^{-1}) - 1$ be the unique canonical cost. Every correct, finite-data, complete certification procedure on the rational class factors* ***uniquely*** *as $\Phi = \mathcal{A} \circ \mathcal{B} \circ \mathcal{P}$, where $\mathcal{P}$ is the J-projection to neutrality, $\mathcal{B}$ is the coercivity bound, and $\mathcal{A}$ is window aggregation. The factorisation order is forced. The three factors are independent. The constants $c_{\min} = 1/2$ and $C_{\text{proj}} = 1$ are optimal among all cost functionals satisfying the composition law.*

No domain-specific input enters at any stage. CPM is not a method one chooses; it is a theorem one proves.

Contents

1 Introduction **2**
- 1.1 The problem . . . 2
- 1.2 What we prove . . . 2
- 1.3 What we do not do . . . 2

2 Axioms and Setup **2**

3 Certification Procedures: Definition and Partial Order **3**

4 The Three Forced Steps **3**
- 4.1 Step $\mathcal{P}$: Projection to neutrality . . . 3
- 4.2 Step $\mathcal{B}$: Coercivity bound . . . 3
- 4.3 Step $\mathcal{A}$: Aggregation . . . 4

5 The Master Theorem **4**

6 Independence of the Three Steps **5**

7 Optimality of Constants **6**

8 Completeness: CPM Decides the Rational Class **6**

9 Constants from the Forced Geometry **7**

10 Discussion **7**
10.1 CPM is not a method 7
10.2 Relationship to the Recognition Stability Audit 7
10.3 The engineering boundary 7

11 Conclusions **7**

1 Introduction

1.1 The problem

Given a unique cost functional J and only finite observational access (a window of eight measurements per cycle), decide whether a configuration $\mathbf{x}$ has zero defect: $J(\mathbf{x}) = 0$.

This is the most basic audit question. We prove there is exactly one way to answer it.

1.2 What we prove

1. **Existence and uniqueness** of an optimal certification procedure (§5).
2. **Forced factorisation** into exactly three steps: projection $\mathcal{P}$, coercivity bound $\mathcal{B}$, aggregation $\mathcal{A}$.
3. **Independence**: none of the three factors is derivable from the other two (§6).
4. **Optimality of constants**: $c_{\min} = 1/2$ and $C_{\text{proj}} = 1$ are the best possible (§7).
5. **Completeness**: CPM decides membership for every configuration in the rational class (§8).

1.3 What we do not do

We reference no specific mathematical problem, no famous conjecture, and no domain-specific estimate. The entire development is intrinsic to the cost functional J and the finite-resolution axiom.

2 Axioms and Setup

Definition 2.1 (Axiom set $\mathfrak{A}$)**.** A *recognition cost system* is a quadruple $(\mathbb{R}_{>0}, J, \sigma, W)$ where:

(A1) **Cost uniqueness.** $J : \mathbb{R}_{>0} \to \mathbb{R}$ is the unique function satisfying: $J(1) = 0$, the composition law $J(xy) + J(x/y) = 2J(x)J(y) + 2J(x) + 2J(y)$, and the calibration $\lim_{t\to 0} 2J(e^t)/t^2 = 1$.

(A2) **Conservation.** For $\mathbf{x} \in (\mathbb{R}_{>0})^n$, define $\sigma(\mathbf{x}) := \sum_{i=1}^n \ln x_i$. Admissible configurations satisfy $\sigma(\mathbf{x}) = 0$.

(A3) **Finite resolution.** The observation window has length $W = 8$. Measurements are W-periodic: only W consecutive values are accessible per cycle.

The *structured set* is $S := \{\mathbf{x} \in (\mathbb{R}_{>0})^n : J(\mathbf{x}) = 0\} = \{(1, \dots, 1)\}$.

Theorem 2.2 (Explicit form of J)**.** *Under* (A1)*,* $J(x) = \frac{1}{2}(x + x^{-1}) - 1$ *for all* $x > 0$*.*

Proof. Set $H(t) := J(e^t) + 1$. The composition law yields d'Alembert's equation $H(t+u) + H(t-u) = 2H(t)H(u)$ with $H(0) = 1$. Setting $t = 0$: $H(u) + H(-u) = 2H(u)$, so H is even. The calibration gives $H''(0) = 1$. Rewriting d'Alembert as $[H(t+u) - 2H(t) + H(t-u)]/u^2 = 2H(t)(H(u) - 1)/u^2$ and taking $u \to 0$: $H''(t) = H(t)$. With $H(0) = 1$, $H'(0) = 0$ (even), the unique solution is $H(t) = \cosh(t)$. □

We write $\phi(t) := J(e^t) = \cosh(t) - 1$ for the log-coordinate form.

Proposition 2.3 (Properties of J that drive everything)**.**

(P1) ϕ *is even:* $\phi(t) = \phi(-t)$*. (Equivalently,* $J(x) = J(x^{-1})$*.)*

(P2) $\phi(t) \geq 0$ *with* $\phi(t) = 0 \iff t = 0$*.*

(P3) $\phi''(t) = \cosh(t) \geq 1$ *for all* t*. (*ϕ *is* 1*-strongly convex.)*

(P4) $\phi(t) \geq t^2/2$ *for all* t*, with equality iff* $t = 0$*.*

(P5) $\phi(t) \leq t^2/2 + t^4/24$ *for* $|t| \leq 1$*.*

Proof. (P1)–(P2): cosh is even and $\cosh(t) \geq 1$. (P3): $\phi''(t) = \cosh(t) \geq \cosh(0) = 1$. (P4): $\cosh(t) = 1 + t^2/2 + t^4/24 + \cdots \geq 1 + t^2/2$. (P5): the first omitted term is $t^6/720 > 0$. □

3 Certification Procedures: Definition and Partial Order

Definition 3.1 (Certification procedure)**.** A *certification procedure* for $(\mathbb{R}_{>0}, J, \sigma, W)$ is a map

$$\Phi : (\mathbb{R}_{>0})^n \longrightarrow \{\texttt{ZERO}, \texttt{NONZERO}, \texttt{INCONCLUSIVE}\}$$

satisfying:
(C1) **Soundness.** $\Phi(\mathbf{x}) = \texttt{ZERO} \implies \mathbf{x} \in S$, and $\Phi(\mathbf{x}) = \texttt{NONZERO} \implies \mathbf{x} \notin S$.
(C2) **Finite data.** $\Phi(\mathbf{x})$ depends on at most $W = 8$ consecutive evaluations per cycle of a finite-state representation of $\mathbf{x}$.
A procedure is *complete on the rational class* if, for every $\mathbf{x}$ whose generating function is rational of known degree, $\Phi(\mathbf{x}) \neq \texttt{INCONCLUSIVE}$.

Definition 3.2 (Optimality ordering)**.** For procedures Φ, Ψ, write $\Phi \succeq \Psi$ if: whenever $\Psi(\mathbf{x}) \neq \texttt{INCONCLUSIVE}$, we have $\Phi(\mathbf{x}) = \Psi(\mathbf{x})$; and there exists $\mathbf{x}$ with $\Phi(\mathbf{x}) \neq \texttt{INCONCLUSIVE}$ but $\Psi(\mathbf{x}) = \texttt{INCONCLUSIVE}$.

A procedure Φ^* is *optimal* if $\Phi^* \succeq \Phi$ for every sound, finite-data procedure Φ.

4 The Three Forced Steps

4.1 Step $\mathcal{P}$: Projection to neutrality

Theorem 4.1 (J-projection)**.** *For every $\mathbf{x} \in (\mathbb{R}_{>0})^n$, the problem*

$$\min\Big\{\sum_{i=1}^{n} \phi(t_i) \ : \ \sum_{i=1}^{n} t_i = -\sigma(\mathbf{x})\Big\}$$

has a unique solution $t_1 = \cdots = t_n = -\sigma(\mathbf{x})/n$. In log-coordinates $y_i = \ln x_i$, the corrected state is

$$\mathcal{P}(\mathbf{y})_i \ = \ y_i - \bar{y}, \qquad \bar{y} := \frac{1}{n}\sum_j y_j. \tag{1}$$

This is the Euclidean orthogonal projection onto $H = \{\mathbf{y} : \sum_i y_i = 0\}$.

Proof. The constraint is linear. The objective $\sum_i \phi(t_i)$ is strictly convex (ϕ is 1-strongly convex by (P3)), so by Jensen's inequality $\frac{1}{n}\sum_i \phi(t_i) \geq \phi(\frac{1}{n}\sum_i t_i)$ with equality iff $t_1 = \cdots = t_n$. The constraint sum is $-\sigma(\mathbf{x})$, so the common value is $-\sigma(\mathbf{x})/n$. Translating to log-coordinates: $y_i' = y_i + t_i = y_i - \bar{y}$. □

Lemma 4.2 (Projection constants)**.** *$\mathcal{P}$ satisfies:*
1. *$\|\mathcal{P}(\mathbf{y})\| \leq \|\mathbf{y}\|$ (nonexpansive; $C_{\mathrm{proj}} = 1$).*
2. *$\mathcal{P}^2 = \mathcal{P}$ (idempotent).*
3. *$\mathcal{P}(\mathbf{y}) = \mathbf{y}$ iff $\mathbf{y} \in H$ (fixes neutral configurations).*

Proof. Orthogonal projections in Euclidean space satisfy all three. □

4.2 Step $\mathcal{B}$: Coercivity bound

Theorem 4.3 (Coercivity)**.** *For all $\mathbf{y} \in \mathbb{R}^n$:*

$$\sum_{i=1}^{n} \phi(y_i) \ \geq \ \frac{1}{2}\|\mathbf{y}\|^2, \tag{2}$$

with equality iff $\mathbf{y} = \mathbf{0}$.

Proof. Sum (P4) over components: $\sum_i \phi(y_i) \ge \sum_i y_i^2/2 = \|\mathbf{y}\|^2/2$. Equality holds iff each $y_i = 0$. ☐

Corollary 4.4 (Reverse coercivity)**.** $\sum_i \phi(y_i) \le \varepsilon \implies \|\mathbf{y}\| \le \sqrt{2\varepsilon}$.

Theorem 4.5 (Tight upper bound)**.** *For $\|\mathbf{y}\| \le 1$:*

$$\sum_{i=1}^{n} \phi(y_i) \;\le\; \frac{1}{2}\|\mathbf{y}\|^2 + \frac{1}{24}\|\mathbf{y}\|_4^4, \tag{3}$$

where $\|\mathbf{y}\|_4^4 = \sum_i y_i^4$. In particular, ϕ and $\|\cdot\|^2/2$ agree to second order at the identity: the coercivity constant $c_{\min} = 1/2$ is tight.

Proof. Sum (P5) over components. ☐

4.3 Step $\mathcal{A}$: Aggregation

Definition 4.6 (Window sums)**.** Partition indices $\{1, \ldots, n\}$ into consecutive blocks of size $W = 8$ (pad with zeros if $8 \nmid n$). The k-th window sum is $W_k := \sum_{j \in \text{block } k} y_j$.

Proposition 4.7 (Window neutrality implies global neutrality)**.** *If $W_k = 0$ for all k, then $\sigma(\mathbf{x}) = \sum_i y_i = 0$.*

Proof. $\sum_i y_i = \sum_k W_k = 0$. ☐

Theorem 4.8 (Finite determination in the rational class)**.** *If $\mathbf{y} = (y_0, y_1, \ldots)$ is the output of a finite-state system with d states, then the generating function $\theta(z) = \sum_n y_n z^n$ is rational of degree $\le d$. In particular, θ is uniquely determined by any $2d+1$ consecutive values, and any global property (e.g., "all $y_n = 0$") is decidable from finite data.*

Proof. Enumerate the states $S = \{1, \ldots, d\}$ with transition matrix A and output vector u. Then $y_n = u^* A^n v$ for initial vector v, so $\theta(z) = u^*(I - zA)^{-1}v$ is a ratio of polynomials of degree $\le d$. A rational function of degree $\le d$ is determined by $2d+1$ values (its d numerator and d denominator coefficients, plus one normalisation). ☐

Corollary 4.9 (Aggregation in the rational class)**.** *If $\mathbf{y}$ lies in the rational class of degree d and $\lceil n/8 \rceil \ge 2d+1$, then the window sums $\{W_0, \ldots, W_{\lceil n/8 \rceil - 1}\}$ determine $\mathbf{y}$ uniquely. In particular, $W_k = 0$ for all k implies $\mathbf{y} = \mathbf{0}$.*

Proof. The window sums are linear functionals of θ evaluated at 8-spaced blocks. For degree $\le d$, $2d+1$ such functionals are sufficient for unique reconstruction. ☐

Proposition 4.10 (Finite sampling alone fails)**.** *Without the rational-class restriction, no finite set of evaluations determines a function globally: for any samples $(z_1, w_1), \ldots, (z_m, w_m)$ and any $a \notin \{z_i\}$, there exists a meromorphic function matching all samples with a pole at a.*

Proof. $f(z) = p(z) + \prod_i (z - z_i)/(z - a)$ where p is the Lagrange interpolant. ☐

5 The Master Theorem

Theorem 5.1 (Unique optimal certification)**.** *Under axioms* (A1)–(A3)*, define*

$$\Phi^*(\mathbf{x}) := \begin{cases} \texttt{ZERO} & \text{if } \mathcal{A} \circ \mathcal{B} \circ \mathcal{P} \text{ certifies } J(\mathbf{x}) = 0, \\ \texttt{NONZERO} & \text{if } \mathcal{A} \circ \mathcal{B} \circ \mathcal{P} \text{ certifies } J(\mathbf{x}) > 0, \\ \texttt{INCONCLUSIVE} & \text{otherwise,} \end{cases} \tag{4}$$

where $\mathcal{P}$ is the J-projection (1)*, $\mathcal{B}$ is the coercivity bound* (2)*, and $\mathcal{A}$ is window aggregation (Corollary 4.9).*

Then:

(I) ***Soundness:*** *Φ^* satisfies* (C1).

(II) ***Completeness:*** *Φ^* is complete on the rational class.*

(III) ***Optimality:*** *Φ^* is optimal (Definition 3.2).*

(IV) ***Uniqueness:*** *Φ^* is the unique optimal procedure.*

(V) ***Forced factorisation:*** *$\Phi^* = \mathcal{A} \circ \mathcal{B} \circ \mathcal{P}$ is the unique order; no permutation of the three steps yields a sound procedure.*

Proof. We prove each claim.

(I) (Soundness). Suppose $\Phi^*(\mathbf{x}) = \texttt{ZERO}$. Then $\mathcal{A}$ has certified that all window sums of the projected sequence $\mathbf{y}' = \mathcal{P}(\mathbf{y})$ vanish. By Corollary 4.9 (finite determination in the rational class), $\mathbf{y}' = \mathbf{0}$, hence $\mathbf{y} \in H$ and $\mathbf{y} = \mathbf{y}' = \mathbf{0}$, hence $\mathbf{x} = (1, \ldots, 1) \in S$.

Suppose $\Phi^*(\mathbf{x}) = \texttt{NONZERO}$. Then $\mathcal{A}$ has found a nonzero window sum. By Proposition 4.7, $\sigma(\mathbf{x}) \neq 0$ or some $y_i \neq 0$. If $\sigma \neq 0$, then $\mathbf{x} \notin S$ (admissibility). If $\sigma = 0$ but some $W_k \neq 0$, then $\mathcal{P}(\mathbf{y}) \neq \mathbf{0}$, so $J(\mathbf{x}) > 0$ by (2), hence $\mathbf{x} \notin S$.

(II) (Completeness). If $\mathbf{x}$ is in the rational class of degree d and $n \geq 8(2d+1)$, then the window sums determine $\mathbf{y}$ uniquely (Corollary 4.9). The procedure can then decide $J(\mathbf{x}) = 0$ or > 0 with certainty, so $\Phi^*(\mathbf{x}) \neq \texttt{INCONCLUSIVE}$.

(III) (Optimality). Let Ψ be any sound, finite-data procedure. We show $\Phi^* \succeq \Psi$. Suppose $\Psi(\mathbf{x}) = \texttt{ZERO}$. Then $\mathbf{x} \in S$ by soundness of Ψ, hence $\Phi^*(\mathbf{x}) = \texttt{ZERO}$ (since $\mathbf{x} \in S$ is detected by any number of window sums). Similarly for $\texttt{NONZERO}$. It remains to show Φ^* resolves strictly more cases than Ψ can. But Φ^* uses the *sharpest* bound at every step:

- $\mathcal{P}$ is the unique cost-minimising projection (Theorem 4.1); any other projection overestimates the correction cost, making the residual a noisier signal.
- $\mathcal{B}$ uses the tight bound $\phi \geq \|\cdot\|^2/2$ (Theorem 4.3 + Theorem 4.5); any looser bound widens the inconclusive zone.
- $\mathcal{A}$ uses full rational reconstruction (Theorem 4.8); any weaker aggregation (e.g., only checking a subset of windows) leaves more cases inconclusive.

Therefore Φ^* resolves every case Ψ resolves plus potentially more.

(IV) (Uniqueness). Suppose Φ^{**} is also optimal. Then $\Phi^* \succeq \Phi^{**}$ and $\Phi^{**} \succeq \Phi^*$, so they agree on every resolved case. By completeness on the rational class, every rational-class input is resolved. For non-rational inputs, both return $\texttt{INCONCLUSIVE}$ (since no finite-data procedure can resolve them by Proposition 4.10). Therefore $\Phi^* = \Phi^{**}$.

(V) (Forced order). We show no permutation of $(\mathcal{P}, \mathcal{B}, \mathcal{A})$ other than $\mathcal{A} \circ \mathcal{B} \circ \mathcal{P}$ is sound.

$\mathcal{P}$ must come first. $\mathcal{B}$ (coercivity) and $\mathcal{A}$ (aggregation) require $\mathbf{y} \in H$ (the neutral hyperplane). On unrestricted $\mathbb{R}^n$, the coercivity bound (2) gives $J \geq \|\mathbf{y}\|^2/2$, but this does not distinguish "admissible with nonzero defect" from "inadmissible (conservation violated)." Without first projecting, a procedure cannot separate these two failure modes and therefore cannot be sound.

$\mathcal{B}$ must come before $\mathcal{A}$. Aggregation $\mathcal{A}$ produces window sums $\{W_k\}$—discrete data. To convert "all $W_k = 0$" into "$J = 0$" requires the coercivity inequality: $\mathbf{y}' = \mathbf{0} \implies J = 0$ is precisely $\|\mathbf{y}'\| = 0 \implies \sum \phi(y_i') = 0$, which is the coercivity bound applied at $\|\mathbf{y}'\| = 0$. Without $\mathcal{B}$, the logical link from window sums to cost is missing. □

6 Independence of the Three Steps

Theorem 6.1 (Independence)**.** *No step of $(\mathcal{P}, \mathcal{B}, \mathcal{A})$ is derivable from the other two. Specifically:*

(a) *$\mathcal{B}$ and $\mathcal{A}$ alone (without $\mathcal{P}$) cannot distinguish admissibility violations from nonzero defect.*

(b) *$\mathcal{P}$ and $\mathcal{A}$ alone (without $\mathcal{B}$) cannot certify $J = 0$ from $\|\mathbf{y}'\| = 0$.*

(c) *$\mathcal{P}$ and $\mathcal{B}$ alone (without $\mathcal{A}$) cannot certify global properties from finite data.*

Proof. **(a)** Consider $\mathbf{y} = (1, 1, \ldots, 1)$ (all components equal to 1). Then $\sigma = n \neq 0$ (inadmissible) but $J(\mathbf{x}) > 0$. Without $\mathcal{P}$, we cannot tell whether the nonzero cost is due to conservation violation or genuine defect. In contrast, $\mathbf{y} = (2, -2, 0, \ldots, 0)$ has $\sigma = 0$ (admissible) and $J > 0$ (genuine defect). The two cases have the same cost profile to $\mathcal{B}$ but different certification outcomes.

(b) Consider the implication "$\|\mathbf{y}'\| = 0 \implies J(\mathbf{x}') = 0$." This uses $\phi(0) = 0$, which is the base case of the coercivity inequality. Without $\mathcal{B}$ (i.e., without knowing that ϕ is non-negative with unique zero at 0), the projection $\mathcal{P}$ could return $\mathbf{0}$ for a configuration that nevertheless has nonzero cost under a different functional. The link $\|\cdot\| = 0 \iff \phi = 0$ is the content of $\mathcal{B}$.

(c) By Proposition 4.10, finite evaluations of $\mathbf{y}$ cannot determine $\mathbf{y}$ globally without the rational-class restriction. Aggregation $\mathcal{A}$ supplies this restriction (via Theorem 4.8). Without it, $\mathcal{P}$ and $\mathcal{B}$ can certify only the sampled points, not the global defect. □

7 Optimality of Constants

Theorem 7.1 (Optimal coercivity constant)**.** *Among all cost functionals $F : \mathbb{R}_{>0} \to \mathbb{R}_{\geq 0}$ satisfying axioms* (A1)*, the coercivity constant*

$$c(F) := \inf_{t \neq 0} \frac{F(e^t)}{t^2/2}$$

satisfies $c(F) = 1$. That is, $c_{\min} = 1/2$ is the largest (best) constant possible: $F(e^t) \geq t^2/2$ for all t, and this is tight.

Proof. By Theorem 2.2, $F = J$ is the unique solution. We have $\phi(t)/\frac{t^2}{2} = \frac{\cosh(t)-1}{t^2/2}$. As $t \to 0$, this ratio $\to 1$ (by L'Hôpital or Taylor). For $t \neq 0$, $\cosh(t) - 1 > t^2/2$ strictly. Therefore $\inf_{t \neq 0} \phi(t)/(t^2/2) = 1$, achieved in the limit $t \to 0$. Hence $c(J) = 1$ and $c_{\min} = 1/2$.

Since J is unique, there is no other F to compare. The constant cannot be improved. □

Theorem 7.2 (Optimal projection constant)**.** *Among all corrections $\mathbf{r}$ that restore neutrality ($\mathbf{x} \odot \mathbf{r} \in \mathcal{M}$), the J-minimising correction is the orthogonal projection (Theorem 4.1) with Lipschitz constant $C_{\text{proj}} = 1$. No correction can have $C_{\text{proj}} < 1$.*

Proof. $C_{\text{proj}} = 1$ because orthogonal projection is nonexpansive (Lemma 4.2). No linear map onto a proper subspace can have Lipschitz constant strictly less than 1 (the orthogonal projection achieves the minimum). □

Corollary 7.3 (The constants are intrinsic to reality)**.** *$c_{\min} = 1/2$ is fixed by the calibration $J''(1) = 1$, which is fixed by the composition law. $C_{\text{proj}} = 1$ is fixed by reciprocal symmetry $J(x) = J(x^{-1})$. Both follow from axiom* (A1) *alone. The certification template and its constants are as inevitable as the cost functional.*

8 Completeness: CPM Decides the Rational Class

Theorem 8.1 (Decision procedure)**.** *For $\mathbf{x}$ in the rational class of known degree d with $n \geq 8(2d+1)$, Φ^* terminates with output `ZERO` or `NONZERO`; it never returns `INCONCLUSIVE`.*

Proof. By Theorem 4.8, $\mathbf{y}$ is determined by $2d+1$ values. With $n/8 \geq 2d+1$ window sums available, $\mathcal{A}$ reconstructs $\mathbf{y}$ exactly. Then $\mathcal{B}$ computes $J(\mathbf{x})$ exactly, and the decision is $J = 0$ or > 0. □

Remark 8.2 (Scope boundary). Outside the rational class, Φ^* may return `INCONCLUSIVE`. This is not a deficiency: by Proposition 4.10, no finite-data procedure can decide membership for arbitrary analytic functions. The rational-class boundary is sharp.

9 Constants from the Forced Geometry

In the full Recognition Science framework, the window length $W = 8$ is not a choice: it is the minimal covering cycle for a D-cube walk with $D = 3$ (forced by linking + gap-45 synchronisation). The complete constant table:

Constant	**Value**	**Origin**	**Optimal?**
$c_{\min}$	$1/2$	$\phi''(0) = 1$ (calibration (A1))	Yes (Thm 7.1)
C_{proj}	1	Orthogonal proj. (reciprocity (P1))	Yes (Thm 7.2)
K_{net}	1	Single covering window	—
W	8	$2^D = 2^3$ (minimal cover)	Forced
L	$\frac{1}{1+\lambda}$	ϕ is 1-strongly convex ((P3))	Forced

No entry in this table is adjustable.

10 Discussion

10.1 CPM is not a method

The word "method" implies choice. The Master Theorem 5.1 proves there is no choice: the template, the order of steps, and the constants are all forced by the axioms. A more accurate name would be the *Coercive Projection Theorem* (CPT) or simply the *Certification Theorem of Canonical Cost*.

10.2 Relationship to the Recognition Stability Audit

The RSA paper proves an *impossibility* certificate: if a candidate forces a sensor pole, the Cayley–Schur pinch excludes it. CPM proves a *membership* certificate: if all window tests pass, the configuration is in S. Together:

$$\text{CPM (membership)} \;+\; \text{RSA (exclusion)} \;=\; \text{complete two-sided audit.}$$

Both derive from J's strict convexity. They are the existence and impossibility faces of the same coin.

10.3 The engineering boundary

Everything in this paper is foundation. What remains for any specific application is engineering: identify the structured set S, supply the domain-specific finite-state model, and run Φ^*. The template requires no domain input.

11 Conclusions

1. $J(x) = \frac{1}{2}(x + x^{-1}) - 1$ is uniquely forced by the composition law, normalization, and calibration.
2. The certification procedure $\Phi^* = \mathcal{A} \circ \mathcal{B} \circ \mathcal{P}$ is the **unique optimal** procedure for deciding $J = 0$ from finite data (Master Theorem 5.1).
3. The factorisation is **forced**: no reordering of the three steps is sound.
4. The three steps are **independent**: none is derivable from the other two (Theorem 6.1).
5. The constants $c_{\min} = 1/2$ and $C_{\text{proj}} = 1$ are **optimal**: no cost functional satisfying the axioms achieves better (Theorems 7.1–7.2).

6. Φ^* is **complete** on the rational class and **sharp** at its boundary: outside the rational class, no finite-data procedure can do better (Theorem 8.1 + Proposition 4.10).

The Coercive Projection Method is not a method. It is a theorem.

References

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

[2] J. Aczél, *Lectures on Functional Equations*, Academic Press, 1966.

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

[4] E. F. Beckenbach and R. Bellman, *Inequalities*, Springer, 1961.

[5] J. Washburn, "The Recognition Stability Audit," RS preprint, 2026.

[6] J. Washburn, "The Algebra of Reality," *Axioms* **15**(2), 90 (2025).