

# Tightness and Independence in a Calibrated d'Alembert Characterization of Reciprocal Comparison Cost

Jonathan Washburn<sup>1</sup>, Sebastian Pardo-Guerra<sup>1</sup>, Megan Simons<sup>1</sup>,  
Anil Thapa<sup>1\*</sup>

<sup>1\*</sup>Recognition Physics Institute, Austin, TX, USA.

\*Corresponding author(s). E-mail(s): [athapa@recognitionphysics.org](mailto:athapa@recognitionphysics.org);

## Abstract

We analyze the hypothesis set behind the calibrated reciprocal comparison cost  $\mathbf{J}(\mathbf{x}) = \frac{1}{2}(\mathbf{x} + \mathbf{x}^{-1}) - \mathbf{1}$ . Earlier work established the uniqueness of this cost under a reciprocal d'Alembert-type composition law and a quadratic calibration. The present paper asks what that hypothesis set really contains. We prove that finite pairwise polynomial closure is equivalent to a multiplicative shifted-cost composition principle, and that this closure hypothesis is logically independent of the remaining comparison axioms by an explicit quartic-logarithmic witness. We also show that the nondegeneracy and calibration clauses are indispensable, formulate the positive-ratio carrier from standard measurement axioms through Hölder's theorem, and record a measurable-regularity variant of the uniqueness theorem. Two Lean-kernel certificates fix the scope of possible extrapolations: no non-trivial dimensionless rational monomial of the speed of light  $c_{\text{light}}$ ,  $\hbar$ , and  $\mathbf{G}$  exists, so an external dimensional unit calibration is unavoidable, and the composition law is the essential input that selects  $\mathbf{J}$ . Thus the main contribution is a tightness and scope theorem, not a claim that the cost or any physical structure follows from weaker assumptions.

**Keywords:** functional equations, d'Alembert equation, reciprocal cost, formal verification, axiom independence, foundations of physics

**MSC Classification:** 39B22; 03B35

# 1 Introduction

How should the mismatch between two positive quantities be priced? Write the comparison as a ratio  $x > 0$ , so that  $x = 1$  means perfect agreement and  $x$  and  $x^{-1}$  describe the same mismatch seen from the two sides. A cost for such comparisons should vanish at the identity, treat  $x$  and  $x^{-1}$  alike, and behave coherently when two comparisons are composed. Sharpened by a nondegeneracy requirement, a closure law, and a calibration, these demands become seven named conditions, (C1)–(C7) of Definition 6.2, and this paper proves that they admit exactly one continuous solution on  $\mathbb{R}_{>0}$ :

$$J(x) = \frac{1}{2}(x + x^{-1}) - 1$$

(Theorem 7.1). Without the final curvature calibration (C7b), the solutions form, after the cost-unit normalization (C7a), the one-parameter hyperbolic family  $\cosh(\lambda \log x) - 1$ ; before that normalization the family carries an additional overall scale,  $A(\cosh(\lambda \log x) - 1)$  with combiner coefficient  $c = 2/A$ . Nothing narrower survives.

The value of a characterization theorem lies in its hypotheses, so we analyze them one at a time. Continuity can be weakened to measurability with local boundedness (Corollary 7.2). The closure hypothesis (C6) is equivalent to an interpretable multiplicative composition law (Theorem 6.3) and is proved logically independent of the remaining axioms (Proposition 6.4). The selection and calibration conditions (C5), (C6), and (C7b) are each shown indispensable by explicit counterexamples (Remark 6.10). The continuum carrier itself is reduced to standard measurement axioms through Hölder’s theorem (Proposition 5.5). The theorem is compact; the tightness suite around it is the mathematical contribution.

The second contribution is a precise statement of scope. The theorem arose within our broader research program (Recognition Science), which reads  $J$  as the cost of an elementary act of comparison and builds physical structure on top of it. For such a construction, precision about inputs is as valuable as the theorem itself: which conclusions follow from the mathematics alone, and which require further physical hypotheses. Three scope results settle this in permanent form: two are certificates checked by the Lean kernel in the same public repository that formalizes the theorem’s calibrated endpoint (the endpoint’s uniqueness theorem assumes a smoothness typeclass that has no instance in the pinned slice; the instance is proved in the program’s full development library and reproduced in prose in Appendix A, as detailed in the Data and code availability statement), and one is proved in this paper.

Section 8 states all three. The first, kernel-checked, fixes the division of labor for dimensionful constants: the exponent matrix of  $(c_{\text{light}}, \hbar, G)$  is nonsingular, so no non-trivial dimensionless monomial in them exists, and no framework whose outputs are pure numbers can produce their SI values without external dimensional unit calibration. The second, kernel-checked, identifies the input that selects  $J$ : symmetry, normalization, continuity, and calibration by themselves admit other costs (an explicit quadratic one), so the composition law is the essential ingredient. The third, proved in Part I, completes the tightness analysis: a quartic-logarithmic cost shows that the closure hypothesis does real work. Section 9 then presents three conditional structures from the wider framework (a golden-ratio scale ladder, a conditional dimension-three argument, and

an eight-tick recognition cycle), each stated with the exact hypotheses on which it rests. What remains open is stated in place and collected in the conclusion.

The paper is organized as follows. Section 2 places the result in the literature and states precisely what is new relative to our published companion papers. Section 3 fixes the claim-status conventions and the standing hypotheses used throughout. The remainder has two parts. Part I (Sections 4–7) contains the mathematics: the recognition-work counting results and the uniqueness theorem with its tightness suite. Part II (Sections 8 and 9) maps the scope: the certificates and the conditional structures. These two parts are independent by design: Part II is not needed for, and cannot strengthen, the theorem of Part I; a reader interested only in the mathematics can stop at the end of Section 7. The proved content of the paper is Part I together with the certificates of Section 8; every physics-facing statement in Part II carries its hypotheses explicitly.

## 2 Related Work and Contribution

The method of Part I has a definite lineage. D’Alembert wrote down the functional equation that now carries his name in 1769, in a study of the composition of forces. Cauchy made the additive equation a template for extracting an exact formula from qualitative constraints plus a regularity assumption. Aczél’s 1966 treatise consolidated a century and a half of such results into a discipline [1]. The same instinct drives the modern reconstructions of quantum theory from information-theoretic postulates [15, 16]: a short list of operational axioms should force a formalism, not merely accommodate it, and a derivation is judged by the independence and interpretability of its axioms as much as by its conclusion. Part I is an entry in the first tradition held to the standard of the second. It offers one composition law, one calibrated solution, and a tightness suite recording what each hypothesis pays for.

In spirit the framework sits among programs that place information, relation, or operational comparison at the base of physics [10–15]; we note the kinship only to set it aside. Unlike those programs, the present paper reconstructs neither Hilbert space, nor the Born rule, nor gauge structure, and neither does the pinned Lean slice it cites. Any such claims elsewhere in the broader program are not used as premises in this paper. This paper proves one functional-equation characterization and fixes, in machine-checked form, what does and does not follow from it.

At the level of the theorem itself, the closest relatives are not foundational programs but the classical functional-equation literature. The classification of continuous solutions of the d’Alembert equation  $H(s+t) + H(s-t) = 2H(s)H(t)$  into constant, cosine, and hyperbolic-cosine branches is due to Aczél [1], with the measurability refinement in Aczél–Dhombres [2]. Stage (iii) of the proof of Theorem 7.1 is exactly this classification, after a logarithmic change of variables and a shift. What the present theorem adds is not a new solution but the reduction (Proposition 6.2): the route-independence ansatz with a symmetric degree- $\leq 2$  polynomial combiner and the identity boundary condition force the combiner  $\Phi(u, v) = 2u + 2v + cuv$ , after which the d’Alembert machinery applies. The result is best read as a calibrated d’Alembert-type characterization with recognition-motivated hypotheses.

The theorem is also distinct from the standard axiomatizations of divergences and losses, such as Bregman divergences, Csiszár  $f$ -divergences, and proper scoring rules [19–21]. Those families are characterized by convexity, monotonicity under coarse-graining, or properness. None of these properties appears among (C1)–(C7), and a generic member of those families fails the two-channel composition law (C3). Conversely, up to the substitution  $x = a/b$ , the cost  $J$  is the symmetric divergence  $\frac{(a-b)^2}{2ab}$ . The theorem can therefore be read as singling out one specific symmetric divergence by composition rather than by convex geometry.

The scope certificates of Section 8 belong to a productive tradition in foundations: theorems that establish exactly which inputs a conclusion requires. Bell’s theorem and the Kochen–Specker theorem [17, 18] are the canonical examples, and dimensional analysis since Bridgman [26] plays the same role for dimensionful constants. What we add is mechanical permanence: two of the three scope results are Lean-kernel certificates at a pinned public commit.

## Strength of the hypotheses

The hypotheses of Theorem 7.1 are strong, and the two load-bearing ones already point at the d’Alembert structure. Finite closure (C6) is equivalent, for some nonzero rescaling  $\nu$ , to the requirement that  $1 + \nu F$  composes multiplicatively across the two channels (Theorem 6.3). In logarithmic coordinates that is precisely the d’Alembert equation. The theorem is therefore a characterization relative to an interpretable composition principle, not a derivation of that principle from anything weaker. The claim we defend is not that  $J$  emerges from weak assumptions. It is that the assumption set is named, its dependency order is explicit, and each selection step has a stated witness. A reader who grants (C6) only in its multiplicative form may read Theorem 7.1 as a calibrated repackaging of the Aczél classification. We accept that reading; the novelty rests on the tightness suite, not on the classification step.

Parts of this material are already published. The uniqueness of the canonical reciprocal cost under a composition law is established in [7], the rigidity of the bilinear combiner family in [8], reciprocal convex costs for ratio matching in [6], the passage from coherent comparison to discrete ledger dynamics in [9], and the geometric setting in [5]. The present paper restates the first two items for self-containment and, as Table 1 records, separates them from the new tightness, independence, and scope results.

Thus the novelty claim is deliberately narrow: the paper is not a second announcement of the uniqueness theorem, but an analysis of the hypothesis set as a whole. Its new contributions are the equivalence of the closure hypothesis with a multiplicative composition law (Theorem 6.3), its logical independence from the remaining axioms (Proposition 6.4), the separate indispensability witnesses for the selection and calibration conditions (Remark 6.10), the measurement-theoretic derivation of the continuum carrier (Proposition 5.5), and the measurable-regularity variant (Corollary 7.2), together with the machine-checked scope of Section 8. The paper is self-contained at the level of these claims: no proof below depends on the companion papers, and no assumption is discharged by citing them.

**Table 1** Relationship of the results used in this paper to the authors’ prior publications

| Result used in this paper                                                   | Status relative to prior work                               |
|-----------------------------------------------------------------------------|-------------------------------------------------------------|
| Canonical-cost uniqueness under composition and calibration                 | Previously published in [7]; restated for self-containment. |
| Degree-two polynomial-combiner rigidity                                     | Previously published in [8]; restated for self-containment. |
| Closure equivalence with multiplicative shifted-cost composition            | Presented here as a new synthesis and theorem.              |
| Logical independence of finite closure from the remaining comparison axioms | Presented here as a new independence witness.               |
| Hölder derivation of the continuum carrier from measurement axioms          | Presented here as a new carrier analysis.                   |
| Measurable-regularity formulation of the uniqueness theorem                 | Presented here as a new variant for this hypothesis set.    |
| Lean-audited scope certificates and formal boundary statements              | Presented here as new scope-control material.               |

**Table 2** Claim-status convention used throughout the paper.

| Status              | Meaning in this paper                                                                                                                                                                                                                                                                                                                                                                                                                        |
|---------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| THEOREM/CERTIFICATE | A theorem or certificate present in the public Lean repository at the pinned commit recorded in the Data and code availability statement, or proved in this paper. Where relevant the text distinguishes four grades: kernel-checked and closed; kernel-checked conditionally on a hypothesis undischarged at the pinned commit (so stated); finite arithmetic rechecked by kernel decide; and proved in this paper without a Lean artifact. |
| NATIVE ALGEBRA      | An identity in recognition-native units, such as $\hbar_R = \varphi^{-5}$ or $G_R = \varphi^5/\pi$ , where $\varphi = (1 + \sqrt{5})/2$ is the golden ratio obtained conditional on hypothesis (SI6) in Section 9. These are unit definitions, not SI derivations (Certificate 8.1).                                                                                                                                                         |
| BRIDGE HYPOTHESIS   | An identification between recognition-native structures and physical quantities.                                                                                                                                                                                                                                                                                                                                                             |
| OPEN PROBLEM        | A stated target with its missing bridge named exactly, in place where it arises.                                                                                                                                                                                                                                                                                                                                                             |

### 3 Conventions and Standing Hypotheses

We distinguish four kinds of claim (Table 2). The distinction matters because the public Lean repository proves a formal recognition spine and a native algebra layer; physics beyond that layer enters only through labeled bridge hypotheses. The physics-facing material below identifies which of these four statuses applies wherever it appears.

One reading rule governs every claim in this paper:

*A conditional result is never inflated into a prediction.*

The dimension-three selection, for example, is conditional on an embedding postulate and is stated as such; wherever a hypothesis is open, it is named where it enters and collected in the conclusion (Section 10).

### 3.1 Standing hypotheses

The argument moves from a logical floor to a reciprocal comparison cost, and from there, under additional hypotheses, to scale structure and physical identifications. Every hypothesis is stated either among the ratio-level conditions (C1)–(C7) of Definition 6.2 or in the wider-framework catalogue SI1–SI11 below; overlaps are cross-referenced explicitly. The entries are not epistemically uniform. Three are genuinely open, load-bearing bridge inputs: the positive-ratio ledger (SI2) (reduced to measurement axioms by Proposition 5.5), the adjacent additive rule (SI6), and the embedding postulate (SI10) (isolated as the single input behind the dimension-three selection). Finite closure (SI3) is different: it is an adopted closure axiom, proved logically independent of the remaining cost axioms by Proposition 6.4, not an open physical bridge. The remaining entries are lighter. Hypothesis (SI1) is motivational and not consumed by the theorem. Hypothesis (SI4) is an ordinary regularity assumption, weakened in Corollary 7.2. Hypothesis (SI5) is a calibration convention plus one curvature condition. Hypotheses (SI7)–(SI9) are satisfiable choices exhibited by explicit construction in the pinned repository. Hypothesis (SI11) is definitional for the comparison-cost reading, though mathematically load-bearing (Remark 6.10). The uniqueness theorem of Part I uses only the ratio-level assumptions (C1)–(C7): normalization, reciprocity, route independence, the closure principle (SI3), regularity (SI4), calibration (SI5), and nondegeneracy (SI11), with (SI2) entering only as the bridge that supplies the cost object (Remark 6.1). It invokes none of SI1 or SI6–SI10. The paper’s title and claims are calibrated accordingly.

- (SI1) Recognition work: dichotomy plus exact additivity over semantically independent configurations (disjoint syntactic supports together with joint satisfiability of consistent parts; Definition 5.2, Section 5).
- (SI2) Positive-ratio ledger: a finite recognition ledger supplies positive intensities whose comparisons are recorded as ratios (Section 6).
- (SI3) Finite pairwise polynomial closure on the cost combiner (Section 6, Proposition 6.2); equivalently, multiplicative two-channel closure of a nonzero rescaling  $1 + \nu F$  (Theorem 6.3).
- (SI4) Continuity on the positive-ratio cost (Theorem 7.1; the measurable-regularity weakening is Corollary 7.2, via Appendix A).
- (SI5) Unit log-curvature calibration at the identity ratio (Section 6.3).
- (SI6) Constant oriented one-rung cost and adjacent additive closure on the resulting geometric ladder (Section 9).
- (SI7) Loop-pair separator: disjoint closed ledger loops admit a non-trivial additive oriented separator (Section 9).
- (SI8) Loop-class rank: the abelianized loop-class group has rank at least two (Section 9).
- (SI9) Rank-one target minimality: the minimal torsion-free target for the separator is identified with  $\mathbb{Z}$  (Section 9).

- (SI10) External topology: ordinary integer linking of one-dimensional closed loops, detected by first homology of the complement via classical Alexander duality (Section 9).
- (SI11) Comparison-cost nondegeneracy: the reciprocal cost is non-negative on  $\mathbb{R}_{>0}$ , vanishes only at the identity ratio, and is unbounded as the log-ratio diverges. This is condition (C5) of Definition 6.2 and is required for the branch selection in Theorem 7.1 (it excludes the constant branch, which has no comparison content, and the cosine branch, whose zeros are not isolated to the identity).

Later statements cite these by name where they enter. The loop, separator, and loop-class vocabulary of (SI7)–(SI10) is not defined at the cost level of Part I; one-line glosses sufficient to read Section 9 are collected at the head of that section, and the full definitions are in the pinned repository and the companion papers. The statements of (SI7)–(SI10) are therefore inputs to conditional structures, not self-contained theorems of this paper.

## 4 The Primitive Floor

Before a cost or a geometry can be discussed we need a logical substrate that supplies content to compare. The substrate adopted here is the minimal one used in this paper: the existence of a single non-trivial distinction on some carrier. Every later layer (recognition work, the reciprocal comparison cost, the scale geometry, and the physical identifications) is built on top of this floor; all added mathematical structure and bridge assumptions are introduced and labeled explicitly.

This section and the next begin the paper’s narrative rather than its proof: they situate the reciprocal cost within a recognition account of comparison, and we mark precisely where the proved mathematics detaches from that account. The uniqueness theorem of Part I does not depend on either section: as Remark 6.1 records, the proofs of Sections 6–7 begin only after the positive-ratio ledger hypothesis (SI2) introduces the cost object, and they invoke none of the definitions or theorems of Sections 4–5. A reader interested only in the proved result may proceed directly to Section 6.

*Definition 4.1* (Distinguishability floor). The primitive floor is the proposition that some carrier admits a non-trivial distinction:

$$\exists K, \exists x, y \in K, \quad x \neq y.$$

The floor is not a physical event. It has no time, no position, and no metric. It is the minimal condition for there being a proposition with content. Without at least one distinction, there is no difference between assertion and denial.

**Remark 4.1** (Three properties of the floor). We record three elementary observations about the floor; each is used downstream in a different role.

- (a) *Specification from distinction.* On an inhabited carrier  $K$ , the distinguishability floor and the existence of a non-trivial specification predicate are interchangeable reformulations of the same datum, and we use them interchangeably throughout the paper. Concretely: if  $x \neq y$  are two distinct elements of  $K$ , the predicate  $P(z) :\Leftrightarrow z = x$  is true of  $x$  and false of  $y$ , so a non-trivial distinction supplies a predicate with at least one positive and one negative witness; conversely, if

- $P(a)$  holds and  $P(b)$  does not, then  $a \neq b$ , so a non-trivial predicate supplies a distinction. This equivalence is elementary and is not used as a load-bearing step in any downstream proof; it clarifies the formal language used to phrase the floor.
- (b) *Minimality*. The floor supplies no topology, metric, continuity hypothesis, number system, time parameter, probability measure, or notion of physical state. Those structures enter only when introduced later. At this stage the only conclusion is the existence of a predicate with one positive and one negative witness on  $K$ .
  - (c) *Self-bootstrap*. To deny the floor is to assert a proposition (“no distinction holds”) distinct from its negation. That assertion exhibits one distinction at the meta-language level. The floor is therefore not falsifiable inside the meta-language. This is not a metaphysical comment; it is the observation that the meta-language is itself non-degenerate. This observation is rhetorical context only: it is not used in any proof in this paper.

The scope of this equivalence should not be overstated. The meta-language in which it is phrased already supplies equality, subsets, and predication, so what is recovered is a non-trivial specification predicate within an ambient logic, not a derivation of logic from nothing. The floor is mathematically trivial and deliberately so: it contributes no arithmetic, cost, probability, time, or topology, and no later section borrows more from it than the existence of one two-point distinction.

## 5 Recognition Work

Distinguishability supplies a logical floor but not yet a numerical cost. We now introduce *recognition work*: the operational cost of performing or maintaining a distinction. The point of this section is simple. Once a cost of maintaining distinctions is admitted at all, it cannot be an arbitrary label on inconsistency; independent contradictions must be counted independently. The positive-ratio cost developed in the next section is a strictly later layer that sits on top of this configuration-level accounting.

The section has three jobs. First, it fixes the accounting of recognition work and proves its rigidity: on a specified generator family the cost is a weighted counting measure (Theorem 5.3). Second, it states the bridge from this configuration-level accounting to positive ratios, structural hypothesis (SI2), and reduces that bridge to standard measurement axioms (Proposition 5.5). Third, it records the recognition-quotient boundary that fixes how the word “bit” is used later. Only the second job feeds Part I.

*Definition 5.1* (Predicate vocabulary and configuration). Fix a carrier  $K$  and a syntactic predicate vocabulary  $\mathcal{A}$ . Each predicate symbol  $P \in \mathcal{A}$  has an interpretation

$$\llbracket P \rrbracket := \{z \in K : P(z)\} \subseteq K.$$

Constraints are literals  $P$  or  $\neg P$  with truth sets  $\llbracket P \rrbracket$  and  $K \setminus \llbracket P \rrbracket$ , respectively. A configuration is a finite family of such literals. Its predicate support is the set of predicate symbols from  $\mathcal{A}$  that occur in the family, ignoring negation. A configuration is consistent if the intersection of the truth sets of its literals is non-empty. Supports are syntactic: two different predicate symbols with the same truth set remain distinct unless they have explicitly been identified in the vocabulary.

*Definition 5.2* (Independent configurations and independent join). Two configurations  $\Gamma$  and  $\Delta$  are *independent* if

- (I1) (*syntactic disjointness*) their predicate supports are disjoint in the syntactic sense of Definition 5.1 (no predicate symbol appears in both); and
- (I2) (*semantic non-interaction*) for every pair of subconfigurations  $\Gamma' \subseteq \Gamma$  and  $\Delta' \subseteq \Delta$ , if  $\Gamma'$  and  $\Delta'$  are each consistent then their union  $\Gamma' \cup \Delta'$  is consistent: every consistent choice on the  $\Gamma$  side is jointly satisfiable with every consistent choice on the  $\Delta$  side.

For independent  $\Gamma$  and  $\Delta$ , the *independent join*  $\Gamma \sqcup \Delta$  is the configuration obtained by taking the union of their constraint sets. A sufficient condition for (I2) is a product structure: if  $K \cong K_1 \times K_2$  and every truth set of a literal of  $\Gamma$  is a cylinder  $A \times K_2$  while every truth set of a literal of  $\Delta$  is a cylinder  $K_1 \times B$ , then witnesses combine coordinatewise and (I2) holds.

**Remark 5.1** (Syntactic disjointness alone is not independence). Condition (I2) is not redundant, and an earlier version of this definition that required only disjoint syntactic supports was inconsistent with the cost axioms below. Take  $K = \{0, 1\}$  with two predicate symbols interpreted as  $\llbracket P \rrbracket = \{0\}$  and  $\llbracket Q \rrbracket = \{1\}$ . Then  $\Gamma = \{P\}$  and  $\Delta = \{Q\}$  have disjoint syntactic supports and are each consistent, yet  $\Gamma \cup \Delta$  is inconsistent ( $\llbracket P \rrbracket \cap \llbracket Q \rrbracket = \emptyset$ ). Under the cost axioms below, (W1) would force  $C(\Gamma \cup \Delta) > 0$  while support-disjoint additivity would force  $C(\Gamma \cup \Delta) = C(\Gamma) + C(\Delta) = 0$ ; no cost function could satisfy both. Disjoint predicate names therefore do not imply that the predicates cannot interact semantically; condition (I2) excludes exactly this failure mode, and the pair  $(\{P\}, \{Q\})$  above is not independent in the sense of Definition 5.2.

**Lemma 5.1** (Consistency factorizes over independent joins). *Let  $\Gamma$  and  $\Delta$  be independent in the sense of Definition 5.2. Then  $\Gamma \sqcup \Delta$  is consistent if and only if  $\Gamma$  and  $\Delta$  are each consistent.*

*Proof.* If  $\Gamma$  and  $\Delta$  are each consistent, condition (I2) applied with  $\Gamma' = \Gamma$  and  $\Delta' = \Delta$  gives a common witness for  $\Gamma \cup \Delta = \Gamma \sqcup \Delta$ . Conversely, a witness for  $\Gamma \sqcup \Delta$  lies in the intersection of all the literal truth sets, hence witnesses  $\Gamma$  and  $\Delta$  separately (adding literals only shrinks the witness set).  $\square$

*Definition 5.3* (Recognition-work cost). A recognition-work cost is a map

$$C : \{\text{finite configurations}\} \rightarrow \mathbb{R}_{\geq 0}$$

satisfying:

- (W1)  $C(\Gamma) = 0$  if and only if  $\Gamma$  is consistent;
- (W2)  $C(\Gamma \sqcup \Delta) = C(\Gamma) + C(\Delta)$  when  $\Gamma$  and  $\Delta$  are independent in the sense of Definition 5.2.

Since the codomain is  $\mathbb{R}_{\geq 0}$ , positivity on inconsistent configurations follows from (W1) and is not listed as a separate axiom. By Lemma 5.1, (W1) and (W2) are mutually consistent on independent joins: when both pieces are consistent the join is consistent and both sides of (W2) vanish; when at least one piece is inconsistent the join is inconsistent and both sides are positive.

The independent-additivity condition is the load-bearing structural postulate. It is not derived from the dichotomy in (W1): indicator costs satisfy (W1) but fail

(W2), and superadditive or subadditive costs would also satisfy (W1) but violate (W2). Postulating exact additivity over independent configurations is the operational claim that semantically non-interacting predicates introduce no positive or negative interaction term in the cost. A systematic comparison with nearby alternatives (bare consistency, monotone scores, and sub- or superadditive costs) is given in [9]; the example below records the essential case.

**Example 5.1** (Two independent contradictions). Let  $K = \{0, 1\} \times \{0, 1\}$ , let  $P$  be interpreted as the cylinder  $\llbracket P \rrbracket = \{0\} \times \{0, 1\}$  (a constraint on the first coordinate only), and let  $Q$  be interpreted as  $\llbracket Q \rrbracket = \{0, 1\} \times \{0\}$  (second coordinate only). Set  $\Gamma = \{P, \neg P\}$  and  $\Delta = \{Q, \neg Q\}$ . The supports are syntactically disjoint, and the cylinder structure realizes the semantic non-interaction condition (I2) of Definition 5.2: any consistent subconfiguration of  $\Gamma$  constrains only the first coordinate and any consistent subconfiguration of  $\Delta$  only the second, so witnesses combine coordinatewise. Each of  $\Gamma$  and  $\Delta$  is inconsistent. The independent join  $\Gamma \sqcup \Delta = \{P, \neg P, Q, \neg Q\}$  uses both predicate families and is also inconsistent.

Under recognition work,  $C(\Gamma) > 0$ ,  $C(\Delta) > 0$ , and

$$C(\Gamma \sqcup \Delta) = C(\Gamma) + C(\Delta) > 0.$$

Under the inconsistency indicator  $I$ ,  $I(\Gamma) = I(\Delta) = I(\Gamma \sqcup \Delta) = 1$ , but

$$I(\Gamma \sqcup \Delta) = 1 \neq 2 = I(\Gamma) + I(\Delta).$$

Recognition work counts two failures; the indicator counts only one. The difference is the operational claim that maintaining two independent contradictions is twice as much work as maintaining one.

Recognition work is therefore a genuine quantitative extension: it counts independent failures separately. The number-valued work functional is not present in bare distinguishability. The conclusion should be stated modestly: once the operational cost of maintaining distinctions is admitted, the dichotomy and additivity conditions rule out the degenerate indicator reading; they postulate, rather than derive, the additive accounting whose consequences on a specified generator family are worked out below. A companion treatment passes from coherent ratio comparison to discrete ledger dynamics with atomic ticks and double-entry postings [9].

*Definition 5.4* (Indecomposable inconsistent configuration). An inconsistent configuration  $\Gamma$  is indecomposable if it cannot be written as  $\Gamma = \Gamma_1 \sqcup \Gamma_2$  for two independent inconsistent configurations  $\Gamma_1, \Gamma_2$  in the sense of Definition 5.2. This definition does not imply that removing any single literal makes  $\Gamma$  consistent; indecomposability is the non-splitting condition above.

For instance, the atomic contradiction  $\{P, \neg P\}$  is indecomposable, while the four-literal join  $\{P, \neg P, Q, \neg Q\}$  of Example 5.1 decomposes as  $\{P, \neg P\} \sqcup \{Q, \neg Q\}$  and is not.

*Definition 5.5* (Specified recognition-work generator join-semilattice). Fix a specified label set  $\mathcal{K}$  and a family  $(\Gamma_k)_{k \in \mathcal{K}}$  of labeled indecomposable (Definition 5.4) inconsistent generators that are *jointly independent*: their predicate supports are pairwise disjoint, and for every finite  $S \subseteq \mathcal{K}$  and every choice of consistent subconfigurations  $\Gamma'_k \subseteq \Gamma_k$

( $k \in S$ ), the union  $\bigcup_{k \in S} \Gamma'_k$  is consistent. Joint independence implies that any two disjoint sub-joins of the family are independent in the sense of Definition 5.2, so (W2) applies inductively along any decomposition; it is automatic in the cylinder (product-carrier) realization of Example 5.1 extended to  $|\mathcal{K}|$  factors. The recognition-work generator join-semilattice  $\mathcal{G}$  is the set of finite independent joins of distinct labeled generators:

$$\mathcal{G} = \left\{ \bigsqcup_{k \in S} \Gamma_k : S \subseteq \mathcal{K}, |S| < \infty \right\},$$

with the empty join interpreted as the empty consistent configuration. If several occurrences of the same generator type are needed, they are represented by distinct labels in  $\mathcal{K}$ ; the independent-join operation itself never repeats a generator with the same predicate support. Thus  $\mathcal{G}$  is the free commutative idempotent join-semilattice on the specified labeled generator family, represented concretely by finite subsets of the label set together with their corresponding formal independent joins.

The remainder of this section establishes uniqueness of canonical decompositions on the specified generator join-semilattice (Lemma 5.2), then applies that uniqueness to prove rigidity of recognition-work costs by an explicit counting formula (Theorem 5.3). The rigidity is a theorem about the chosen free join-semilattice, not a derivation of that semilattice from bare distinguishability.

**Lemma 5.2** (Uniqueness of the canonical decomposition). *On the specified generator join-semilattice  $\mathcal{G}$  of Definition 5.5, the support set  $S(\Gamma)$  indexing any independent-join decomposition  $\Gamma = \bigsqcup_{k \in S(\Gamma)} \Gamma_k$  is intrinsic to  $\Gamma$ : if  $\Gamma = \bigsqcup_{k \in S} \Gamma_k = \bigsqcup_{k \in S'} \Gamma_k$ , then  $S = S'$ .*

*Proof.* Each  $\Gamma_k$  is an inconsistent generator, so its underlying constraint set uses at least one predicate because an empty configuration is consistent. Hence the predicate support  $\text{supp}(\Gamma_k)$  is non-empty for every  $k$ . By Definition 5.5, the supports  $\{\text{supp}(\Gamma_k)\}_{k \in \mathcal{K}}$  are pairwise disjoint. The predicate support of  $\bigsqcup_{k \in S} \Gamma_k$  is the disjoint union  $\bigcup_{k \in S} \text{supp}(\Gamma_k)$ . Equality of  $\bigsqcup_{k \in S} \Gamma_k$  and  $\bigsqcup_{k \in S'} \Gamma_k$  as configurations forces equality of these disjoint unions; since the  $\text{supp}(\Gamma_k)$  are non-empty and pairwise disjoint, this forces  $S = S'$ . Explicitly, writing  $U := \bigcup_{k \in S} \text{supp}(\Gamma_k)$ , the index set is recovered as  $S = \{k \in \mathcal{K} : \text{supp}(\Gamma_k) \subseteq U\}$ : for  $k \in S$  the inclusion is immediate, while for  $k \notin S$  the support  $\text{supp}(\Gamma_k)$  is non-empty and, by pairwise disjointness, disjoint from every  $\text{supp}(\Gamma_{k'})$  with  $k' \in S$ , hence disjoint from  $U$  and not contained in it. Thus  $U$  determines  $S$ , and  $U = \bigcup_{k \in S'} \text{supp}(\Gamma_k)$  likewise determines  $S'$ , so  $S = S'$ .  $\square$

With uniqueness of decompositions in hand, the additive formula  $C(\Gamma) = \sum_{k \in S(\Gamma)} c_k$  used in the theorem below is well-defined: any two representations of the same  $\Gamma$  as an independent join of labeled generators agree on the index set  $S(\Gamma)$  and hence on the sum.

**Theorem 5.3** (Recognition-work costs on a specified generator family). *Fix the recognition-work generator join-semilattice  $\mathcal{G}$  of Definition 5.5. Let  $C$  satisfy (W1) and (W2) on configurations in  $\mathcal{G}$ . Then every  $\Gamma \in \mathcal{G}$  has a unique canonical independent-join decomposition  $\Gamma = \bigsqcup_{k \in S(\Gamma)} \Gamma_k$  with  $S(\Gamma) \subseteq \mathcal{K}$  finite (unique by Lemma 5.2),*

and

$$C(\Gamma) = \sum_{k \in S(\Gamma)} c_k, \quad c_k := C(\Gamma_k) > 0.$$

In particular, any two recognition-work costs that agree on the generators agree on all of  $\mathcal{G}$ .

*Proof.* By (W1),  $C$  vanishes on consistent configurations, and since the codomain is  $\mathbb{R}_{\geq 0}$ , every inconsistent generator has strictly positive cost. By (W2),  $C$  is additive over independent joins, and induction extends this to any finite independent decomposition in  $\mathcal{G}$ . By Lemma 5.2, the support set  $S(\Gamma)$  is uniquely determined by  $\Gamma$ . Repeated additivity then gives the well-defined formula  $C(\Gamma) = \sum_{k \in S(\Gamma)} C(\Gamma_k)$ . Setting  $c_k = C(\Gamma_k)$  proves the displayed formula. Positivity  $c_k > 0$  on every generator is the (W1) clause for inconsistent configurations, so any two recognition-work costs agreeing on the generator family  $\{\Gamma_k\}_{k \in \mathcal{K}}$  agree on all of  $\mathcal{G}$ .  $\square$

Thus the formal result has a simple interpretation: on the specified generator join-semilattice, recognition work is a weighted counting measure, and the positive weights  $c_k = C(\Gamma_k)$  are the only freedom once the generator family is fixed.

**Remark 5.2** (Non-vacuity of the cost axioms). The class of recognition-work costs on  $\mathcal{G}$  is non-empty. Realize the generator family on a product carrier  $K = \prod_{k \in \mathcal{K}} K_k$  with each  $\Gamma_k$  constraining only the  $k$ th coordinate through cylinder truth sets, as in Example 5.1; joint independence then holds by coordinatewise combination of witnesses. For any choice of weights  $c_k > 0$ , the weighted counting measure  $C(\Gamma) := \sum_{k \in S(\Gamma)} c_k$  satisfies (W1) on  $\mathcal{G}$  (the empty join is consistent with cost 0; every non-empty join contains an inconsistent generator and has positive cost) and (W2) by construction. The axioms are therefore consistent on the specified domain; on unrestricted vocabularies without the semantic non-interaction condition they would not be (Remark 5.1).

This section does not yet derive the reciprocal cost  $J$ . It derives the existence and rigidity of a cost framework once recognition work and its two operational constraints are admitted. The structural results above are elementary: they amount to the universal property of a free commutative idempotent join-semilattice together with additivity. Their value is in making explicit exactly how much is postulated before any ratio structure appears. The specific positive-ratio cost is the next layer.

Independent additivity is the operational content of saying that two independent distinctions have both been made: if each failure requires work to distinguish, performing both cannot cost the same as performing one unless the second distinction is operationally free, which is the indicator reading just rejected. The condition is also deliberately weaker than a full measure structure. It is imposed only on independent joins; nothing in the paper assumes additivity for overlapping predicate families, where correlations or shared witnesses may change the cost.

**Proposition 5.4** (Unit freedom in recognition work). *If  $C$  is a recognition-work cost and  $\mu > 0$ , then  $\mu C$  is also a recognition-work cost. Thus recognition work determines a cost scale only up to a positive unit calibration before the positive-ratio functional equation is imposed.*

*Proof.* Zero sets are unchanged by multiplication by  $\mu > 0$ , positivity on inconsistent configurations is preserved, and independent additivity is preserved because

$$(\mu C)(\Gamma \sqcup \Delta) = \mu C(\Gamma \sqcup \Delta) = \mu C(\Gamma) + \mu C(\Delta).$$

□

This unit freedom is expected. A cost framework before calibration is like a length geometry before choosing meters. Later, the unit log-curvature calibration fixes the positive-ratio scale. For a comparison of recognition work with the nearby alternatives (no cost, the indicator, arbitrary monotone scores, and sub- or superadditive costs, each of which either loses compositional content or claims an interaction term with no referent on the predicate side) we refer to [9]; recognition work sits between Boolean consistency and arbitrary scoring, as the constrained passage from predicate structure to positive-ratio cost.

## 5.1 From recognition work to ratio comparison

Recognition work in Section 5 assigns costs to finite configurations of predicates. Ratio comparison in Section 6 assigns costs to positive ratios  $x = a/b$  between comparable intensities. The passage is structural hypothesis (SI2); it is not derived from independent additivity on configurations alone.

One operational reading of SI2 is that a recognition entry carries many discrete independent specification predicates  $N_a$  and  $N_b$  representing two descriptions of a single quantity. If the cost of maintaining these specifications scales linearly according to Theorem 5.3, the positive-ratio ledger hypothesis defines intensities such as  $a$  and  $b$  and tracks the dimensionless mismatch by  $x = a/b \in \mathbb{R}_{>0}$ . This reading is heuristic; SI2 is the bridge assumption that introduces positive intensities and ratios. The functional behavior of  $F(x)$  then requires an independent route-independence postulate on how these ratios compose. Configurations that only assert  $a \neq b$  without magnitudes remain at the recognition-work layer. Independent additivity over semantically independent configurations does not, by itself, force

$$F(xy) + F(x/y) = \Phi(F(x), F(y)).$$

That law is an additional route-independence hypothesis on ratio composition: the aligned product  $xy$  and the anti-aligned ratio  $x/y$  are the two canonical two-step assemblies, and their total cost must depend only on  $F(x)$  and  $F(y)$ . The mapping from discrete ledger dynamics to coherent ratio comparison is developed in [9]; here we state only the ratio-level inputs needed for the reciprocal composition law.

**Proposition 5.5** (Hölder representation of the comparison group). *Structural hypothesis (SI2) need not be posited in the bare form “comparisons are recorded as positive ratios  $x = a/b \in \mathbb{R}_{>0}$ ”; its content splits into standard axioms of extensive measurement. Suppose the comparable magnitudes, taken up to the comparisons between them, carry: an associative and commutative composition of comparisons; a neutral element (the null comparison of a magnitude with itself); an inverse for each comparison (the same comparison read in the opposite direction); and a linear order,*

translation-invariant under composition, that is Archimedean (for any two positive comparisons  $g, h > 0$ , some  $n \in \mathbb{N}$  satisfies  $ng > h$ ). Then the comparison structure is an Archimedean linearly ordered abelian group  $G$ , and by Hölder's theorem [3, 4] it embeds order-isomorphically, uniquely up to one positive scale factor, into  $(\mathbb{R}, +)$ . Writing  $t$  for the image of a comparison and  $x := e^t$ , the comparisons are realized as a subgroup of the multiplicative group  $(\mathbb{R}_{>0}, \cdot)$ ; not, in general, all of  $\mathbb{R}_{>0}$ ; with  $x = 1$  the neutral comparison,  $x \mapsto x^{-1}$  the inverse, and reciprocity (C2) the evenness of the cost under  $t \mapsto -t$ . If, in addition,  $G$  is non-trivial and its order is dense (strictly between any two comparisons lies a third) and Dedekind complete (every non-empty order-bounded set of comparisons has a least upper bound), then the embedding is onto and the carrier is exactly  $\mathbb{R}_{>0}$ .

*Proof.* The four listed properties are the axioms of an Archimedean linearly ordered abelian group. Hölder's theorem gives an order-embedding  $\iota : G \rightarrow (\mathbb{R}, +)$ , unique up to a positive multiplicative constant; transporting along  $t \mapsto e^t$  realizes  $G$  as a subgroup of  $(\mathbb{R}_{>0}, \cdot)$  with neutral element 1 and inverse  $x \mapsto x^{-1}$ , and a cost that does not distinguish a comparison from its reverse is invariant under  $t \mapsto -t$ , i.e. satisfies  $F(x) = F(x^{-1})$ . Hölder's theorem asserts an embedding, not an isomorphism onto  $\mathbb{R}$ : the image can be a proper subgroup (for example the cyclic ladder generated by a single comparison).

For the onto clause, suppose  $G$  is non-trivial, densely ordered, and Dedekind complete. A non-trivial subgroup of  $(\mathbb{R}, +)$  is either cyclic or dense in  $\mathbb{R}$ ; a cyclic image would make the order of  $G$  discrete, contradicting density, so  $\iota(G)$  is dense in  $\mathbb{R}$ . Fix  $r \in \mathbb{R}$ . The set  $S := \{g \in G : \iota(g) < r\}$  is non-empty and bounded above (both by density of  $\iota(G)$  in  $\mathbb{R}$ ), so by Dedekind completeness it has a supremum  $s \in G$ . If  $\iota(s) < r$ , density supplies  $g \in G$  with  $\iota(s) < \iota(g) < r$ , so  $g \in S$  exceeds  $s$ , a contradiction; if  $\iota(s) > r$ , density supplies  $g$  with  $r < \iota(g) < \iota(s)$ , and this  $g$  is an upper bound of  $S$  smaller than  $s$ , again a contradiction. Hence  $\iota(s) = r$ , so  $\iota(G) = \mathbb{R}$  and the carrier is exactly  $\mathbb{R}_{>0}$ .  $\square$

**Remark 5.3** (What Proposition 5.5 does and does not settle). This reduces (SI2) to named measurement axioms rather than a bespoke postulate naming  $\mathbb{R}_{>0}$  directly, and it locates exactly where the continuum enters. The ordered-group axioms force the carrier to be order-isomorphic to some subgroup of  $\mathbb{R}_{>0}$  (possibly a proper one, such as a discrete cyclic ladder), so what they exclude is any comparison structure not representable inside the positive reals; only the two further axioms of order density and Dedekind completeness force the full continuum carrier  $\mathbb{R}_{>0}$  itself. The continuum clause is load-bearing for Part I: without density and completeness the carrier may be a proper subgroup, for example a discrete cyclic carrier  $\{r^n : n \in \mathbb{Z}\}$  for some  $r > 1$ , on which the continuity hypothesis (C4) of Theorem 7.1 has no force; the theorem characterizes  $J$  on the continuum carrier only. What remains open is one level deeper: a derivation, from the configuration-level recognition work of Section 5 alone, that comparable magnitudes carry this ordered, dense, complete comparison structure in the first place. The recognition-work functional of Theorem 5.3 is an additive counting measure and does not by itself supply a translation-invariant Archimedean order, let alone density or completeness; supplying these remains open.

**Example 5.2** (Minimal two-entry ledger). Suppose one ledger carries ratio  $x = a/b$  and a second carries  $y = c/d$  on independent entry types. Composing the comparisons in the aligned orientation gives overall mismatch  $xy$ ; composing with one orientation reversed gives  $x/y$ . Route independence requires

$$F(xy) + F(x/y) = \Phi(F(x), F(y))$$

for some combiner  $\Phi$ . For  $x = 2$ ,  $y = 3$ , the left side is  $F(6) + F(2/3)$ ; the right side depends only on  $F(2)$  and  $F(3)$ . No configuration-level sum of recognition-work costs produces this constraint without the ratio hypothesis and the two-channel postulate.

## 5.2 Recognition quotient and gauge boundary

The public recognition-core files supply an important boundary condition. A single Boolean recognition coordinate is an atomic split, but it is not a complete code for an arbitrary state space. Let  $\mathcal{F}$  be a family of recognizers  $f : X \rightarrow C$ . The recognition equivalence relation is

$$x \sim_{\mathcal{F}} y \iff \forall f \in \mathcal{F}, f(x) = f(y).$$

The physical quotient is  $X/\sim_{\mathcal{F}}$ , and gauge freedom means indistinguishability by the selected recognition family.

**Remark 5.4** (Repository recognition boundary). The public recognition layer of the repository proves, in machine-checked form, the following boundary statements.

- (i) Two states project to the same physical quotient exactly when all selected recognizers agree.
- (ii) Gauge freedom means that no available recognition act distinguishes the two representatives.
- (iii) A recognition signature is complete exactly when the recognizer family separates points.
- (iv) One Boolean coordinate is not, by itself, a complete code for a multi-coordinate state.
- (v) Adding recognizers refines the quotient.

Thus later uses of a “bit” must be read as uses of an atomic recognition act. A complete physical description requires a separating family of recognizers or a bridge hypothesis asserting that the chosen signature is complete for the sector under discussion. The sole later use of this vocabulary is the eight-tick cycle of Section 9, whose “recognizer family of  $D$  bits” names  $D$  atomic recognition acts in the sense fixed here, not a claim about quantum information.

## 6 The Reciprocal Composition Law

Section 5 established that recognition work on finite configurations is a positive, additive counting measure over indecomposable inconsistencies (Theorem 5.3). The freedom that remains is the choice of positive generator weights. Section 5.1 states how positive intensities and ratios are introduced; that step is structural hypothesis (SI2). Once it

is accepted, the cost of a comparison is a function on positive ratios  $x > 0$ , the identity ratio  $x = 1$  carries zero cost, and the reciprocal ratio  $x^{-1}$  carries the same cost as  $x$ . The question that drives this section becomes: what function on  $\mathbb{R}_{>0}$  is singled out by requiring that two-step comparisons composed in both orientations depend only on the costs of the individual steps?

The section proceeds in three moves. First, it defines a reciprocal comparison cost through the conditions (C1)–(C7) and records what each condition excludes. Second, it reduces the composition combiner to a one-parameter bilinear form (Proposition 6.2). Third, it recasts the closure condition carrying that reduction as a multiplicative two-channel composition law (Theorem 6.3) and proves it independent of the remaining axioms (Proposition 6.4). The uniqueness theorem itself is assembled in Section 7.

**Remark 6.1** (Mathematical independence from Section 5). The proofs in this section and in Section 7 do not invoke Definitions 5.1–5.5 or Theorem 5.3 from Section 5. The bridge between the two layers is structural hypothesis (SI2) (the positive-ratio ledger hypothesis), which introduces a cost  $F : \mathbb{R}_{>0} \rightarrow \mathbb{R}$  as a formally independent object. Without (SI2), the recognition-work results of Section 5 supply no constraints on the functional form of  $F$ .

**Remark 6.2** (Necessity certificates). The independence just stated is complemented by two necessity results, which identify exactly what must be added to the floor to reach  $J$ . First, kernel-checked at the pinned commit (Certificate 8.2): the floor’s own equality-derived cost provably fails the multiplicative composition consistency that the reciprocal-cost theorem requires. Thus the canonical equality-induced cost does not itself supply the analytic composition layer; the multiplicative carrier structure introduced by (SI2) is an additional input on this route. Second, proved below (Proposition 6.4): the quartic-log cost witnesses that (C1)–(C5) do not imply finite polynomial closure (C6). This is the precise sense in which the hypothesis set is tight: any claim that the chain from distinction to  $J$  has been completed must first dislodge these necessity and independence witnesses.

*Definition 6.1* (Finite recognition ledger). A finite recognition ledger is a finite family of nonzero recognition work entries

$$\mathcal{L} = \{(a_i, b_i)\}_{i=1}^N, \quad a_i, b_i \in \mathbb{R}_{>0},$$

where  $a_i$  and  $b_i$  are two positive intensities assigned to the same kind of recognition entry under two compatible descriptions. The ledger records only the dimensionless ratios

$$x_i := a_i/b_i \in \mathbb{R}_{>0}.$$

The identity comparison is  $x_i = 1$ .

**Remark 6.3** (The ledger hypothesis). Definition 6.1 is not derived from the predicate algebra alone. The recognition-work results in Section 5 show how independent finite configurations carry additive cost. The additional hypothesis here is that comparable ledger entries carry positive intensities and that their mismatch is represented by the ratio of those intensities. Once this hypothesis is accepted, a ratio  $x > 0$  measures mismatch between two comparable magnitudes, and the identity ratio is  $x = 1$ . The passage from coherent comparison on ratios to a full discrete ledger formalism is carried out in [9].

*Definition 6.2* (Reciprocal comparison cost). A *reciprocal comparison cost* is a function  $F : \mathbb{R}_{>0} \rightarrow \mathbb{R}$  satisfying:

- (C1) normalization:  $F(1) = 0$ ;
- (C2) reciprocity:  $F(x) = F(x^{-1})$ ;
- (C3) route independence: there exists a function  $\Phi : F(\mathbb{R}_{>0}) \times F(\mathbb{R}_{>0}) \rightarrow \mathbb{R}$  with

$$F(xy) + F(x/y) = \Phi(F(x), F(y)) \quad \text{for all } x, y \in \mathbb{R}_{>0};$$

the well-definedness of  $\Phi$  as a single-valued function on the cost image is spelled out in Remark 6.4 immediately after the definition;

- (C4) regularity:  $F$  is continuous on  $\mathbb{R}_{>0}$ ;
- (C5) comparison-cost nondegeneracy (structural hypothesis (SI11)):  $F(x) \geq 0$  for all  $x \in \mathbb{R}_{>0}$ ,  $F(x) = 0$  if and only if  $x = 1$ , and  $F(e^t) \rightarrow +\infty$  as  $|t| \rightarrow \infty$ .

Conditions (C1)–(C4) record the recognition-composition data: normalization at the identity, reciprocal symmetry, route independence, and continuity. Condition (C5) makes the comparison-cost interpretation explicit as a labeled hypothesis: it forbids the cost from being identically zero (no comparison content), from vanishing on more than the identity ratio (no unique minimum), or from remaining bounded as  $|\log x| \rightarrow \infty$  (extreme mismatch costs fail to grow without bound). The d’Alembert equation derived below has continuous real solution branches that satisfy (C1)–(C4) but fail (C5); those branches are recorded as alternatives rather than counted as reciprocal comparison costs. A reciprocal comparison cost is *calibrated* if, in addition,

- (C6) finite pairwise polynomial closure with non-degenerate bilinear branch ((SI3)): the combiner  $\Phi$  is the restriction to  $F(\mathbb{R}_{>0}) \times F(\mathbb{R}_{>0})$  of a symmetric real polynomial in two variables of total degree at most two, and the bilinear coefficient satisfies  $c \neq 0$ . The sign of  $c$  is not constrained by (C6): the unit convention (C7a) below fixes only the magnitude  $|c| = 2$ , and comparison-cost nondegeneracy (C5) then excludes the negative branch  $c = -2$  via Remark 6.7, leaving the positive branch  $c = 2$ .
- (C7) canonical calibration, consisting of two logically distinct steps:
  - (C7a) *Bilinear unit convention*. For either nonzero bilinear branch, rescale  $F$  by the positive factor  $|c|/2$ . The normalized representative then has mixed coefficient  $2 \operatorname{sgn}(c) \in \{-2, 2\}$ . This fixes the magnitude of the coefficient without selecting its sign: the positive branch has  $c = 2$ , while the negative branch has  $c = -2$ .
  - (C7b) *Unit log-curvature condition*. With the unit fixed by (C7a), require

$$\lim_{t \rightarrow 0} \frac{F(e^t)}{t^2/2} = 1,$$

equivalently  $F(e^t) = \frac{1}{2}t^2 + o(t^2)$  in logarithmic coordinates. Within the positive hyperbolic family retained by (C5), this fixes  $\lambda = 1$ . This condition is *not* a consequence of (C1)–(C7a): it selects one specific curvature profile from the one-parameter family allowed on the retained branch. The limit is well posed on every continuous branch classified under (C1)–(C6). In

the sign-neutral unit  $|c| = 2$ , the nonconstant representatives have log-coordinate forms  $\pm(\cosh(\lambda t) - 1)$  or  $\pm(1 - \cos(\lambda t))$ , with limit  $\pm\lambda^2$ ; the constant branch has limit 0. On the positive hyperbolic branch, (C7b) selects  $\lambda = 1$ . Thus (C7b) is a calibration condition on the representative called calibrated, not an independent existence-of-limit assumption on an arbitrary  $F$ .

A function  $F : \mathbb{R}_{>0} \rightarrow \mathbb{R}$  is *calibrated* if it satisfies (C1)–(C5), (C6), (C7a), and (C7b); equivalently, since (C5) forces  $F \geq 0$ , the codomain may be taken to be  $\mathbb{R}_{\geq 0}$ .

**Remark 6.4** (Well-definedness of  $\Phi$  in (C3)). Condition (C3) requires that  $\Phi$  be a function rather than a multi-valued relation. This means: whenever  $(F(x_1), F(y_1)) = (F(x_2), F(y_2))$ , the corresponding right-hand sides agree,  $F(x_1 y_1) + F(x_1 / y_1) = F(x_2 y_2) + F(x_2 / y_2)$ . Reciprocity (C2) automatically forces this on the canonical level sets  $\{x, x^{-1}\} \times \{y, y^{-1}\}$ , since both inversions only exchange the two summands  $F(xy)$  and  $F(x/y)$ ; on any larger level sets that a particular  $F$  may carry, (C3) imposes the same constancy as an additional condition.

**Example 6.1** (A reciprocal cost for which (C3) fails). Condition (C3) genuinely excludes functions, and the mechanism is exactly the larger-level-set constancy of Remark 6.4. Let  $g(t) := t^2(|t| - 2)^2$  and  $F(x) := g(\log x)$ . Then  $F$  is continuous,  $F(1) = 0$ , and  $F(x) = F(x^{-1})$  since  $g$  is even, so (C1), (C2), and (C4) hold. But  $F$  has the non-reciprocal level coincidence  $F(1) = F(e^2) = 0$ , and route independence fails on it: taking  $y = e$  (so  $F(y) = g(1) = 1$  in both cases),

$$F(e^{0+1}) + F(e^{0-1}) = 2g(1) = 2, \quad F(e^{2+1}) + F(e^{2-1}) = g(3) + g(1) = 9 + 1 = 10,$$

while the would-be arguments  $(F(x), F(y)) = (0, 1)$  agree. No single-valued  $\Phi$  exists, so  $F$  is not a reciprocal comparison cost. (C3) is therefore a substantive restriction even before the polynomial closure (C6) is imposed.


The route-independence condition (C3), in its most primitive form, requires that the cost of composing two reciprocal ratios depend only on the costs of the two inputs. The closure condition (C6) and calibration (C7a)–(C7b) are the additional structural hypotheses needed to extract a unique cost from this law, and the comparison-cost nondegeneracy (C5) is the hypothesis that selects the hyperbolic branch among the d’Alembert solutions. After the bilinear reduction, the proof uses the rescaling  $\hat{F} = (c/2)F$  to display the canonical unit; in the definition above, however, the curvature condition is stated directly on the cost representative being called calibrated.

## 6.1 Route independence and the composition law

A positive-ratio comparison has two canonical two-step assemblies. If one comparison has ratio  $x$  and a second has ratio  $y$ , the aligned assembly has ratio  $xy$  and the anti-aligned assembly has ratio  $x/y$ . The Recognition Composition Law is the additional route-independence assumption that a label-blind comparison cost depends on this two-channel composite only through the costs of the two inputs. Under that assumption there is a two-variable combiner  $\Phi$  with

$$F(xy) + F(x/y) = \Phi(F(x), F(y)).$$

Condition (C3) includes the well-definedness of  $\Phi$  on the cost image. Reciprocity guarantees the required invariance under  $x \leftrightarrow x^{-1}$  and  $y \leftrightarrow y^{-1}$ , since those inversions only exchange the two terms on the left. A general reciprocal cost may have larger level sets than the inversion pairs; when it does, route independence requires  $F(xy) + F(x/y)$  to be constant on those larger level sets as well. For the calibrated solution  $J$  the level sets are exactly the reciprocal pairs. Outside the image  $F(\mathbb{R}_{>0}) \times F(\mathbb{R}_{>0})$ ,  $\Phi$  is unconstrained; only its restriction to that image enters the functional equation.



**Fig. 1** Two-channel composition of ratio comparisons. Aligned composition of steps  $x$  and  $y$  yields  $xy$ ; anti-aligned composition yields  $x/y$ . Route independence identifies the total channel cost  $F(xy) + F(x/y)$  with a combiner  $\Phi(F(x), F(y))$  depending only on the step costs.

Figure 1 displays the assembly. The two compositions  $xy$  (aligned) and  $x/y$  (anti-aligned) are the two products of  $x$  and  $y$  that depend on the orientation of  $y$  alone, not on the order in which the two steps are applied. Choosing this two-channel pair as the canonical pairwise composition is a structural assumption: it expresses the requirement that a label-blind comparison cost be sensitive to relative orientation ( $y$  versus  $y^{-1}$ ) but symmetric under the exchange of step labels. Higher-order or asymmetric compositions (e.g.  $\{xy, x/y^2\}$ ) would model different ledger conventions and lie outside the two-channel hypothesis adopted here. Route independence is then the statement that the sum  $F(xy) + F(x/y)$  (the total cost of running both orientations) is a function of the individual step costs alone. The diagram makes explicit what Example 5.2 checks algebraically: the left side can depend on  $x$  and  $y$  only through  $F(x)$  and  $F(y)$ , not through separate coordinate labels.

The published d’Alembert-inevitability theorem proves the rigidity of this form under continuity, non-constancy, identity normalization, and a polynomial combiner of degree at most two [8]. The operator reading used here is identity, symmetry, totality, and composition consistency on positive ratios. For completeness we state the required polynomial form here.

**Remark 6.5** (Symmetry of  $\Phi$ ). The combiner  $\Phi$  in (C3) is automatically symmetric in its two arguments. Indeed,  $F(xy) = F(yx)$  trivially, while reciprocity (C2) gives  $F(x/y) = F((x/y)^{-1}) = F(y/x)$ . Hence  $F(yx) + F(y/x) = F(xy) + F(x/y)$ , so  $\Phi(F(y), F(x)) = \Phi(F(x), F(y))$  for all  $x, y \in \mathbb{R}_{>0}$ . The symmetric polynomial ansatz used in Proposition 6.2 is therefore not an additional hypothesis beyond (C2)–(C3); it is a consequence.

**Lemma 6.1** (Identity boundary from normalization). *Let  $F : \mathbb{R}_{>0} \rightarrow \mathbb{R}$  satisfy  $F(1) = 0$  and the route-independence relation  $F(xy) + F(x/y) = \Phi(F(x), F(y))$  for some two-variable  $\Phi$ . Then*

$$\Phi(u, 0) = 2u \quad \text{for every } u \text{ in the range of } F.$$

*Proof.* Set  $y = 1$ . Then  $xy = x$  and  $x/y = x$ , so

$$F(x) + F(x) = \Phi(F(x), F(1)).$$

Since  $F(1) = 0$ , this gives  $\Phi(F(x), 0) = 2F(x)$ . Letting  $u = F(x)$  and allowing  $x$  to range over  $\mathbb{R}_{>0}$  gives the claim on the range of  $F$ .  $\square$

**Proposition 6.2** (Finite pairwise polynomial closure). *Let  $F : \mathbb{R}_{>0} \rightarrow \mathbb{R}$  be continuous and non-constant with  $F(1) = 0$ , and assume route independence (C3) holds with combiner  $\Phi$ . (Continuity and non-constancy are used only to guarantee that  $\text{Range}(F)$  contains at least three distinct values; the conclusion holds verbatim under that weaker hypothesis, which is the form invoked by the measurable variant Corollary 7.2.) If  $\Phi$  extends to a symmetric real polynomial  $\tilde{\Phi} : \mathbb{R}^2 \rightarrow \mathbb{R}$  of total degree at most two such that  $\tilde{\Phi}(u, v) = \Phi(u, v)$  on  $F(\mathbb{R}_{>0}) \times F(\mathbb{R}_{>0})$ , then the restriction satisfies*

$$\Phi(u, v) = 2u + 2v + cuv$$

on  $F(\mathbb{R}_{>0}) \times F(\mathbb{R}_{>0})$ , for some constant  $c \in \mathbb{R}$ . The polynomial-closure conclusion alone does not fix the sign of  $c$ : non-degeneracy is the additional condition  $c \neq 0$ , and the positive bilinear branch is the further branch choice  $c > 0$ .

*Proof.* Write the symmetric polynomial extension as

$$\tilde{\Phi}(u, v) = a + b(u + v) + d(u^2 + v^2) + cuv,$$

where  $a, b, c, d \in \mathbb{R}$  (a symmetric real polynomial of total degree at most two is a linear combination of  $1$ ,  $u + v$ ,  $u^2 + v^2$ , and  $uv$ ). Lemma 6.1 gives  $\tilde{\Phi}(u, 0) = 2u$  for every  $u \in \text{Range}(F)$ , hence

$$a + bu + du^2 = 2u \quad \text{for every } u \in \text{Range}(F).$$

Since  $\mathbb{R}_{>0}$  is connected and  $F$  is continuous, the image  $F(\mathbb{R}_{>0})$  is a connected subset of  $\mathbb{R}$ , hence an interval. Because  $F$  is non-constant and  $F(1) = 0$ , this interval is non-degenerate and contains 0. The polynomial  $p(u) := a + bu + du^2 - 2u$  vanishes at every point of the non-degenerate interval  $F(\mathbb{R}_{>0})$ , hence has infinitely many roots; a polynomial of degree at most two with more than two roots is identically zero, forcing  $a = 0$ ,  $b = 2$ ,  $d = 0$ . The combiner reduces to  $\Phi(u, v) = 2u + 2v + cuv$  on  $F(\mathbb{R}_{>0}) \times F(\mathbb{R}_{>0})$ . Only three distinct points of  $\text{Range}(F)$  are in fact used: a degree- $\leq 2$  polynomial vanishing at three distinct values is identically zero. Continuity enters solely to supply them (the connected image is a non-degenerate interval, hence infinite); the

measurable variant Corollary 7.2 instead supplies three distinct values from  $0 = F(1)$  together with the unbounded-growth clause of (C5), so the polynomial reduction holds without continuity. This proves only the polynomial form and leaves  $c \in \mathbb{R}$ . If the non-degeneracy clause of (C6) is imposed, then  $c \neq 0$ . If, in addition, comparison-cost nondegeneracy (C5) is imposed, the negative branch  $c < 0$  is excluded by Remark 6.7. The remaining positive branch  $c > 0$  can then be put in the canonical mixed-coefficient unit  $c = 2$  by the calibration convention of Section 6.3.  $\square$

The status of (C6) is settled by two results proved later in this section: it is equivalent to a multiplicative two-channel composition law (Theorem 6.3), and it is logically independent of (C1)–(C5) (Proposition 6.4). The degree-two polynomial form above is the computational face of that composition principle, not an independent stipulation; the remarks and calibration that follow may be read as motivation for adopting it.

The next three remarks record, in turn, the sense in which the polynomial extension of  $\Phi$  is unique, how (C5) excludes the  $c < 0$  branch, and the role of the two named hypotheses in the proof of Proposition 6.2.

**Remark 6.6** (A posteriori uniqueness of the polynomial extension). Hypothesis (C6) of Definition 6.2 requires only the *existence* of a symmetric polynomial extension of  $\Phi$  to  $\mathbb{R}^2$  of total degree at most two; uniqueness is not part of the hypothesis. The proof of Proposition 6.2 shows that any such extension is forced to take the form

$$\tilde{\Phi}(u, v) = 2u + 2v + cuv \quad \text{on } F(\mathbb{R}_{>0}) \times F(\mathbb{R}_{>0})$$

with a single real parameter  $c \in \mathbb{R}$ . For a fixed non-constant  $F$  and fixed combiner  $\Phi$ , that polynomial extension is already unique: two extensions can differ only in the coefficient of  $uv$ , and evaluation at any nonzero  $u \in F(\mathbb{R}_{>0})$  gives  $(c - c')u^2 = 0$ . What remains unfixed across admissible costs is the value and sign of  $c$ . The non-degeneracy clause of (C6) removes the additive case  $c = 0$ ; the sign-neutral unit convention (C7a) normalizes either remaining branch to  $c = \pm 2$ ; and comparison-cost nondegeneracy (C5) excludes  $c = -2$ , leaving  $c = 2$ .

**Remark 6.7** (Exclusion of the  $c < 0$  branch). The case  $c < 0$  admits, after the unified rescaling  $\hat{F} = (|c|/2)F$  (the sign-neutral unit convention of (C7a), agreeing with (1) on  $c > 0$ ), the canonical equation

$$\hat{F}(xy) + \hat{F}(x/y) = 2\hat{F}(x) + 2\hat{F}(y) - 2\hat{F}(x)\hat{F}(y),$$

obtained by substituting  $F = (2/|c|)\hat{F}$  and  $c = -|c|$  and clearing the common factor  $2/|c|$ . Setting  $H = 1 - \hat{F}$  converts this to the d'Alembert equation  $H(s+t) + H(s-t) = 2H(s)H(t)$  in log coordinates, with  $H(0) = 1$  and  $H$  even. The classification of continuous solutions of d'Alembert's equation invoked here (constant, cosine, hyperbolic cosine) is the classical one, proved independently of any sign bookkeeping in Appendix A and Stage (iii) of Theorem 7.1; this remark and Table 3 only apply (C5) to the three already-classified branches, so no circularity arises: the order of dependence is d'Alembert classification first, the (C7a) magnitude normalization  $|c| = 2$  second, and (C5) exclusion of the  $c = -2$  branch third. Branch-by-branch verification that

none of the resulting three continuous solutions satisfies (C5) is given by the bottom three rows of Table 3. The  $c < 0$  combiner branch is excluded by (C5).

**Remark 6.8** (Role of finite polynomial closure). The two named hypotheses in Proposition 6.2 are (i) symmetric polynomial form of total degree at most two and (ii) non-degeneracy of the mixed term. Dropping (i) admits continuous combiners that are not finite pairwise polynomials. Dropping (ii) collapses to the additive branch with  $c = 0$ , whose unique unit-log-curvature-normalized solution is  $F(x) = \frac{1}{2}(\log x)^2$  rather than  $J$ . The paper takes (i) as a named closure axiom (structural hypothesis (SI3)) and (ii) as the structural choice of the reciprocal branch. This closure is not derived from the distinguishability floor.

The restriction to a polynomial combiner of degree at most two (structural hypothesis (SI3)) can be given a physical reading. We offer it as motivation only, not as justification. If  $\Phi(u, v)$  were purely linear ( $\Phi = 2u + 2v$ ), the two channels would be completely decoupled, and the cost geometry would be logarithmic,  $F(x) \propto (\log x)^2$ . The bilinear term  $cuv$  is the lowest-order interaction: the cost of validating channel  $x$  is modulated by the validation status of channel  $y$ . Truncating at degree two says that multi-channel interactions are pairwise, and that higher-order non-linearities (e.g.  $u^2v^2$ ) vanish in the asymptotic limit of thin ledgers. The analogy with local field theories of bounded polynomial degree, invoked below, is likewise suggestive rather than probative: many pairwise physical laws are non-polynomial or contain higher powers, so nothing in this reading derives the degree bound. The hypothesis stands or falls as a named axiom. A sharper, unit-free statement of the same hypothesis (that a nonzero rescaling  $1 + \nu F$  composes multiplicatively across the two channels) is given in Remark 6.9 and Theorem 6.3; that is the form in which we recommend (C6) be assessed.

Structural hypothesis (SI3) requires that the cost of a pairwise composite close after finitely many pairwise interactions (constant, linear, and bilinear terms in the cost variables), in the same spirit as a local field theory whose action density is a polynomial of bounded degree at a point. Here the “fields” are  $F(x)$  and  $F(y)$ , and the bound is total degree two, excluding genuinely higher-degree or non-polynomial combiners. Example 6.2 gives a reciprocal cost  $F_4(x) = (\log x)^4$  that is continuous and symmetric under  $x \mapsto x^{-1}$  and satisfies route independence, yet induces a combiner with a  $12\sqrt{ab}$  term, not a degree-two polynomial in  $a = F_4(x)$  and  $b = F_4(y)$ . The present framework does not adopt  $F_4$ ; a rival program that keeps route independence but drops finite pairwise polynomial closure would live in that larger class.

This proposition is where a broad class of continuous comparison algebras is narrowed to the Recognition Composition Law family. Without finite pairwise polynomial closure, more general continuous combiners exist; with that closure, the bilinear reciprocal family is forced.

## 6.2 Examples outside finite pairwise closure

Section 6.1 and Fig. 1 already fixed the aligned and anti-aligned assemblies  $xy$  and  $x/y$  and the route-independence ansatz  $F(xy) + F(x/y) = \Phi(F(x), F(y))$ . The examples below show continuous reciprocal costs that satisfy route independence yet violate finite pairwise polynomial closure (structural hypothesis (SI3)). They are the concrete examples of what changes if that closure is dropped while keeping the two-channel law.

**Example 6.2** (Quartic-log comparison). Let

$$F_4(x) = (\log x)^4.$$

Then

$$F_4(xy) + F_4(x/y) = (\log x + \log y)^4 + (\log x - \log y)^4.$$

Writing  $a = F_4(x)$  and  $b = F_4(y)$ , the induced combiner is  $\Phi(a, b) = 2a + 2b + 12\sqrt{ab}$  on  $[0, \infty)^2$  (from  $(s+t)^4 + (s-t)^4 = 2s^4 + 12s^2t^2 + 2t^4$  with  $a = s^4$ ,  $b = t^4$ ). It is continuous but agrees with no polynomial of total degree at most two: if a polynomial  $P$  of degree  $\leq 2$  agreed with  $\Phi$  on the image, then fixing  $b = 1$  the polynomial  $q(a) := P(a, 1) - 2a - 2$  would satisfy  $q(a) = 12\sqrt{a}$  for all  $a \geq 0$ , hence  $q(a)^2 = 144a$ , forcing  $2 \deg q = 1$ , which is impossible. Thus  $F_4$  satisfies route independence as a reciprocal comparison cost while failing finite pairwise polynomial closure.

The non-polynomial combiner shape  $\Phi(u, v) = 2u + 2v + 12\sqrt{uv}$  is excluded by structural hypothesis (SI3): it is continuous and symmetric on the cost image but is not a polynomial of total degree at most two. By contrast,  $\exp(u+v) - 1$  is not an admissible (C1)–(C3) combiner at all, since it fails the identity boundary  $\Phi(u, 0) = 2u$ .

The example above isolates a failure of finite pairwise polynomial closure (C6). The next example records four representative exclusions across the load-bearing clauses (C5)–(C7), including the hyperbolic miscalibration witness in item (d). It is not a Boolean independence table.

**Example 6.3** (Four exclusions illustrating the load-bearing clauses (C5)–(C7)). The following candidates locate the nearby failure modes of the clauses that do selection work. They are not a Boolean independence table; each item shows what one or more specific clauses exclude. Under the sign-neutral convention (C7a), item (c) becomes an exact witness for the indispensability of (C5).

- (a) *Additive cost.*  $F(x) = \frac{1}{2}(\log x)^2$  is continuous, reciprocally symmetric, satisfies route independence with combiner  $\Phi(u, v) = 2u + 2v$  (the  $c = 0$  degenerate bilinear branch), is non-negative, vanishes only at  $x = 1$ , and satisfies  $F(e^t) = t^2/2 \rightarrow \infty$  as  $|t| \rightarrow \infty$ . It therefore satisfies (C5) and the (C7b) curvature limit  $F(e^t)/(t^2/2) \rightarrow 1$ , but its combiner has  $c = 0$ , violating the non-degeneracy clause of (C6) and leaving (C7a)’s nonzero-bilinear unit convention inapplicable. Theorem 7.1(i) excludes this branch because the bilinear branch (C6) requires  $c \neq 0$ .
- (b) *Quartic-log cost.*  $F_4(x) = (\log x)^4$  of Example 6.2 is continuous, reciprocally symmetric, non-negative with unique zero at  $x = 1$ , and unbounded as  $|\log x| \rightarrow \infty$ , so it satisfies (C5). It is, however, not calibrated in the sense of (C7) (the leading log-Taylor coefficient is 0, not  $1/2$ ), and, more substantively, its induced combiner is not in the finite pairwise polynomial-of-degree- $\leq 2$  class, so it violates (C6). Theorem 7.1(i) excludes  $F_4$  at the combiner stage, before any calibration question is asked.
- (c) *Cosine cost.*  $F_{\cos}(x) = 1 - \cos(\log x)$  is continuous and reciprocally symmetric. It satisfies the  $c = -2$  mixed-coefficient equation, hence (C6), and is already in the sign-neutral unit of (C7a). Moreover  $\lim_{t \rightarrow 0} F_{\cos}(e^t)/(t^2/2) = 1$ , so it satisfies (C7b). It is non-negative, but its zero set is the infinite discrete subgroup  $\log x \in 2\pi\mathbb{Z}$ , not the singleton  $\{1\}$ , and it is bounded above by 2, so it fails

exactly (C5). This is the direct witness that comparison-cost nondegeneracy is indispensable.

- (d) *Miscalibrated hyperbolic cost.* For  $\lambda > 0$ ,  $\lambda \neq 1$ ,  $F_\lambda(x) = \cosh(\lambda \log x) - 1$  satisfies (C1)–(C6) and is already in the mixed-coefficient-2 unit of (C7a), but  $\lim_{t \rightarrow 0} F_\lambda(e^t)/(t^2/2) = \lambda^2 \neq 1$ , so it fails exactly (C7b). This is the direct witness that the curvature calibration is load-bearing after branch selection.

The four exclusions correspond to distinct failure mechanisms: (C6) excludes (a) at its non-degeneracy clause and (b) at its polynomial-class clause, (C5) excludes (c), and (C7b) excludes (d). Conversely, every continuous function satisfying (C1)–(C4) outside the canonical-unit hyperbolic family  $\cosh(\lambda \log x) - 1$  must fail at least one of (C5), (C6), or (C7). These are representative exclusions rather than a full independence table. The cost  $F_4$  of (b) is the same object as in Example 6.2, and the cosine branch (c) is the  $c < 0$  representative of Remark 6.7. More broadly, arbitrary Bregman or information divergences lie outside the present framework unless they happen to satisfy the reciprocal two-channel composition law of Definition 6.2.

### 6.3 Calibration of the mixed coefficient

The exclusions above locate the boundary of the retained family; what remains is to fix its one free parameter, which on the positive bilinear branch is a matter of choosing units.

Proposition 6.2 leaves a one-parameter bilinear family

$$F(xy) + F(x/y) = 2F(x) + 2F(y) + cF(x)F(y).$$

For  $c > 0$ , fix the cost unit by defining

$$\hat{F} := (c/2)F. \tag{1}$$

This is the canonical rescaling used throughout the paper; the unified  $|c|$ -form  $\hat{F} = (|c|/2)F$  of Remark 6.7 agrees with (1) on the positive branch ( $c > 0$ ) and extends it to the negative branch. Substitution in the bilinear law gives

$$\hat{F}(xy) + \hat{F}(x/y) = 2\hat{F}(x) + 2\hat{F}(y) + 2\hat{F}(x)\hat{F}(y),$$

i.e.  $\hat{F}$  satisfies the same law with mixed coefficient 2. The choice  $c = 2$  is therefore not a fitted physical number. It is the unit in which the bilinear branch has unit mixed coefficient, and the later log-curvature calibration makes the same unit choice at the identity.

**Remark 6.9** (Finite closure is a multiplicative composition principle). Hypothesis (C6) is easy to misread as an unmotivated truncation (“why total degree two and not three?”). Given (C1)–(C4) it is equivalent to a closure principle that carries no degree bookkeeping at all. With the canonical representative  $\hat{F}$  of (1), add 2 to both sides of the mixed-coefficient-2 law and set  $H := 1 + \hat{F}$ :

$$(1 + \hat{F}(xy)) + (1 + \hat{F}(x/y)) = 2(1 + \hat{F}(x))(1 + \hat{F}(y)),$$

$$\text{i.e.} \quad H(xy) + H(x/y) = 2H(x)H(y).$$

Thus the shifted cost  $H = 1 + \hat{F}$  composes multiplicatively across the two channels: the aligned and anti-aligned assemblies combine through a single product  $H(x)H(y)$ . The display is the  $c > 0$  branch (the one retained by (C5)), where  $\nu := c/2 > 0$  and  $H = 1 + \hat{F}$ ; for  $c < 0$  the same computation with the negative rescaling  $\nu = c/2 < 0$  gives  $H = 1 + \nu F = 1 - (|c|/2)F$ , recovering the convention of Remark 6.7. Conversely, any non-constant  $F$  for which some nonzero rescaling  $H = 1 + \nu F$  obeys this product law has combiner  $2u + 2v + 2\nu uv$ , of total degree two with nonzero bilinear coefficient. Hence (C6) is equivalent to the statement that there is a nonzero unit in which  $1 + \nu F$  is two-channel multiplicative (exactly the d'Alembert closure solved in Section 7); the precise two-sided statement is Theorem 6.3. The substantive content of finite closure is therefore not the integer “two” but the requirement that the two channels combine as a product of shifted costs; combiners of genuinely higher degree are precisely those for which no multiplicative shift exists. This does not make (C6) a theorem: multiplicative two-channel closure remains an assumption, recorded as structural hypothesis (SI3). But it replaces a bare degree bound with an interpretable composition law, and it is the form in which we recommend the hypothesis be weighed.

**Theorem 6.3** (Equivalent forms of finite closure). *Let  $F$  satisfy (C1)–(C4) and be non-constant. The following are equivalent.*

- (a) Polynomial closure (C6). *The combiner  $\Phi$  of (C3) is the restriction to  $F(\mathbb{R}_{>0}) \times F(\mathbb{R}_{>0})$  of a symmetric real polynomial of total degree at most two with nonzero bilinear coefficient.*
- (b) Multiplicative closure. *For some nonzero rescaling  $\nu \in \mathbb{R} \setminus \{0\}$ , the shifted cost  $H := 1 + \nu F$  satisfies  $H(xy) + H(x/y) = 2H(x)H(y)$  for all  $x, y \in \mathbb{R}_{>0}$ . On the positive bilinear branch  $c > 0$  retained by (C5) one may take  $\nu = c/2 > 0$ , recovering  $H = 1 + \hat{F}$  with (1); on the  $c < 0$  branch  $\nu = c/2 < 0$ , and  $H = 1 + \nu F = 1 - (|c|/2)F$  recovers the convention of Remark 6.7. The rescaling cannot in general be taken positive: for  $c < 0$  no  $\nu > 0$  works, since the induced bilinear coefficient  $2\nu$  must equal  $c$ .*
- (c) d'Alembert form. *In logarithmic coordinates the even function  $H(t) := 1 + \nu F(e^t)$  (with  $\nu$  as in (b)) solves the d'Alembert equation  $H(s+t) + H(s-t) = 2H(s)H(t)$  with  $H(0) = 1$  and  $H$  non-constant.*

*Proof.* (a) $\Rightarrow$ (b): by Proposition 6.2 the combiner is  $\Phi(u, v) = 2u + 2v + cuv$  with  $c \neq 0$ . Set  $\nu := c/2 \neq 0$  and  $G := \nu F$ ; multiplying the route-independence law by  $\nu$  gives  $G(xy) + G(x/y) = 2G(x) + 2G(y) + 2G(x)G(y)$ , and adding 2 to both sides gives, with  $H = 1 + G = 1 + \nu F$ , the relation  $H(xy) + H(x/y) = 2H(x)H(y)$ . No sign condition on  $c$  is needed. (b) $\Leftrightarrow$ (c): substitute  $x = e^s$ ,  $y = e^t$ ; evenness of  $H$  is reciprocity (C2), and  $H(0) = 1 + \nu F(1) = 1$  by normalization (C1). (c) $\Rightarrow$ (a): with  $G := H - 1 = \nu F$ , the d'Alembert identity unwinds to  $G(xy) + G(x/y) = 2G(x) + 2G(y) + 2G(x)G(y)$ ; dividing by  $\nu$  gives the combiner  $\Phi(u, v) = 2u + 2v + 2\nu uv$  for  $F$ , a symmetric polynomial of total degree two with nonzero bilinear coefficient  $c = 2\nu$ , which is form (a).  $\square$

**Proposition 6.4** (Finite closure is logically independent of the remaining axioms). *Hypothesis (C6) cannot be derived from (C1)–(C5):*

**Table 3** Continuous d'Alembert branches classified by the signs of the bilinear coefficient and curvature parameter

| sgn( $c$ ) | $\kappa$ | Cost $F$ (representative)                      | (C5)                             | Curvature limit                     |
|------------|----------|------------------------------------------------|----------------------------------|-------------------------------------|
| +          | 0        | $F \equiv 0$                                   | $\times$ no comparison           | $\times$ limit = 0                  |
| +          | −        | $F(e^t) = (2/c)(\cos \lambda t - 1) \leq 0$    | $\times$ negative                | $\times$ limit = $-2\lambda^2/c$    |
| +          | +        | $F(e^t) = (2/c)(\cosh \lambda t - 1) \geq 0$   | $\checkmark$                     | $\checkmark$ if $\lambda^2 = c/2$   |
| −          | 0        | $F \equiv 0$                                   | $\times$ no comparison           | $\times$ limit = 0                  |
| −          | −        | $F(e^t) = (2/ c )(1 - \cos \lambda t) \geq 0$  | $\times$ bounded, infinite zeros | $\checkmark$ if $\lambda^2 =  c /2$ |
| −          | +        | $F(e^t) = (2/ c )(1 - \cosh \lambda t) \leq 0$ | $\times$ negative                | $\times$ limit = $-2\lambda^2/ c $  |

- (i) The quartic-log cost  $F_4(x) = (\log x)^4$  (Example 6.2) satisfies (C1)–(C5) and is route-independent with the continuous symmetric combiner  $\Phi(u, v) = 2u + 2v + 12\sqrt{uv}$ , which is not a polynomial of degree at most two. Hence (C1)–(C5) do not imply the polynomial-closure clause of (C6).
- (ii) The additive cost  $F(x) = \frac{1}{2}(\log x)^2$  satisfies (C1)–(C5) with polynomial combiner  $\Phi(u, v) = 2u + 2v$  ( $c = 0$ ). Hence (C1)–(C5) together with polynomial closure do not imply the nonzero-bilinear clause  $c \neq 0$ .

Consequently (C1)–(C5) do not force (C6); it is an independent hypothesis relative to those conditions.

*Proof.* For (i),  $F_4$  is continuous, even in  $\log x$ , vanishes only at  $x = 1$ , and is unbounded, so (C1)–(C5) hold; the displayed  $\Phi$  is computed in Example 6.2, and  $\sqrt{uv}$  is not a polynomial in  $(u, v)$ . For (ii),  $\frac{1}{2}(\log x)^2$  is the additive branch, which satisfies (C1)–(C5) and whose combiner is  $2u + 2v$  by direct substitution. Each item exhibits a function satisfying every named axiom except the targeted clause of (C6), proving independence.  $\square$

Proposition 6.4 answers an open question in the negative: finite closure is not a redundant consequence of the comparison axioms but a genuine, irreducible selection principle. The most that can be done constructively (and what Theorem 6.3 does) is to restate it as the multiplicative two-channel composition law for a nonzero rescaling  $1 + \nu F$ , replacing an opaque degree bound with an interpretable closure axiom.

The remaining curvature parameter  $\kappa := H''(0)$  of the rescaled cost is fixed by the explicit log-curvature condition (C7b). The three real-branch alternatives ( $\kappa = 0$ ,  $\kappa < 0$ ,  $\kappa > 0$ ) of the resulting d'Alembert ODE, together with the branch selection by (C5) and the calibration  $\lambda = 1$  by (C7b), are classified inside the proof of Theorem 7.1; Table 3 summarizes the branch behavior across the bilinear coefficient  $c$  and curvature  $\kappa$ .

**Remark 6.10** (Non-redundancy of (C5), (C6), and (C7b)). A natural question is whether (C5) (comparison-cost nondegeneracy) is supplied by (C6) together with the two calibration clauses (C7a)–(C7b). Under the sign-neutral convention (C7a), the continuous d'Alembert branches available after (C6) form six candidate families, indexed by the sign of the bilinear coefficient  $c$  and by the sign of  $\kappa := H''(0)$  in the constant/cosine/hyperbolic trichotomy classified inside the proof of Theorem 7.1. Table 3 records the branch behavior at general  $c$ ; in the (C7a) unit  $|c| = 2$ , the curvature limits specialize to  $\pm\lambda^2$ .

In the sign-neutral unit  $c = -2$ , the row  $\text{sgn}(c) = -$ ,  $\kappa < 0$  contains the exact witness  $F(x) = 1 - \cos(\log x)$ . It satisfies (C1)–(C4), (C6), (C7a), and (C7b), but fails (C5): it is bounded and has zeros on the infinite set  $\log x \in 2\pi\mathbb{Z}$ . Thus (C5) is not supplied by closure and calibration. The three selection requirements perform distinct work in the dependency order:

- (C6) removes the additive branch  $c = 0$ ; its unit-log-curvature-normalized solution is  $F(x) = \frac{1}{2}(\log x)^2$ .
- (C5) removes the negative-bilinear cosine witness above, which satisfies the remaining displayed conditions.
- (C7b) removes the miscalibrated positive hyperbolic family  $F_\lambda(x) = \cosh(\lambda \log x) - 1$  with  $\lambda \neq 1$ , which satisfies (C1)–(C6), (C7a), and (C5).

The intersection  $\{(C1)–(C5), (C6), (C7a), \text{ and } (C7b) \text{ all hold}\}$  is the single calibrated solution  $J(x) = \frac{1}{2}(x + x^{-1}) - 1$ .

## 7 Uniqueness of the Cost

**Theorem 7.1** (Calibrated unique reciprocal cost). *Let  $F$  be a reciprocal comparison cost in the sense of Definition 6.2, i.e.  $F$  satisfies (C1)–(C5).*

- (i) *Under the basic conditions (C1)–(C5) together with (C6) (finite pairwise polynomial closure), Proposition 6.2 gives the polynomial form  $\Phi(u, v) = 2u + 2v + cuv$  on the cost image. The non-degeneracy clause of (C6) gives  $c \neq 0$ , and the comparison-cost nondegeneracy condition (C5) excludes  $c < 0$ ; hence the retained reciprocal branch has  $c > 0$  and*

$$F(xy) + F(x/y) = 2F(x) + 2F(y) + cF(x)F(y).$$

- (ii) *For the canonical representative  $\hat{F} = (c/2)F$ , the continuous real solutions of the resulting mixed-coefficient-2 equation split into constant, cosine, and hyperbolic branches. The comparison-cost branch, selected by (C5), is the one-parameter positive hyperbolic family*

$$\hat{F}_\lambda(x) = \cosh(\lambda \log x) - 1, \quad \lambda > 0.$$

*Thus no unique  $J$  is selected before the explicit calibration condition (C7).*

- (iii) *If  $F$  is calibrated in the sense of (C7), then the unit convention (C7a) fixes  $|c| = 2$  and the branch selection by (C5) (a standing hypothesis of this theorem) leaves  $c = 2$ , while the log-curvature condition (C7b) fixes  $\lambda = 1$ . Hence the unique calibrated reciprocal comparison cost is*

$$J(x) = \frac{1}{2}(x + x^{-1}) - 1.$$

*Proof. Strategy.* The argument proceeds in four stages.

- (i) Polynomial reduction. Proposition 6.2 narrows the combiner to  $\Phi(u, v) = 2u + 2v + cuv$  with  $c \neq 0$ ; comparison-cost nondegeneracy (C5) and Remark 6.7 then exclude the  $c < 0$  branch.

- (ii) Unit rescaling. The canonical convention (1) puts the retained  $c > 0$  branch in mixed-coefficient 2 form.
- (iii) Reduction to d'Alembert. In logarithmic coordinates with  $H := \widehat{F}(e^{\cdot}) + 1$ , the equation becomes the d'Alembert form  $H(s+t) + H(s-t) = 2H(s)H(t)$ , with the three continuous real solution branches *constant*, *cosine*, and *hyperbolic*.
- (iv) Branch selection and calibration. The hyperbolic branch is selected by (C5), and the unit log-curvature condition (C7b) fixes  $\lambda = 1$ , yielding  $J(x) = \frac{1}{2}(x+x^{-1}) - 1$ .

We now execute each stage in turn.

*Stage (i): polynomial reduction.* Part (i) is Proposition 6.2: under (C1)–(C4) and (C6), with the non-constancy hypothesis of the proposition supplied by (C5) (the growth clause  $F(e^t) \rightarrow +\infty$  forces  $F$  non-constant), the polynomial extension takes the form  $\Phi(u, v) = 2u + 2v + cuv$  on the cost image. The non-degeneracy clause of (C6) gives  $c \neq 0$ . The exclusion of  $c < 0$  by (C5) uses only the d'Alembert classification of Stage (iii), applied through Remark 6.7, and is completed there; the retained reciprocal branch has  $c > 0$ .

*Stage (ii): unit rescaling.* Apply the canonical rescaling (1),  $\widehat{F} := (c/2)F$ ; on the retained  $c > 0$  branch the positive factor  $c/2$  equals  $|c|/2$ , so this rescaling is exactly the sign-neutral unit convention (C7a) (the negative branch  $c = -2$  having been excluded by (C5) in Stage (i), consistently with the dependency order of Remark 6.10). Then  $\widehat{F}$  satisfies the same route-independence law with mixed coefficient 2:

$$\widehat{F}(xy) + \widehat{F}(x/y) = 2\widehat{F}(x) + 2\widehat{F}(y) + 2\widehat{F}(x)\widehat{F}(y).$$

Because  $c > 0$ , this rescaling preserves all five clauses of (C1)–(C5): non-negativity (positive scaling), the isolated zero at  $x = 1$  (scaling does not move zeros), reciprocity (preserved pointwise), continuity, and unbounded growth ( $F(e^t) \rightarrow +\infty$  implies  $(c/2)F(e^t) \rightarrow +\infty$  since  $c/2 > 0$ ). The polynomial extension and its bilinear coefficient are scaled accordingly, so (C6) is preserved with new mixed coefficient 2.

*Stage (iii): reduction to d'Alembert and ODE classification.* Substitute log coordinates and shift by one:

$$G(t) := \widehat{F}(e^t), \quad H(t) := G(t) + 1.$$

**Reduction to d'Alembert.** With  $x = e^s$  and  $y = e^t$ , the canonical law becomes

$$G(s+t) + G(s-t) = 2G(s)G(t) + 2G(s) + 2G(t).$$

Adding 1 to both sides and using  $H = G + 1$ ,

$$H(s+t) + H(s-t) = G(s+t) + 1 + G(s-t) + 1 = 2G(s)G(t) + 2G(s) + 2G(t) + 2.$$

Expanding  $2H(s)H(t) = 2(G(s)+1)(G(t)+1) = 2G(s)G(t) + 2G(s) + 2G(t) + 2$ . The two sides agree:

$$H(s+t) + H(s-t) = 2H(s)H(t). \tag{2}$$

This is the d'Alembert functional equation.

**Initial conditions.** Normalization gives  $G(0) = 0$  and  $H(0) = 1$ . Reciprocity gives  $G(-t) = G(t)$ , so  $H$  is even. Before calibration, the value  $H''(0)$  is the remaining curvature parameter of the canonical representative.

**Regularity.** By hypothesis  $F$  is continuous on  $\mathbb{R}_{>0}$ , so  $H$  is continuous on  $\mathbb{R}$ . Equation (2) together with continuity of  $H$  and  $H(0) = 1$  forces  $H$  to be smooth, by the integral-bootstrap argument recorded in Appendix A. That appendix also records the standard functional-equation strengthening in which measurability plus local boundedness can replace continuity; the theorem keeps continuity as its stated regularity hypothesis, and the weaker route is consumed only by the explicitly external Corollary 7.2 below.

**ODE reduction.** Let  $D_t$  denote partial differentiation in  $t$ . Applying  $D_t^2$  to (2):

$$H''(s+t) + H''(s-t) = 2H(s)H''(t).$$

Evaluating at  $t = 0$ :

$$2H''(s) = 2H(s)H''(0).$$

Hence

$$H''(s) = H''(0)H(s). \quad (3)$$

**ODE classification.** Equation (3) is a linear second-order ODE with constant coefficient  $\kappa := H''(0)$ . Together with the even initial data  $H(0) = 1$ ,  $H'(0) = 0$  its unique solution is

$$H(s) = \begin{cases} \cosh(\sqrt{\kappa}s) & \kappa > 0, \\ 1 & \kappa = 0, \\ \cos(\sqrt{-\kappa}s) & \kappa < 0. \end{cases}$$

Before calibration, these are the constant, cosine, and hyperbolic real branches. The classification of continuous solutions of the d'Alembert equation (2) into precisely these three families is the standard result of Aczél [1] (see also Aczél and Dhombres [2] for the measurability-only refinement recorded in Appendix A). On the hyperbolic branch, writing  $\lambda = \sqrt{\kappa} > 0$  gives  $\hat{F}(e^s) = \cosh(\lambda s) - 1$ .

*Stage (iv): branch selection and calibration.* **Branch selection.** The comparison-cost nondegeneracy condition (C5) selects the positive-curvature hyperbolic branch. The constant branch ( $\kappa = 0$ ) gives  $H \equiv 1$ ,  $\hat{F} \equiv 0$ , and hence (since  $c \neq 0$  and  $F = (2/c)\hat{F}$ )  $F \equiv 0$ , which has no comparison content; this violates two clauses of (C5) simultaneously: the unique-zero clause (every  $x$  is a zero of  $F$ ) and the unbounded-growth clause ( $F(e^t) \equiv 0 \not\rightarrow +\infty$  as  $|t| \rightarrow \infty$ ). It is also incompatible with the non-constancy hypothesis of Proposition 6.2, since  $F \equiv 0$  is a constant solution. We attribute the exclusion to (C5) because the non-degeneracy clause of (C6) is a statement about the bilinear coefficient  $c$ , not about  $F$ , and is consistent with the trivial solution  $F \equiv 0$  on the degenerate cost image  $F(\mathbb{R}_{>0}) = \{0\}$ . The cosine branch ( $\kappa < 0$ ) gives  $\hat{F}(e^s) = \cos(\sqrt{-\kappa}s) - 1 \leq 0$ , with zeros at  $\sqrt{-\kappa}s \in 2\pi\mathbb{Z}$ ; this violates the non-negativity and unique-zero clauses of (C5). This exclusion is the positive-curvature

analogue of the negative-bilinear-branch exclusion in Remark 6.7. Thus the retained branch is

$$\widehat{F}(e^s) = \cosh(\lambda s) - 1, \quad \lambda > 0.$$

**Calibration and assembly.** If  $F$  is calibrated in the sense of (C7), then the cost unit has already been chosen so that the bilinear coefficient is  $c = 2$ . Hence  $\widehat{F} = F$ . The log-curvature condition in (C7) gives

$$F(e^t) = \frac{1}{2}t^2 + o(t^2).$$

On the hyperbolic branch,

$$F(e^t) = \cosh(\lambda t) - 1 = \frac{\lambda^2}{2}t^2 + O(t^4),$$

so  $\lambda^2 = 1$ . Since  $\lambda > 0$ ,  $\lambda = 1$ . Substituting  $t = \log x$ :

$$F(x) = \cosh(\log x) - 1 = \frac{1}{2}(e^{\log x} + e^{-\log x}) - 1 = \frac{1}{2}(x + x^{-1}) - 1.$$

This is  $J(x)$ . □

**Consequences and variants.** The remainder of this section collects the material that accompanies the theorem without entering its proof: the measurable-regularity variant (Corollary 7.2), an existence check (Remark 7.1), the separation of the three scaling parameters (Remark 7.2), and the basic identities and local expansion of  $J$ .

**Corollary 7.2** (Measurable variant of the uniqueness theorem). *In Theorem 7.1 the continuity hypothesis (C4) may be weakened to:  $F$  is Lebesgue measurable and locally bounded on  $\mathbb{R}_{>0}$ . All remaining hypotheses are unchanged and remain in force; in particular every clause of (C5) is still assumed, and its unbounded-growth clause becomes load-bearing in a new way, replacing connectedness of the image in the polynomial-reduction step (see the proof). The conclusions of parts (i)–(iii) are unchanged. The proof of this variant consumes one external ingredient, the measurable d’Alembert classification (Lemma A.1, recorded from Aczél–Dhombres [2]), so its status is “theorem modulo a cited standard result”; the headline theorem of the paper remains the continuity version.*

*Proof.* Only two steps of the proof of Theorem 7.1 use the regularity of  $F$ . First, Stage (i) (Proposition 6.2) uses continuity solely to guarantee that  $\text{Range}(F)$  is large, via connectedness of the image; but the vanishing of the degree- $\leq 2$  polynomial  $p(u) = a + bu + du^2 - 2u$  on  $\text{Range}(F)$  forces  $p \equiv 0$  as soon as the range contains three distinct values, and (C1) together with the unbounded-growth clause of (C5) already supplies infinitely many:  $0 = F(1)$  lies in the range, and  $F(e^t) \rightarrow +\infty$  gives values exceeding every bound. Second, Stage (iii) replaces the continuity bootstrap by Lemma A.1: measurability and local boundedness of  $F$  transfer to  $H(t) = 1 + (c/2)F(e^t)$  along the diffeomorphism  $t \mapsto e^t$  (which preserves Lebesgue measurability in both directions and maps compacta to compacta), and the lemma then yields continuity,

smoothness, and the constant/cosine/hyperbolic classification. Branch selection by (C5) and calibration by (C7) are pointwise and proceed exactly as in Stages (ii) and (iv).  $\square$

**Remark 7.1** (Existence of the calibrated cost). Theorem 7.1 is a *uniqueness* theorem: it characterizes  $J$  as the only function satisfying (C1)–(C7). For the characterization to be non-vacuous,  $J$  itself must satisfy all seven conditions, and it does.

(C1)  $J(1) = \frac{1}{2}(1 + 1) - 1 = 0$ .

(C2)  $J(x^{-1}) = \frac{1}{2}(x^{-1} + x) - 1 = J(x)$ .

(C3) Using the addition formula  $\cosh(A + B) + \cosh(A - B) = 2 \cosh A \cosh B$  with  $A = \log x$ ,  $B = \log y$ :

$$\begin{aligned} J(xy) + J(x/y) &= (\cosh(\log x + \log y) - 1) + (\cosh(\log x - \log y) - 1) \\ &= 2 \cosh(\log x) \cosh(\log y) - 2 \\ &= 2(J(x) + 1)(J(y) + 1) - 2 \\ &= 2J(x) + 2J(y) + 2J(x)J(y). \end{aligned}$$

The combiner is  $\Phi(u, v) = 2u + 2v + 2uv$ , a symmetric polynomial with mixed coefficient  $c = 2$ .

(C4)  $J(x) = \frac{1}{2}(x + x^{-1}) - 1$  is continuous on  $\mathbb{R}_{>0}$ .

(C5) By the AM–GM inequality,  $\frac{1}{2}(x + x^{-1}) \geq 1$  with equality iff  $x = 1$ , so  $J(x) \geq 0$  with unique zero at  $x = 1$ . Also  $J(e^t) = \cosh t - 1 \rightarrow +\infty$  as  $|t| \rightarrow \infty$ .


(C6) The combiner  $\Phi(u, v) = 2u + 2v + 2uv$  is a symmetric polynomial of total degree two with non-degenerate bilinear term ( $c = 2 \neq 0$ ).

(C7a) The mixed coefficient is already  $c = 2$ , so no unit rescaling is needed.

(C7b)  $\lim_{t \rightarrow 0} J(e^t)/(t^2/2) = \lim_{t \rightarrow 0} (\cosh t - 1)/(t^2/2) = 1$ .

Hence the class characterized by Theorem 7.1 is non-empty.

This is the central mathematical result under the stated hypotheses. The rest of the paper uses  $J$  as calibrated cost structure, not as a fitted loss. The same cost appears in the reciprocal-cost uniqueness and ratio-matching literature [6, 7, 9].



**Fig. 2** Calibrated reciprocal comparison cost  $J(x) = \frac{1}{2}(x + x^{-1}) - 1 = \cosh(\log x) - 1$  in log-ratio coordinates. The curve is symmetric under  $x \mapsto x^{-1}$ , strictly positive away from the identity  $x = 1$ , and unbounded as  $|\log x| \rightarrow \infty$ .

Figure 2 plots  $J(x)$  against  $\log x$ . The curve is symmetric under  $x \mapsto x^{-1}$  because  $J$  depends only on the log-ratio  $t = \log x$ . The unique zero at  $x = 1$  ( $t = 0$ ) is the identity comparison; every other ratio carries strictly positive cost. Near the identity the graph follows the local quadratic law  $J(e^t) = \frac{1}{2}t^2 + O(t^4)$ , which is the unit log-curvature calibration in graphical form (Corollary 7.4). Away from the identity the cost rises without bound as  $x \rightarrow 0^+$  or  $x \rightarrow \infty$ , so extreme mismatches are infinitely expensive under the calibrated law. The hyperbolic cosine shape is the branch selected by  $H''(0) > 0$  in the proof; the cosine branch would dip below zero for small  $|t|$  and is incompatible with a non-negative reciprocal comparison cost.

**Remark 7.2** (Three scaling symbols). The paper uses three scaling parameters and keeps them disjoint. The *recognition-work unit scaling*  $C \mapsto \mu C$  with  $\mu > 0$  is the cost rescaling of Proposition 5.4; it rescales costs on finite configurations before any ratio structure is admitted. The *closure rescaling*  $\nu = c/2$  of Theorem 6.3 shifts the ratio-level cost into the multiplicative d'Alembert form  $H = 1 + \nu F$ . The *log-curvature parameter*  $\lambda > 0$  of the hyperbolic branch  $F_\lambda(x) = \cosh(\lambda \log x) - 1$  (Theorem 7.1) fixes the curvature of the reciprocal comparison cost in log coordinates at the identity. These act at different layers:  $\mu$  rescales configuration-level costs;  $\nu$  rescales the ratio-level cost for closure;  $\lambda$  selects the hyperbolic branch after the positive bilinear unit is fixed. Throughout this paper  $\mu$  never appears as a closure or curvature parameter,  $\nu$  never appears as a recognition-work factor, and  $\lambda$  never appears as either.

**Corollary 7.3** (Basic identities for the calibrated cost). *For every  $x > 0$ ,*

$$J(x) = \frac{(x-1)^2}{2x} = \frac{1}{2} \left( \sqrt{x} - \frac{1}{\sqrt{x}} \right)^2.$$

*Consequently  $J(x) \geq 0$ ,  $J(x) = 0$  if and only if  $x = 1$ ,  $J(x) = J(x^{-1})$ , and  $\lim_{x \rightarrow 0^+} J(x) = +\infty$ .*

*Proof.* The two displayed identities are algebraic expansions of  $\frac{1}{2}(x+x^{-1})-1$ . Positivity and the unique zero follow from the square form. Reciprocity is immediate from the first definition. Divergence at zero follows from the  $x^{-1}/2$  term.  $\square$

**Corollary 7.4** (Local quadratic law). *In logarithmic coordinates,*

$$J(e^t) = \cosh t - 1 = \frac{1}{2}t^2 + \frac{1}{24}t^4 + O(t^6).$$

*Proof.* Substitute  $x = e^t$  and use the Taylor series of  $\cosh$  at the origin.  $\square$

**Nearby alternatives at the functional-equation layer.** Four representative exclusions across (C5)–(C7), including the hyperbolic miscalibration witness in Example 6.3, item (d), together with an additive log-squared cost, the quartic-log cost  $F_4(x) = (\log x)^4$ , and a cosine cost  $F_{\cos}(x) = 1 - \cos(\log x)$ , are tabulated in Example 6.3 in Section 6.2. The claim of Theorem 7.1 is therefore not that  $J$  is the only mathematically reasonable comparison cost, but that it is the unique cost compatible with the named reciprocal recognition constraints (C1)–(C7).

## 8 The Scope of the Theorem: Two Kernel Certificates and One Independence Witness

This section states three results that fix the exact scope of the theorem of Part I: what it yields on its own, and what requires further input. Two are theorems in the Lean 4 proof assistant, verified by the kernel at the pinned public commit described in the Data and code availability statement, and the semantics stated here follow the formal statements exactly. The third is the independence witness of Proposition 6.4, proved in Part I. The three have different logical characters, and we keep them distinct: a division-of-labor result for dimensionful constants, a necessity result for the composition law, and an independence witness for the closure hypothesis. All three are *local*: each settles one named question exactly, and physical extrapolation remains open along every route not named. The kernel certificates of this section are closed at their stated hypotheses. The finite arithmetic facts exposed by the verification target are rechecked by kernel decide; they are distinct from the two scope certificates below. Exactly two results are kernel certificates: Certificate 8.1 (dimensional rigidity) and Certificate 8.2 (necessity of the composition law). The repository also checks the signature of the reciprocal-cost uniqueness theorem conditional on the typeclass `AczelSmoothnessPackage`; that conditional signature is not counted here as a closed kernel certificate, and the missing instance is discussed in the Data and code availability statement. The conditional structures of Section 9—the golden-ratio ladder, the conditional dimension-three argument, and the eight-tick cycle—are *not* kernel-derived physical conclusions: their machine-checked components certify stated definitional or arithmetic implications only. For dimension three this is the implication  $(D - 2 = 1) \Rightarrow D = 3$ , with Alexander duality entering as a definition rather than as a kernel-checked theorem (Section 9); the genuinely topological input  $H_1(S^1; \mathbb{Z}) \cong \mathbb{Z}$  has theorem declarations available in the pinned import closure, while Alexander duality and its wiring into the linking argument remain outside the pinned formalization.

*Certificate 8.1* (Dimensional rigidity of  $(c_{\text{light}}, \hbar, G)$ ). In the (M, L, T) basis the exponent vectors of the speed of light  $c_{\text{light}}$ ,  $\hbar$ , and  $G$  are  $(0, 1, -1)$ ,  $(1, 2, -1)$ , and  $(-1, 3, -2)$ ; the matrix they form has determinant  $-2 \neq 0$ , so the only dimensionless monomial  $c_{\text{light}}^a \hbar^b G^g$  is the trivial one ( $a = b = g = 0$ ). Consequently no framework whose outputs are pure numbers can produce the SI value of  $c_{\text{light}}$ ,  $\hbar$ , or  $G$  without an external dimensional unit calibration. (Machine-checked as `no_dimensionless_combination`: for rational exponents  $a, b, g$ , if  $a \dim(c_{\text{light}}) + b \dim(\hbar) + g \dim(G) = 0$  componentwise, then  $a = b = g = 0$ .)

The mathematics of this certificate is classical: it is the dimensional-homogeneity constraint of Bridgman [26], and the nonsingularity of the  $(c_{\text{light}}, \hbar, G)$  exponent matrix is precisely why Planck units exist. What the kernel-checked form adds is permanence: the reading is fixed once and applies uniformly. Our framework carries native-unit identities  $\hbar_R = \varphi^{-5}$  and  $G_R = \varphi^5/\pi$  (dimensionless assignments in recognition units for the action and curvature normalizations, not SI measurements), and Certificate 8.1 settles their meaning as a matter of kernel-checked arithmetic: the native identities fix dimensionless algebraic structure in recognition units; conversion to SI values requires a dimensionful native-to-SI unit map, and that map is measured, not derived. (Fixing

all three of  $c_{\text{light}}$ ,  $\hbar$ ,  $G$  in SI requires anchors fixing all three unit conversions; in the post-2019 SI, where  $c_{\text{light}}$  and  $\hbar$  are exact by definition, a single measured anchor completes the map.) At the dimensionless level the identities contribute the paired exponent structure ( $-5$  in the action identification,  $+5$  in the curvature identification). The certificate also yields a sharp criterion for any future “derivation of  $\hbar$ ”: the construction must make  $\varphi^{-5}$  load-bearing, so that a different rung would give a different observable. (At the pinned commit the criterion is open: the spectral module that writes a Schrödinger-form evolution on the eight-tick carrier fixes its eigenvalues by phase-matching, so the value of  $\hbar$  cancels identically.) Dimensionless quantities (mass ratios, coupling values, degree counts) remain the natural targets throughout.

*Certificate 8.2* (Necessity of the composition law). The composition law is a necessary input: two kernel-checked results establish its necessity. First, the equality-induced (Hamming-type) cost on  $(\mathbb{R}_{>0}, \cdot)$  with any nonzero weight fails composition consistency: no combiner  $P$  exists with  $C(xy, 1) + C(x/y, 1) = P(C(x, 1), C(y, 1))$  (machine-checked as `equality_cost_insufficient_for_recognition`). A cost supporting the composition analysis must therefore respect the carrier’s multiplicative structure, and that requirement is substantive, not definitional. Second, the ledger-derived constraints, regularity, and unit curvature calibration do not force the d’Alembert identity: the quadratic cost  $G(t) = t^2/2$  is even, vanishes at 0, is continuous, and has unit second derivative at the origin, yet violates the d’Alembert-type identity (machine-checked as `aczel_hypothesis_refuted`, with the quadratic cost as explicit witness). Symmetry, normalization, continuity, and calibration therefore leave the cost underdetermined; the multiplicative closure law equivalent to (C6) supplies the selection that singles out  $J$ .

A caution on attribution accompanies Certificate 8.2. Aczél’s classification runs in one direction only: it classifies solutions *of* the d’Alembert equation. No classical result asserts that evenness, normalization, continuity, and calibration force the equation, and the certificate proves that no such result can exist. The uniqueness theorem of Part I is therefore conditional on (C1)–(C7), including the load-bearing closure hypothesis, and is not unconditionally forced from the floor. Remark 6.2 in Part I records the same boundary from the paper’s side.

**Proposition 8.1** (Independence witness for the closure hypothesis). *The quartic-logarithmic cost  $F_4(x) = (\log x)^4$  is continuous, reciprocally symmetric, normalized, and satisfies the nondegeneracy condition (C5), but no finite pairwise polynomial closure (C6) holds for it; hence (C6) is logically independent of (C1)–(C5) and cannot be derived from them (Example 6.2 and Proposition 6.4 in Part I). This is an independence witness, not a pathology: it proves the hypothesis set of Theorem 7.1 is tight, in the sense that deleting the closure hypothesis admits explicit alternative costs.*

The two certificates and the independence witness compose into a single division of labor. Certificate 8.1 fixes what dimensionful constants require: an external dimensional unit calibration beyond the dimensionless structure. Certificate 8.2 fixes what  $J$  requires: the composition law, so every downstream structure carries (C6) as an explicit hypothesis. Proposition 8.1 shows that this hypothesis is irreducible. What stands is the calibrated characterization theorem of Part I, relative to its named hypotheses, and any physical extrapolation now knows exactly which bridge it must supply. Section 9 states the conditional structures built on this basis; what remains open is collected in Section 10.

## 9 Conditional Outlook from the Wider Framework

Three structural constructions from our wider framework extend the theorem toward physical structure, in conditional form. For each we state what is proved, on which hypotheses, and what would upgrade it. Full developments are in the pinned repository and the companion papers; nothing below is consumed by Part I, and nothing below is asserted unconditionally of nature: each structure is a theorem conditional on its named bridge hypotheses. Three facts about the formalization, themselves machine-checked, are used below: the “closed-observable layer” means the repository fragment containing only closed recognition observables and their algebra, before ledger dynamics or external physical bridges are added; (i) additive closure is not forced by that layer alone; (ii) the realization layer for the eight-tick cycle is predicate-level, not a cell-complex proof; and (iii) the homology computation  $H_1(S^1; \mathbb{Z}) \cong \mathbb{Z}$  has theorem declarations available in the pinned import closure, but it is not yet wired as a premise of the linking argument, whose Alexander-duality step currently enters as a definition.

*Terminology and status.* Read this section as definitional scope, not as new self-contained theorems. The constructions below are stated relative to primitive notions of the wider framework whose full definitions live in the pinned repository and the companion papers; Part I does not define them, and no statement in this section is self-contained at the cost level of Sections 6–7. For readability we gloss the load-bearing terms once. A *ledger loop* is a closed oriented walk of recognition entries. A *loop-pair separator* ((SI7)) is an additive, orientation-respecting  $\mathbb{Z}$ -valued functional on loops that is nonzero on a suitable pair of disjoint closed loops. The *loop-class group* is the abelian group of closed loops modulo the separator-trivial ones; its *rank* ((SI8)) is the rank of that abelianization. The *minimal torsion-free target* ((SI9)) is the group  $\mathbb{Z}$  into which the separator factors. The *embedding postulate* ((SI10)) identifies this abstract  $\mathbb{Z}$ -valued separator with geometric integer linking of loops in space. Each paragraph below is to be read as “conditional on these named inputs, the following holds”, with the inputs themselves neither proved here nor claimed of nature.

### The golden-ratio scale ladder (conditional on (SI6))

On any geometric ladder  $\{s_0 r^n\} \subset \mathbb{R}_{>0}$  with  $r > 1$ , the  $J$ -cost of one rung is constant (an elementary property of  $J$ ’s log-symmetry), and imposing adjacent additive closure  $s_n + s_{n+1} = s_{n+2}$  forces  $r^2 = r + 1$ , hence  $r = \varphi = (1 + \sqrt{5})/2 \approx 1.618$ , the positive root of that equation (the negative root is excluded by  $r > 1$ ). This is an algebraic identity given hypothesis (SI6); it derives the golden ratio from additive closure, not from  $J$  alone, and (SI6) is an assumption about which ladders are realized. The identity and the uniqueness of the self-similar ratio are machine-checked, and a machine-checked obstruction shows that the closed-observable layer of the framework cannot by itself supply hypothesis (SI6). The upgrade path is to derive adjacent additive closure from ledger dynamics; this is open.

### Dimension three (conditional on (SI7)–(SI10))

Under the separator package (SI7)–(SI9), the framework’s loop-separator structure exists recognition-internally in every dimension;  $D = 3$  enters *only* through the

embedding postulate (SI10), which identifies the abstract  $\mathbb{Z}$ -valued separator with spatial integer linking of closed loops, after which classical Alexander duality—the theorem relating the reduced homology of the complement of an embedded sphere in  $S^D$  to the reduced cohomology of the sphere itself [27]—forces  $D = 3$  uniquely (non-trivial circle linking requires  $\tilde{H}^{D-2}(S^1; \mathbb{Z}) \neq 0$ , hence  $D - 2 = 1$ ). The reduction isolating (SI10) as the single remaining input is a structural result of the framework, and the postulate is open: we know of no recognition-internal argument that discharges it. The reduction locates exactly where any such argument would have to act. One disclosure on the formal side is needed. The machine-checked statement encodes the classical cohomological fact (Alexander duality) as a definition, so the checked theorem is the arithmetic implication  $(D - 2 = 1) \Rightarrow D = 3$ . The genuinely topological ingredient,  $H_1(S^1; \mathbb{Z}) \cong \mathbb{Z}$ , has theorem declarations available in the pinned import closure (the CircleWindingChain modules, theorems circleH1ZIsoInt\_holds and circleH1ZNonzero\_unconditional). The residual formal gap is therefore exactly two-fold: a Lean formalization of Alexander duality (absent from Mathlib and from the library), and the wiring of the  $H_1$  computation in as a premise of the linking argument. The residual physical bridge is the embedding postulate (SI10) alone. Same-sector linking arithmetic alone permits every odd  $D \geq 3$ ; restricting to loop-loop linking is the additional topological input. Concretely, two disjoint oriented closed submanifolds of dimensions  $p$  and  $q$  in  $S^D$  carry a well-defined linking number precisely when their dimensions are complementary,  $p + q = D - 1$ ; for two loops ( $p = q = 1$ ) this forces  $D - 1 = 2$ , i.e.  $D = 3$ . Permitting equal-dimensional self-linking  $p + p = D - 1$  instead admits every odd  $D = 2p + 1 \geq 3$ , which is exactly why restricting the linked objects to one-dimensional loops is a genuine additional input rather than a consequence of the linking arithmetic alone. The formalization therefore certifies the conditional argument; it is not an independent proof that space is three-dimensional.

### **The eight-tick cycle (conditional on binary atomicity, occupancy, and $D = 3$ ).**

Given binary-atomic recognition acts, a fully occupied separating recognizer family of  $D$  bits (Section 5.2), and  $D = 3$ , the minimal complete re-identification cycle has period exactly  $2^D = 8$ . The machine-checked statement consumes the dimension package as a hypothesis, and an explicit Gray-code cycle—a cyclic listing of the  $2^D$  bit-strings in which consecutive entries differ in exactly one bit—exhibits the period; the circle-realization layer of the formalization is predicate-level rather than geometric. The period is relocated to those three inputs, not removed: binary atomicity and full occupancy are substantive assumptions, and  $D = 3$  is itself conditional, as above.

## **10 Conclusion**

This paper delivers two things. The first is a theorem: a calibrated d’Alembert-type characterization selecting  $J(x) = \frac{1}{2}(x + x^{-1}) - 1$  as the unique calibrated continuous reciprocal cost. The theorem is hypothesis-tight in a concrete sense. Its closure hypothesis is proved equivalent to an interpretable composition principle and logically independent of the remaining axioms. Its carrier is reduced to measurement axioms.

Its measurable-regularity variant is also given. The load-bearing conditions (C5), (C6), and (C7b) are each individually indispensable.

The second is a precisely drawn scope. Two kernel-checked certificates and one paper-proved independence witness establish what each further step requires: an external dimensional unit calibration for dimensionful constants, the composition law for  $J$ , and the closure principle as an independent axiom. Beyond the certificates, three conditional structures from the wider framework (the golden-ratio scale ladder, the conditional dimension-three argument, and the eight-tick cycle) are stated with the exact hypotheses on which they rest. None of the three is claimed as kernel-derived: each is a theorem conditional on its named bridge hypotheses, and the dimension-three formalization certifies only the arithmetic implication  $(D - 2 = 1) \Rightarrow D = 3$  with Alexander duality assumed as a definition.

What stands is one theorem; what remains are four named bridges, each stated where it arises and collected here as the forcing targets of the program. (1) The continuum residue of Proposition 5.5: derive order density and Dedekind completeness from recognition work alone (Remark 5.3). (2) The closure principle: derive (C6) from something weaker; Proposition 6.4 proves it cannot come from the remaining comparison axioms, so any derivation must supply new structure. (3) Adjacent additive closure ((SI6)): derive it from ledger dynamics; a machine-checked obstruction shows the closed-observable layer alone cannot supply it. (4) The embedding postulate ((SI10)): find a recognition-internal warrant; the reduction of Section 9 locates exactly where such an argument must act. Hilbert space, the Born rule, the Standard Model gauge group, and measured constants are not claimed or used in this paper.

A remark on method is in order, because it is transferable. Machine-checked proof entered mathematics through affirmative landmarks: Gonthier verified the four-color theorem and later the Feit–Thompson odd-order theorem, Hales completed the Flyspeck verification of the Kepler conjecture, and in the Liquid Tensor Experiment Scholze asked the Lean community to check a theorem whose proof he did not fully trust [22–25]. We use the same instrument to keep a foundational framework aligned with its theorems. Every machine-checked claim identified as such in this paper is pinned to a public commit, and the two scope certificates of Section 8 are checked by the Lean kernel. Results proved only in this paper are identified separately. This prevents the formal claims and their limits from drifting as the surrounding prose evolves. One exception is stated explicitly rather than hidden: the slice’s uniqueness theorem assumes the smoothness typeclass `AczelSmoothnessPackage`, which has no instance at the pinned commit; that instance is discharged in the program’s full development library and reproduced in prose in Appendix A, as recorded in the Data and code availability statement. The kernel does not know what a lepton is, and it cannot tell us whether a bridge hypothesis is true of the world. What it can do is refuse to let the distance between a definition and its intended meaning disappear into confident writing.

Whether these open questions close is not prejudged here. What the paper establishes is one theorem, and a scope around that theorem precise enough to be checked by a machine.

**Acknowledgements.** The authors thank the members of the Recognition Physics Institute for their valuable feedback on the manuscript and the Lean framework.

## Statements and Declarations

**Funding.** The authors declare that no funds, grants, or other financial support were received for the research, authorship, or publication of this article.

**Competing interests.** The authors declare that they have no competing interests.

**Ethics approval and consent to participate.** Not applicable. This is a theoretical and mathematical study; it involves no human participants, animal subjects, or human-derived data.

**Consent for publication.** Not applicable.

**Data and code availability.** No datasets were generated or analyzed during this study. Every machine-checked statement cited in this paper is checked at source commit 6b7861b031f229ffe3be33f735c2db2fff3612e9 of the public Lean repository <https://github.com/jonwashburn/r> (MIT license), through the verification target and its import closure, under toolchain leanprover/lean4:v4.27.0-rc1 with mathlib4 pinned by the committed manifest; the build command is `lake build Paper`. The later metadata commit f31798f53fa6ce45659abcf66184327f1f357403 records this same paper-slice pin in the repository README and paper index; it is not used here as the code pin. An archival deposit with DOI will accompany the version of record. The pinned repository is a public verification snapshot and import closure; it is not the program’s full development library. The full library is larger and in places strictly stronger: it contains, at the time of writing, the smoothness instance corresponding to Appendix A. The theorem declarations for the circle-homology facts named in Section 9 are available in the pinned import closure, while Alexander duality and its wiring into the linking argument remain outside the pinned formalization. Claims in this paper are scoped to the pinned verification snapshot so that every cited formal theorem signature is checkable by one build command; where the full library is stronger, the text says so explicitly. One scope result is not a Lean artifact: the closure-independence witness of Section 8 is proved in Part I (Proposition 6.4). Three caveats apply. A theorem name is not evidence; the formal content is the exact statement and hypotheses at the pinned commit. Verification certifies only that the stated definitions imply the stated theorems, not physical identifications built on top of them. Finite arithmetic facts exposed by the verification target are rechecked by kernel decide; the two scope certificates of Section 8 are closed at their stated hypotheses. The smoothness typeclass assumed by the slice’s uniqueness theorem (AczelSmoothnessPackage) has no instance in the pinned slice; the slice omits it for dependency control. The instance is proved in the program’s full development library (Cost/AczelProof.lean, via the smoothness bootstrap theorem dAlembert\_contDiff\_top), and Appendix A of this paper reproduces the argument in prose. One naming note: the repository’s uniqueness theorem is named `law_of_logic_forces_jcost`. “Forces” there means forcing given the full hypothesis set (C1)–(C7), including the load-bearing closure law, exactly as in Theorem 7.1; Certificate 8.2 of this paper is the kernel-checked proof that the equality-induced floor cost does not satisfy the required composition consistency.

**Author contributions.** The authors contributed according to the CRediT taxonomy as follows. *Conceptualization:* J. Washburn. *Formal analysis:* J. Washburn, S. Pardo-Guerra, M. Simons, and A. Thapa. *Investigation:* J. Washburn, S. Pardo-Guerra, M. Simons, and A. Thapa. *Methodology:* J. Washburn, S. Pardo-Guerra, M. Simons, and A. Thapa. *Project administration:* S. Pardo-Guerra and M. Simons. *Software:* J. Washburn and A. Thapa. *Supervision:* J. Washburn. *Validation:* J. Washburn and A. Thapa. *Visualization:* S. Pardo-Guerra, M. Simons, and A. Thapa. *Writing (original draft):* J. Washburn. *Writing (review and editing):* J. Washburn, S. Pardo-Guerra, M. Simons, and A. Thapa. All authors read and approved the final manuscript.

## Appendix A Smoothness Bootstrap for d'Alembert

The main text uses the classical continuous classification of d'Alembert's equation. This appendix records the smoothness idea in a form sufficient for review. It also records a standard regularity strengthening often used in functional-equation theory: measurability plus local boundedness can replace continuity. The main theorem in this paper does not rely on that weaker route; it keeps continuity as the stated hypothesis.

### Step 1: smoothness bootstrap under continuity

Let  $H : \mathbb{R} \rightarrow \mathbb{R}$  be continuous,  $H(0) = 1$ , and

$$H(s+t) + H(s-t) = 2H(s)H(t).$$

Integrate in  $s$  over an interval  $[a, b]$ :

$$\int_a^b H(s+t) ds + \int_a^b H(s-t) ds = 2H(t) \int_a^b H(s) ds.$$

Changing variables in the two integrals gives

$$\int_{a+t}^{b+t} H(u) du + \int_{a-t}^{b-t} H(u) du = 2H(t) \int_a^b H(s) ds.$$

We claim there exists  $[a, b]$  with  $\int_a^b H(s) ds \neq 0$ . Since  $H$  is continuous and  $H(0) = 1 > 0$ , there exists  $\delta > 0$  such that  $H(s) > 0$  for all  $s \in (-\delta, \delta)$ . Setting  $a = 0$  and  $b = \delta/2$  gives  $\int_0^{\delta/2} H(s) ds > 0$ . Thus the required interval exists and  $I := \int_a^b H(s) ds > 0$ .

Dividing both sides of the integration identity by  $2I \neq 0$ ,

$$H(t) = \frac{1}{2I} \left( \int_{a+t}^{b+t} H(u) du + \int_{a-t}^{b-t} H(u) du \right).$$

The right side expresses  $H(t)$  as a sum of definite integrals of the continuous function  $H$  over intervals whose endpoints depend continuously (in fact, affinely) on  $t$ . By

the fundamental theorem of calculus applied to each sliding-window integral,  $H$  is differentiable, with

$$H'(t) = \frac{1}{2I} [H(b+t) - H(a+t) - H(b-t) + H(a-t)].$$

Since  $H$  is continuous, the difference  $H(b+t) - H(a+t) - H(b-t) + H(a-t)$  is a continuous function of  $t$ , so  $H'$  is continuous and  $H \in C^1(\mathbb{R})$ . The same identity now exhibits  $H'(t)$  as a finite sum of values of  $H$  at affine functions of  $t$ ; since  $H \in C^1$ ,  $H'$  is differentiable, so  $H \in C^2(\mathbb{R})$ . Iterating,  $H \in C^k(\mathbb{R})$  for every  $k \geq 1$ , hence  $H \in C^\infty(\mathbb{R})$ . Continuous solutions are therefore smooth, and the ODE classification in the main text applies.

Differentiating the d'Alembert equation twice in one variable and setting the other variable to zero yields

$$H''(t) = H''(0)H(t).$$

The solutions are constant, cosine, or hyperbolic cosine depending on the sign of  $H''(0)$ . The comparison-cost nondegeneracy (condition (C5) of Definition 6.2) selects the hyperbolic case with positive curvature, leaving the parameter  $H''(0) = \lambda^2 > 0$  until the unit log-curvature calibration axiom (C7) sets  $H''(0) = 1$ .

## Step 2: weakening continuity to measurability

The regularity hypothesis in the d'Alembert classification can be weakened from continuous to Lebesgue measurable and locally bounded, with no change in the conclusion. This is the standard strengthening of Aczél-class functional-equation theorems and is given in Aczél and Dhombres [2]; we record only the statement here and do not reproduce the proof.

**Lemma A.1** (Measurable d'Alembert classification, external). *Let  $H : \mathbb{R} \rightarrow \mathbb{R}$  be Lebesgue measurable and locally bounded, with  $H(0) = 1$  and*

$$H(s+t) + H(s-t) = 2H(s)H(t) \quad \forall s, t \in \mathbb{R}.$$

*Then  $H$  is continuous on  $\mathbb{R}$ , hence (by Step 1) smooth, and equals one of  $\cosh(\sqrt{\kappa} \cdot)$ , the constant function 1, or  $\cos(\sqrt{-\kappa} \cdot)$ , depending on the sign of  $\kappa = H''(0)$ .*

The main theorem of this paper (Theorem 7.1) assumes continuity, not measurability. The measurable route is consumed in exactly one place: Corollary 7.2, which weakens (C4) to measurability plus local boundedness and whose proof imports Lemma A.1 from the cited literature rather than proving it here. No other statement in the paper depends on this lemma.

## References

- [1] J. Aczél, *Lectures on Functional Equations and Their Applications*, Academic Press, 1966.

- [2] J. Aczél and J. Dhombres, *Functional Equations in Several Variables*, Encyclopedia of Mathematics and its Applications **31**, Cambridge University Press, 1989.
- [3] O. Hölder, “Die Axiome der Quantität und die Lehre vom Mass,” *Berichte über die Verhandlungen der Königlich Sächsischen Gesellschaft der Wissenschaften zu Leipzig, Math.-Phys. Classe* **53**, 1–64 (1901).
- [4] D. H. Krantz, R. D. Luce, P. Suppes, and A. Tversky, *Foundations of Measurement, Volume I: Additive and Polynomial Representations*, Academic Press, 1971.
- [5] J. Washburn, M. Zlatanović, and E. Allahyarov, “Recognition Geometry,” *Axioms* **15**(2), 90 (2026). <https://doi.org/10.3390/axioms15020090>.
- [6] J. Washburn and A. Rahnamai Barghi, “Reciprocal convex costs for ratio matching: Axiomatic characterization,” *Axioms* **15**(2), 151 (2026). <https://doi.org/10.3390/axioms15020151>.
- [7] J. Washburn and M. Zlatanović, “Uniqueness of the canonical reciprocal cost,” *Mathematics* **14**(6), 935 (2026). <https://doi.org/10.3390/math14060935>.
- [8] J. Washburn, M. Zlatanović, and E. Allahyarov, “The d’Alembert Inevitability Theorem,” *Mathematics* **14**(8), 1386 (2026). <https://doi.org/10.3390/math14081386>.
- [9] S. Pardo-Guerra, A. Thapa, M. Simons, and J. Washburn, “Coherent Comparison as Information Cost: Axiomatic Foundations for Discrete Ledger Dynamics,” *Foundations* **6**(2), 17 (2026). <https://doi.org/10.3390/foundations6020017>.
- [10] J. A. Wheeler, “Information, physics, quantum: The search for links,” in W. H. Zurek (ed.), *Complexity, Entropy, and the Physics of Information*, Addison-Wesley, 1990.
- [11] M. Tegmark, “The mathematical universe,” *Foundations of Physics* **38**, 101–150 (2008).
- [12] A. Connes, *Noncommutative Geometry*, Academic Press, 1994.
- [13] C. Rovelli, “Relational quantum mechanics,” *International Journal of Theoretical Physics* **35**, 1637–1678 (1996).
- [14] E. Verlinde, “On the origin of gravity and the laws of Newton,” *J. High Energy Phys.* **04** (2011), 029. [https://doi.org/10.1007/JHEP04\(2011\)029](https://doi.org/10.1007/JHEP04(2011)029).
- [15] G. Chiribella, G. M. D’Ariano, and P. Perinotti, “Informational derivation of quantum theory,” *Phys. Rev. A* **84**, 012311 (2011). <https://doi.org/10.1103/PhysRevA.84.012311>.

- [16] L. Masanes and M. P. Müller, “A derivation of quantum theory from physical requirements,” *New Journal of Physics* **13**, 063001 (2011). <https://doi.org/10.1088/1367-2630/13/6/063001>.
- [17] J. S. Bell, “On the Einstein Podolsky Rosen paradox,” *Physics* **1**, 195–200 (1964).
- [18] S. Kochen and E. P. Specker, “The problem of hidden variables in quantum mechanics,” *Journal of Mathematics and Mechanics* **17**, 59–87 (1967).
- [19] L. M. Bregman, “The relaxation method of finding the common point of convex sets and its application to the solution of problems in convex programming,” *USSR Computational Mathematics and Mathematical Physics* **7**(3), 200–217 (1967). [https://doi.org/10.1016/0041-5553\(67\)90040-7](https://doi.org/10.1016/0041-5553(67)90040-7).
- [20] I. Csiszár, “Information-type measures of difference of probability distributions and indirect observations,” *Studia Scientiarum Mathematicarum Hungarica* **2**, 299–318 (1967).
- [21] T. Gneiting and A. E. Raftery, “Strictly proper scoring rules, prediction, and estimation,” *Journal of the American Statistical Association* **102**(477), 359–378 (2007). <https://doi.org/10.1198/016214506000001437>.
- [22] G. Gonthier, “Formal proof—the four-color theorem,” *Notices of the American Mathematical Society* **55**(11), 1382–1393 (2008).
- [23] G. Gonthier, A. Asperti, J. Avigad, Y. Bertot, C. Cohen, F. Garillot, S. Le Roux, A. Mahboubi, R. O’Connor, S. O. Biha, I. Pasca, L. Rideau, A. Solovyev, E. Tassi, and L. Théry, “A machine-checked proof of the odd order theorem,” in *Interactive Theorem Proving*, Lecture Notes in Computer Science **7998**, Springer, 163–179 (2013). [https://doi.org/10.1007/978-3-642-39634-2\\_14](https://doi.org/10.1007/978-3-642-39634-2_14).
- [24] T. Hales et al., “A formal proof of the Kepler conjecture,” *Forum of Mathematics, Pi* **5**, e2 (2017). <https://doi.org/10.1017/fmp.2017.1>.
- [25] P. Scholze, “Liquid Tensor Experiment,” *Experimental Mathematics* **31**(2), 349–354 (2022). <https://doi.org/10.1080/10586458.2021.1926016>.
- [26] P. W. Bridgman, *Dimensional Analysis*, Yale University Press, 1922.
- [27] A. Hatcher, *Algebraic Topology*, Cambridge University Press, 2002. <https://pi.math.cornell.edu/~hatcher/AT/AT.pdf>.