foundations

MDPI

Article

Coherent Comparison as Information Cost: Axiomatic Foundations for Discrete Ledger Dynamics

Sebastian Pardo-Guerra *, Anil Thapa, Megan Simons and Jonathan Washburn

Recognition Physics Institute, Austin, TX 78701-1634, USA; athapa@recognitionphysics.org (A.T.); msimons@recognitionphysics.org (M.S.); jon@recognitionphysics.org (J.W.)
* Correspondence: sebas@recognitionphysics.org

Abstract

We develop an information-theoretic, cost-first framework for discrete dynamics in which the primitive operation is ratio-based comparison. Given two quantities compared via their ratio $x = a/b$, we assign a cost $F(x)$ measuring deviation from equilibrium ($x = 1$). Adopting a reciprocal d'Alembert composition law motivated by coherent chaining, together with quadratic calibration at unity, uniquely determines a reciprocal comparison cost $J(x) = \frac{1}{2}(x + x^{-1}) - 1$. Taking J as input, we model recognition events as deterministic updates on directed graphs recorded in a minimal ledger. Minimality (no intra-tick ordering metadata) together with non-commutativity of events implies atomic ticks: at most one event per tick. With conservation, pairwise locality, and quantization in $\delta\mathbb{Z}$, each event is recorded as a balanced double-entry posting. For graphs with cycles, assuming time-aggregated cycle closure over a finite clearing horizon, we show that cleared cycle closure is equivalent to path-independence and that the cumulative flow admits a scalar potential on each connected component (unique up to additive constant) via a discrete Poincaré lemma. On hypercube graphs Q_d, atomic single-edge updates impose a 2^d-tick minimal period for timestamp-unique coverage, realized by cyclic Gray codes (explicitly for $d = 3$). The framework links ratio-based cost functions, conservative graph flows, and discrete potential theory through explicitly stated axioms and structural assumptions.

Keywords: cost functional; d'Alembert equation; information ledger; discrete time

1. Introduction

What is the most primitive form of information processing? At its most basic level, any information-processing system must be able to *distinguish*—to tell one entity from another by registering difference and enabling comparison. We argue that this elementary act of comparison, when formalized under minimal and internally consistent requirements, already fixes far more structure than is usually assumed. Indeed, it uniquely determines a quantitative notion of cost that governs all subsequent dynamics.

In this framework, *comparison is the primitive operation*, and every comparison necessarily carries an intrinsic information-theoretic cost. Rather than postulating spacetime, fields, or particles as fundamental objects, we begin from a single operational question: if two quantities are compared via their ratio $x = a/b$ what form of comparison cost is selected once a coherent composition law, and local calibration are imposed?

This operational perspective builds on the axiomatic framework of Recognition Geometry [1], in which recognition is formalized as a mapping whose quotient structure gives rise to observable space, while that framework establishes the relevant topological foundations

Academic Editor: Vilém Novák

Received: 9 February 2026
Revised: 24 March 2026
Accepted: 6 May 2026
Published: 8 May 2026

of recognition, the present work shifts focus to its information-theoretic consequences. In particular, we investigate the cost functional that emerges from imposing coherence, composability, and consistency on comparison-based recognition processes.

Now, from a classical information theory perspective, the cost of encoding a distribution is measured by entropy or code-length [2,3], while the cost of distinguishing two distributions is measured by divergences such as Kullback–Leibler and its generalizations [4–7]. These notions are tightly linked to log-likelihood ratios, exponential families, and the geometry of statistical models [7,8]. We extend this perspective to *ratio-based comparison*: when comparing quantities a and b, the relevant observable is their ratio $x = a/b$, and the cost $F(x)$ quantifies an information-theoretic penalty for deviation from equilibrium ($x = 1$, perfect balance). This choice isolates a scale-free comparison primitive (multiplicative rather than additive), aligning naturally with log coordinates and likelihood-ratio thinking.

The perspective developed here is also relevant to a broader class of applied recognition problems, including those arising in artificial intelligence. In AI and machine learning, recognition systems—whether for image classification, speech processing, anomaly detection, expert aggregation, or model/sensor fusion—are routinely evaluated through divergence-based loss functions that quantify the cost of distinguishing one distribution from another [5,7]. The ratio $x = a/b$ of predicted to reference quantities, likelihoods, or confidence weights is precisely the operative observable in such systems, and measures of recognition quality are naturally expressed in terms of how far this ratio deviates from unity. The cost functional $J(x)$ derived below belongs to the f-divergence family [5] and is structurally related to the χ^2-divergence used in density-ratio estimation and hypothesis testing. In this sense, the coherence requirements imposed on comparison cost provide an axiomatic foundation for a specific class of divergence measures that appear organically in recognition tasks, though the present work remains foundational rather than algorithmic.

Further, this cost is not arbitrary. Coherent chaining of ratio comparisons motivates a composition law in which the direct comparison ratio $xy = a/c$ and the associated cross-ratio $x/y = ac/b^2$ enter symmetrically. Adopting this reciprocal d'Alembert composition law, together with normalization and quadratic calibration, yields the canonical reciprocal cost

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

(see ([9] Theorem 1)). In this sense, the cost functional constitutes the cornerstone of the framework: it is not postulated ad hoc, but fixed once the axioms are chosen.

Once the cost is established, we ask: how do recognition events propagate through a system? We model this via a *discrete ledger*—a sequential record of state updates. The ledger is not an additional structure; it emerges as the minimal way to encode recognition events under three constraints: (i) deterministic updates (the next state depends only on the current state and the event), (ii) minimality (no intra-tick ordering metadata), and (iii) lossless recording (no information is discarded). From these constraints, together with the non-commutativity of events (Axiom L2b) and additional structural assumptions (conservation, pairwise locality, quantization), we obtain: **atomic ticks** (at most one event per tick; Theorem T2), **double-entry postings** (balanced debit–credit pairs from conservation and pairwise locality), and **quantized units** (discrete postings in $\delta\mathbb{Z}$; Theorem T8). Recognition events thus induce a discrete dynamics on a directed graph, with each event posting a signed increment $\pm\delta$ on exactly two nodes.

Regarding graphs with cycles, the atomic single-edge events generically create transient circulation (net flux around closed loops). To recover a scalar potential representation—enabling path-independent summation of costs—we impose a *time-aggregated* cycle-closure (no-arbitrage/clearing) hypothesis: after netting flows over a finite clearing window, the cumulative flux around every cycle vanishes. Under this hypothesis, we prove that cycle closure is equivalent to path-independence (Theorem T3), and that the cumulative edge flow admits a unique scalar potential on each connected component, up to an additive constant (Theorem T4). This structure mirrors classical potential theory, but with explicit time-aggregation accounting for the impulse nature of atomic events. Lastly, when the recognition structure is a hypercube graph Q_d, atomicity imposes a minimal period constraint: to visit all 2^d vertices without repetition requires at least 2^d ticks (Theorems T6–T7). For $d = 3$, a Gray-code Hamiltonian cycle provides an explicit 8-tick realization.

The remainder of this paper proceeds as follows. Section 2 motivates the d'Alembert composition law from coherence principles and establishes the canonical cost functional as the unique solution. Section 3 develops the ledger framework, deriving atomic ticks (Theorem T2), double-entry postings, quantized units (Theorem T8), cycle-flux conservation (Theorem T3), scalar potentials (Theorem T4), and minimal period bounds (Theorems T6–T7). Section 4 synthesizes the framework and discusses its scope.

2. Motivation: From Coherent Comparison to the d'Alembert Composition Law

The framework rests on a simple but profound insight: if recognition involves comparison, and comparison has a cost, then coherent composition strongly constrains the admissible cost structure. This section explains why the reciprocal d'Alembert equation is a natural axiomatic choice for ratio-based comparison.

2.1. The Primacy of Comparison

At its most fundamental level, recognition is a relational act: one entity recognizes another. This recognition involves some form of comparison—measuring similarity, difference, or correspondence. Here, we formalize this by asking: if we compare two quantities by their ratio $x = a/b$, what "cost" or "defect" does this comparison incur?

Throughout this framework, we restrict to positive scalar quantities $a, b > 0$, so that the ratio $x = a/b$ takes values in $\mathbb{R}_{>0}$. This is a deliberate modeling choice, appropriate when a and b represent magnitudes—such as intensities, probabilities, lengths, or signal levels—for which only scale matters and not sign. The set $(\mathbb{R}_{>0}, \cdot)$ forms a multiplicative abelian group closed under inversion ($x \mapsto x^{-1}$, corresponding to reversing the comparison), and it admits the isomorphism $\ln : (\mathbb{R}_{>0}, \cdot) \to (\mathbb{R}, +)$, making log-coordinates canonical. These properties are what make ratio-based comparison and the reciprocity condition $F(x) = F(x^{-1})$ most naturally formulated on this domain. The restriction is not derived from the Recognition Geometry framework [1], which does not constrain the algebraic structure of configuration spaces; it is an explicit modeling commitment of the present work. Comparison of signed quantities, vectors, or elements of more general algebraic structures is outside the scope of this paper and constitutes a natural direction for future generalization.

Central to this framework is the idea that the comparison cost $F(x)$ should depend only on the ratio x itself, not on the absolute magnitudes of a and b. This reflects the intuition that recognition is fundamentally about *relationships*, not absolute values. Notably, the normalization $F(1) = 0$—perfect balance is free—is not imposed as a separate axiom but follows from the composition law itself: substituting $x = y = 1$ into Axiom A1 (formally stated in Section 3) forces $F(1) \in \{0, -1\}$, and non-negativity of F excludes the constant solution $F \equiv -1$ (see [9]).

https://doi.org/10.3390/foundations6020017

2.2. Coherent Composition: Motivating the d'Alembert Form

The goal of this subsection is to motivate—not derive—Axiom A1 (the d'Alembert composition law, stated formally in Section 3). The coherence arguments below show that the d'Alembert equation is the natural and self-consistent choice for a symmetric, ratio-based composition law; they do not constitute a proof that this form is the *only* possible one. Alternative choices (e.g., $F(xy) = F(x) + F(y)$, yielding a logarithmic cost) are possible and lead to different frameworks. The d'Alembert form is adopted as a primitive axiom precisely because it isolates reciprocal symmetry and, together with Axiom A2, uniquely fixes the cost functional J (Theorem T5).

The key insight motivating this choice comes from requiring that comparisons compose coherently. For instance, suppose we compare a to b, obtaining ratio $x = a/b$ with cost $F(x)$. Then, we compare b to c, obtaining ratio $y = b/c$ with cost $F(y)$. What should be the total cost when these comparisons are chained? We seek a composition law relating the cost of the direct comparison a to c (ratio $xy = a/c$) to the individual comparison costs $F(x)$ and $F(y)$. The *core coherence requirement* is that $F(xy)$ should be expressible as a function of $F(x)$ and $F(y)$ alone—i.e., $F(xy) = G(F(x), F(y))$ for some function G. This is the precise operational content of coherent chaining: the cost of the composed comparison is predictable from the costs of its constituent steps, without needing to know a, b, and c individually.

We then make a separate modeling choice to augment this requirement. Rather than constraining $F(xy)$ alone, we write a symmetric equation involving *both* $F(xy)$ and the cross-ratio cost $F(x/y)$, where $x/y = (a/b)/(b/c) = ac/b^2$:

$$F(xy) + F(x/y) = \Psi(F(x), F(y)) \tag{1}$$

for some symmetric function Ψ. This augmentation is *not* operationally required by the chaining argument above—it is a deliberate mathematical modeling choice. Including $F(x/y)$ alongside $F(xy)$ produces an equation that is symmetric under exchange $x \leftrightarrow y$ and under inversion $x \mapsto x^{-1}$, naturally enforces the reciprocity property $F(x) = F(x^{-1})$, and leads to the d'Alembert functional equation, whose solutions are well-characterized (see the Remark below and [10,11]) and whose connection to the cosine equation provides additional structural justification. The simpler requirement $F(xy) = G(F(x), F(y))$ alone—without the cross-ratio term—would yield a different and less constrained class of solutions, such as the logarithmic cost $F(xy) = F(x) + F(y)$, which does not encode reciprocal symmetry.

On the other end, we adopt the modeling choice of restricting the right-hand side to a symmetric bilinear function of $F(x)$ and $F(y)$:

$$\Psi(F(x), F(y)) = \alpha F(x)F(y) + \beta F(x) + \beta F(y) + \gamma, \tag{2}$$

where α, β, γ are constants to be determined. This is not the only consistent choice for Ψ, but it is the simplest one compatible with the following natural conditions:

- **Symmetry:** Interchanging $x \leftrightarrow y$ (equivalently $(xy, x/y) \mapsto (xy, y/x)$) leaves the equation invariant, forcing equal coefficients β for $F(x)$ and $F(y)$.
- **Interaction structure:** The quadratic term $\alpha F(x)F(y)$ captures nonlinear interaction between comparison costs, while the linear terms $\beta(F(x) + F(y))$ represent additive contributions.
- **Simplicity:** This is the most general low-order polynomial form respecting symmetry.

Now, to determine the coefficients, α, β, and γ, we impose consistency constraints:

Normalization consistency: Setting $x = y = 1$ in (1) with (2) gives

$$F(1) + F(1) = \alpha F(1)^2 + 2\beta F(1) + \gamma.$$

This, with $F(1) = 0$ (which follows from the composition law, as noted in Section 3), immediately yields $\gamma = 0$.

Scaling at unity: For $x = e^\epsilon$ with small ϵ, Axiom A2 (quadratic calibration) specifies that $F(e^\epsilon) \sim \epsilon^2/2 + O(\epsilon^4)$. Expanding both sides of (1) to second order in ϵ (with $y = e^\epsilon$ for symmetry), we obtain:

$$\text{LHS:} \quad F(e^{2\epsilon}) + F(1) = \frac{(2\epsilon)^2}{2} + 0 + O(\epsilon^4) = 2\epsilon^2 + O(\epsilon^4),$$
$$\text{RHS:} \quad \alpha\left(\frac{\epsilon^2}{2}\right)^2 + 2\beta\frac{\epsilon^2}{2} + O(\epsilon^4) = \beta\epsilon^2 + O(\epsilon^4).$$

Matching coefficients: $\beta = 2$.

Selecting α: Unlike γ and β, the coefficient α is not fixed by Axiom A2 alone within the present paper. Setting $y = x$ in (1) gives

$$F(x^2) = \alpha F(x)^2 + 4F(x),$$

which constrains the scaling of F under squaring, but does not uniquely fix α without additional input. The theory of iterated functional equations shows that $\alpha = 2$ is the unique value for which the resulting equation admits continuous solutions consistent with integer-power iteration $F(x^n)$ for all n [10,11]. This result lies outside the scope of the present paper; we adopt $\alpha = 2$ on the basis of this external result, treating it as a natural and well-supported selection rather than a derivation from coherence principles alone.

The considerations above—coherent chaining, symmetric bilinear interaction, and the selection $\alpha = 2$ from iterated functional equation theory—collectively motivate the following composition law, with $\alpha = 2$, $\beta = 2$, and $\gamma = 0$:

$$F(xy) + F(x/y) = 2F(x)F(y) + 2F(x) + 2F(y). \tag{3}$$

This is the d'Alembert functional equation in multiplicative form, which is adopted as Axiom A1 in Section 3. It encodes coherent composition: the cost of comparing a to c via intermediate steps ($F(xy) + F(x/y)$) is expressible as a specific combination of the individual costs $F(x)$ and $F(y)$. The form combines multiplicative interaction ($2F(x)F(y)$) with additive contributions ($2F(x) + 2F(y)$), reflecting that comparisons can reinforce each other nonlinearly while also accumulating linearly. For mathematical background on this functional equation, see [10,11].

Remark 1 (Connection to probability theory and stable distributions)**.** *In log coordinates $t = \ln x$, defining $G(t) = F(e^t)$, Axiom A1 reduces to the classical d'Alembert cosine equation $H(t+s) + H(t-s) = 2H(t)H(s)$, where $H = G + 1$, whose continuous solutions are $H(t) = \cosh(\alpha t)$. This equation is the same functional identity that governs the real parts of characteristic functions of symmetric stable distributions in probability theory [12,13]: many limit distributions arising in central-limit-type theorems are characterized precisely by functional equations of this form on their characteristic functions. Axiom A2 (quadratic calibration, $G''(0) = 1$) plays the role of a moment condition in this context, selecting $\alpha = 1$ and thus uniquely fixing $J(x) = \cosh(\ln x) - 1$, in the same way that moment or normalization conditions pick out a specific stable law from its functional equation.*

The formal statements of Axioms A1 and A2, and their roles in Theorem T5, are given in Section 3.

2.3. Calibration, Uniqueness, and Implications

The d'Alembert equation alone does not uniquely determine the cost function. We require one additional constraint to fix the scale. Axiom A2 (Quadratic calibration, formally stated in Section 3) specifies that in log coordinates $t = \ln x$, the cost has unit quadratic behavior at unity:

$$\lim_{t \to 0} \frac{2F(e^t)}{t^2} = 1.$$

If F is twice differentiable at unity, this is equivalent to normalizing the second derivative of the log-lift $G(t) = F(e^t)$ via $G''(0) = 1$. This condition fixes the overall scale of the cost.

Together, these two axioms (A1: Composition Law, A2: Quadratic calibration) uniquely determine the cost functional ([9]). This is the content of Theorem T5, the keystone theorem for the cost-first development, whose unique solution is:

$$J(x) = \frac{1}{2}(x + x^{-1}) - 1. \tag{4}$$

Now, once the cost functional is established, a cascade of implications follows. The cost function $J(x)$ has the property that $J(x) \geq 0$ with equality only when $x = 1$. In other words, perfect balance ($x = 1$) corresponds to zero cost, while any deviation incurs a positive penalty. Further, as $x \to 0^+$ or $x \to \infty$, the cost diverges: $J(x) \to \infty$. We formalize the corresponding boundary consequence as Theorem T1 (Boundary divergence / Meta-Principle) in the next subsection. The cost function also exhibits *reciprocity*: $J(x) = J(x^{-1})$ for all $x > 0$. This symmetry is compatible with representing recognition events in reversible pairs, leading to the balanced (double-entry) posting rule, obtained from explicit ledger-model assumptions (conservation, no sources/sinks, and pairwise-local events), as developed below.

3. Mathematical Framework

This section develops the full ledger-based mathematical framework underlying the results of this manuscript. The keystone cost-uniqueness theorem (Theorem T5) is stated here and proved in [9]. Taking this result as an input, we derive all subsequent ledger–structure results from explicit axioms and structural assumptions. Throughout, the logical structure adheres to a *cost-first* foundation, illustrated in Figure 1. We also provide the taxonomy of foundational elements in the cost-first framework in (Table 1).

Primitive Axioms (A1, A2) $\to$ **T5** (Cost Unique, Keystone)
$\to$ **Bal predicate** (zero-cost equilibrium) & **Exists predicate** (finite-cost admissibility)
$\to$ **T1** (boundary divergence / Meta-Principle)
$\to$ **Ledger Structure** $\to$ **T2** (Atomic Tick)
$\to$ **T8** (Ledger Units) $\to$ **T3** (Cycle flux; under time-aggregated cycle closure) $\to$ **T4** (Potential)
$\to$ **T6–T7** (Minimal period) $\to$ **D=3** (conditional).

Figure 1. High-level cost-first dependency chain used in this manuscript. Arrows represent chronological and conceptual priority: the cost functional J is established first and motivates the ledger framework, but the ledger axioms (L1–L2b) and structural assumptions are logically independent of A1–A2 and would yield the same ledger results for any reasonable cost function. The specific properties of J (reciprocity, boundary divergence, unique zero) connect to the ledger framework conceptually—through the Bal and Exists predicates and the compatibility of double-entry with $J(x) = J(x^{-1})$—rather than derivationally.

Table 1. Taxonomy of foundational elements in the cost-first framework. Status: **Axiom** = primitive assumption; **Assumption** = structural modeling choice; **Derived** = proven from axioms; **Conditional** = requires additional hypotheses.

Category	**Elements**	**Status**
Primitive Cost Axioms	A1 (Composition Law): $F(xy) + F(x/y) = 2F(x)F(y) + 2F(x) + 2F(y)$	Axiom
	A2 (Quadratic calibration): $\lim_{t\to 0} \frac{2F(e^t)}{t^2} = 1$	Axiom
Ledger Axioms	L1: Deterministic state-update semantics ($S_{t+1} = U(S_t, \sigma_t)$)	Axiom
	L2: Minimality of ledger structure (no ordering metadata)	Axiom
	L2b: Non-commutativity of events	Axiom
Derived Theorems	T5: Cost Uniqueness: $J(x) = \frac{1}{2}(x + x^{-1}) - 1$	Derived (A1, A2)
	Perfect balance: $\mathrm{Bal}(x) \iff J(x) = 0 \iff x = 1$	Derived (from T5)
	Finite-cost admissibility: $\mathrm{Exists}(x) \iff J(x) < \infty$	Derived (from T5)
	T1 (Boundary divergence): $J(0^+) \to \infty$	Derived (from T5)
	T2: Atomic Tick (at most one event per tick)	Derived (L1 + L2 + L2b)
	T3: Equivalence of cycle closure and path-independence	Derived (conditional)
	T4: Potential Uniqueness for cleared cumulative flow	Derived (from T3)
	T6: Minimal period 2^d (eight ticks for $d = 3$)	Derived (from T2)
	T7: Coverage Lower Bound (pigeonhole principle)	Derived (from T6)
	T8: Ledger Units (algebraic structure $\Delta \simeq \mathbb{Z}$)	Derived (trivially)
Structural Assumptions	Conservation principle: Total balance invariant per tick	Assumption
	No external sources/sinks	Assumption
	Pairwise locality of events	Assumption
	Quantization in $\delta\mathbb{Z}$ with no torsion	Assumption
	Time-aggregated cycle closure (clearing/no-arbitrage)	Assumption
Conditional Results	Dimension selection $d = 3$ (linking + gap-45 + $2^d = 8$)	Conditional
Definitions	Recognition event: $(a, b) \in A \times B$	Definition
	Ledger state: $S_t \in \mathcal{S}$	Definition
	Tick: Minimal temporal unit for one state update	Definition
	Per-tick edge increment: $\delta\Delta(e, t) \in \delta\mathbb{Z}$	Definition
	Recognition structure: Directed graph $G = (X, E)$	Definition

3.1. The Primitive Foundation: Cost Functional and Uniqueness (T5)

We begin with the keystone theorem establishing the unique cost functional, which serves as the foundational input for all subsequent constructions. This result isolates the minimal structure required to fix the functional form of the cost and precedes the introduction of any discrete dynamics. Accordingly, we distinguish between two classes of axioms. The first are the *cost axioms* (Axioms A1–A2), which determine the form of the cost functional J through coherence and consistency requirements on comparison. The second are the *ledger axioms* (Axioms L1–L2), introduced in Section 3.3, which govern the discrete update structure once the cost is fixed. This separation clarifies the logical architecture of the framework: the cost axioms determine *what* is being minimized, while the ledger axioms determine *how* that structure evolves.

We now state the two cost axioms:

Axiom A1 (Composition Law). For all $x, y > 0$:

$$F(xy) + F(x/y) = 2F(x)F(y) + 2F(x) + 2F(y). \quad (5)$$

This is the d'Alembert functional equation in multiplicative form. As motivated in Section 2, this form is the natural symmetric, bilinear composition law for ratio-based costs: it respects the multiplicative structure of ratios, supports reciprocal symmetry, and is consistent with coherent chaining of comparisons. The supporting arguments for this choice—including the selection $\alpha = 2$ from the theory of iterated functional Equations [10,11]—are presented as motivation in Section 2, not as a derivation from A2 alone.

Axiom A2 (Quadratic calibration). In log coordinates $t = \ln x$, define $G(t) = F(e^t)$. We require

$$\lim_{t\to 0} \frac{2G(t)}{t^2} = \lim_{t\to 0} \frac{2F(e^t)}{t^2} = 1.$$

The latter fixes the overall scale (and, when G is twice differentiable at 0, coincides with $G''(0) = 1$).

Considering the above, we have the next:

Theorem 1 (Theorem T5: Cost Uniqueness)**.** *Let $F : \mathbb{R}_{>0} \to \mathbb{R}$ satisfy:*

1. ***Composition Law (A1):*** $F(xy) + F(x/y) = 2F(x)F(y) + 2F(x) + 2F(y)$
2. ***Quadratic calibration at unity (A2):*** $\lim_{t\to 0} \frac{2F(e^t)}{t^2} = 1$

Then, F is uniquely determined:

$$F(x) = \frac{1}{2}(x + x^{-1}) - 1. \tag{6}$$

We denote this unique cost functional by J.

Proof. See ([9] Theorem 1) for a self-contained proof. □

Some of the key properties of the cost functional are:

- **Reciprocity:** $J(x) = J(x^{-1})$ for all $x > 0$;
- **Non-negativity:** $J(x) \geq 0$ with equality iff $x = 1$;
- **Normalization**: $F(1) = 0$;
- **Boundary divergence:** $J(x) \to \infty$ as $x \to 0^+$ or $x \to \infty$.

In particular, the boundary divergence is *not assumed*—it is a consequence of the unique functional form. This is why the Meta-Principle is derived, not primitive.

Remark 2 (Ulam–Hyers stability of the cost functional)**.** *The uniqueness established in Theorem T5 raises a natural companion question: if a function $\widetilde{F}$ approximately satisfies Axiom A1 (within some tolerance $\varepsilon > 0$), is $\widetilde{F}$ necessarily close to J? This is the* Ulam–Hyers stability problem *for the d'Alembert Equation [10,11]: a function is stable if approximate solutions are always near exact ones. Stability results of this type would provide robustness guarantees for the cost-first framework under perturbation—ensuring that recognition processes which only approximately satisfy the coherence requirement still produce costs close to J—and the connection to the stability theory of functional equations arising in limit theorems of probability theory [12,13] makes this a natural avenue for future investigation.*

3.1.1. Unique Zero-Cost Configuration (Law of Existence)

The unique cost functional immediately yields the *Law of Existence*:

Definition 1 (Perfect Balance predicate)**.** *A configuration $x > 0$* is in perfect balance *(in the present sense) if and only if its defect collapses to zero, that is,*

$$Bal(x) \iff J(x) = 0. \tag{7}$$

The predicate $\mathrm{Bal}(x)$ is a *technical consistency predicate* in this framework, not a claim about ontological presence. Note that only $x = 1$ satisfies $\mathrm{Bal}(x)$ (zero defect). Indeed, starting from the definition of J:

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

we have that

$$\begin{aligned} J(x) = 0 &\iff \frac{1}{2}\left(x + x^{-1}\right) - 1 = 0 \\ &\iff x + x^{-1} = 2 \\ &\iff x^2 - 2x + 1 = 0 \qquad \text{(multiply by } x > 0) \\ &\iff (x-1)^2 = 0 \\ &\iff x = 1. \end{aligned}$$

In this way, configurations with $x \neq 1$ have $J(x) > 0$ and are *recognizable* precisely because their nonzero defect is quantifiable and enables comparison and composition via the ledger rules.

Definition 2. *We say that a configuration* exists *in this framework if and only if it satisfies finite-cost admissibility (existence in the present sense).*

$$\text{Exists}(x) \;:\!\iff\; J(x) < \infty.$$

Hence, every configuration $x > 0$ *exists* in this sense, as it has finite cost. In contrast, the boundary regimes $x \to 0^+$ and $x \to \infty$ correspond to cost blow-up and are thus excluded from admissible configurations. In the next subsection, we show that these divergent limits admit a natural interpretation as "nothingness" within the present framework, marking the breakdown of recognizability rather than the presence of any physical or informational configuration.

3.1.2. Properties of the Cost Function

Near equilibrium ($x = 1$), the cost function exhibits quadratic behavior. Indeed, let $x = e^{\epsilon}$ for small ϵ. Then,

$$J(e^{\epsilon}) = \frac{1}{2}(e^{\epsilon} + e^{-\epsilon}) - 1 = \cosh(\epsilon) - 1 = \frac{\epsilon^2}{2} + \frac{\epsilon^4}{24} + \cdots \approx \frac{1}{2}\epsilon^2, \tag{8}$$

locally approximating one-half the squared log-distance: $J(e^{\epsilon}) \approx \frac{1}{2}\epsilon^2 = \frac{1}{2}d_{\log}^2$, where $d_{\log}(a,b) = |\ln a - \ln b|$ is the Riemannian distance in log-coordinates. This quadratic local behaviour is characteristic of divergences in information geometry [7]: J is not itself a metric (it is a single-argument cost function, or equivalently a two-argument function in which numerator and denominator play asymmetric roles, and no triangle inequality is established), but its leading-order expansion near equilibrium recovers the squared log-scale distance, consistent with J belonging to the f-divergence family [5]. This local quadratic structure ensures well-behaved cost accumulation near equilibrium.

Remark 3 (Self-similar scale and the golden ratio)**.** *The following observation is illustrative and stands apart from the deductive development; it is not used in any subsequent result. Consider now the recurrence equation $x_{n+1} = 1 + 1/x_n$ as a simple self-similar update rule. The fixed points satisfy $x = 1 + 1/x$, yielding the quadratic equation*

$$x^2 - x - 1 = 0 \quad \Rightarrow \quad \phi = \frac{1 + \sqrt{5}}{2} \approx 1.618.$$

At ϕ, the additive (self) and reciprocal (other) components balance. Further, the recognition cost evaluates to

$$J(\phi) = \frac{1}{2}\left(\phi + \frac{1}{\phi}\right) - 1 = \phi - \frac{3}{2} \approx 0.118,$$

and by reciprocity, $J(\phi) = J(\phi^{-1})$, so both ϕ and its reciprocal $1/\phi \approx 0.618$ lie at the same cost. Thus, under this update rule, ϕ marks a natural self-similar scale at which the cost function exhibits a special symmetry, reflecting an exact balance between additive and reciprocal contributions. This symmetry identifies ϕ as a distinguished reference scale intrinsic to ratio-based comparison, even in the absence of any assumed dynamics. Once such a reference point is identified, it is natural to reparametrize deviations multiplicatively relative to this scale, thereby passing from ratio space to logarithmic coordinates. Within this parametrization, the quantity $J_{bit} = \ln \phi \approx 0.481$ emerges as a convenient log-scale reference associated with the self-similar fixed point. In applications, it may be interpreted as a characteristic scale for multiplicative deviations measured in log-coordinates. However, no claim is made here that ϕ (or J_{bit}) is forced without additional dynamical or self-similarity hypotheses beyond the cost axioms. More generally, each metallic ratio defined by $x^2 - nx - 1 = 0$ ($n \geq 1$) is a fixed point of the recurrence $x_{n+1} = n + 1/x_n$ and yields an analogous cost value $J(\phi_n) = \phi_n - (1 + n/2)$; the singling out of $\phi = \phi_1$ here reflects the specific recurrence chosen, not a selection forced by the cost axioms.

Example 1 (Recognition cost across regimes)**.** *To ground the cost functional concretely, consider a recognition system comparing a measured signal level a to a reference level b. The fidelity ratio $x = a/b$ quantifies how closely the measured signal matches the reference: $x = 1$ corresponds to perfect match (zero cost), $x < 1$ to a weak or attenuated signal, and $x > 1$ to an amplified or overconfident signal. Table 2 evaluates $J(x) = \frac{1}{2}(x + x^{-1}) - 1$ across representative regimes.*

The quadratic approximation $J(x) \approx \frac{1}{2}(\ln x)^2$ near $x = 1$ can be verified numerically: for $x = 0.9$, $\frac{1}{2}(\ln 0.9)^2 = \frac{1}{2}(0.1054)^2 \approx 0.006$, matching the tabulated value. In an AI classification context, setting $x = p$ (predicted confidence for the true class, with reference $p_{\text{ref}} = 1$) shows how J assigns a negligible penalty to confident correct predictions ($p \approx 1$, $J \approx 0$) and an escalating, unbounded penalty as confidence collapses toward the rare-event boundary ($p \to 0^+$).

Table 2. Evaluation of $J(x) = \frac{1}{2}(x + x^{-1}) - 1$ across representative fidelity regimes. Reciprocity $J(x) = J(x^{-1})$ is visible in the paired rows: underconfidence and overconfidence incur equal cost at reciprocal ratios. Rare-event regimes ($x \to 0^+$ or $x \to \infty$) correspond to unbounded cost (Theorem T1).

$x = a/b$	Regime	$J(x)$	Notes
0.01	Rare event/near-absent signal	49.005	$J \to \infty$ as $x \to 0^+$ (Theorem T1)
0.10	Severe mismatch	4.050	
0.50	Moderate deviation	0.250	
0.90	Near-perfect (weak)	0.006	Quadratic regime: $\frac{1}{2}(\ln x)^2 \approx 0.006$
1.00	Perfect balance	0	Unique zero (Theorem T5)
$1/0.9 \approx 1.11$	Near-perfect (strong)	0.006	$J(x) = J(x^{-1})$: equal cost to $x = 0.9$
2.00	Moderate deviation (over)	0.250	$J(x) = J(x^{-1})$: equal cost to $x = 0.5$
10.0	Severe overconfidence	4.050	$J(x) = J(x^{-1})$: equal cost to $x = 0.1$
100	Rare event/extreme overconfidence	49.005	$J \to \infty$ as $x \to \infty$ (Theorem T1)

3.2. Boundary Divergence

The first direct consequence of establishing the cost functional J in Theorem T5 is its divergence at the boundary limits. This behavior is not imposed by assumption, but is a consequence of the structure of the derived cost functional.

Remark 4 (Relation to the recognition quotient in Recognition Geometry)**.** *At its most fundamental level, recognition is a relational act: one entity recognizes another. As formalized in [1], a recognizer $R : \mathcal{C} \to \mathcal{E}$ partitions the configuration space into indistinguishable equivalence classes (resolution cells), via the induced equivalence relation $c \sim c' \iff R(c) = R(c')$, and the observable space is the corresponding quotient $\mathcal{C}/\sim$. From this viewpoint, a single recognition act is naturally represented as an ordered pair $(c, R(c)) \in \mathcal{C} \times \mathcal{E}$, i.e., an element of a Cartesian* *product, and the* realized *recognitions form the graph of R, a subset of* $\mathcal{C} \times \mathcal{E}$*. In this way, one can also interpret a* recognition event *as an ordered pair* $(a, b) \in A \times B$*, where set A is the recognizer and set B the recognized. In this case, we write* $\text{Recognition}(A, B) = A \times B$ *for the set of all recognition events from set A to set B. If either set is empty, then* $\text{Recognition}(A, B) = \emptyset$*.*

Theorem 2 (Theorem T1: Boundary divergence)**.** *The cost function diverges at the boundaries:*

$$\lim_{x \to 0^+} J(x) = +\infty \quad \text{and} \quad \lim_{x \to \infty} J(x) = +\infty. \tag{9}$$

These limits lie outside the finite-cost regime of the model.

Proof. From Theorem T5, $J(x) = \frac{1}{2}(x + x^{-1}) - 1$. Rewrite this as

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

As $x \to 0^+$, the term $\frac{1}{2x} \to +\infty$ while $\frac{x}{2} \to 0$ and the constant -1 remains finite. Therefore:

$$\lim_{x \to 0^+} J(x) = +\infty.$$

By the reciprocity property $J(x) = J(x^{-1})$, the limit $x \to \infty$ is equivalent to $x^{-1} \to 0^+$, hence $\lim_{x \to \infty} J(x) = +\infty$ as well. □

Remark 5 (Set-theoretic interpretation of the $x \to 0^+$ limit)**.** *If we interpret* $x = a/b$ *where a and b represent availability or cardinality measures* $a := \mu(A)$ *and* $b := \mu(B)$ *(with* $\mu(\emptyset) = 0$ *and* $\mu(\cdot) > 0$ *for nonempty domains), then the limit* $x \to 0^+$ *with b fixed corresponds to* $\mu(A) \to 0^+$*, i.e., the recognizer domain A approaching emptiness. In the limiting case* $\mu(A) = 0$ *(equivalently* $A = \emptyset$*), the Cartesian product* $A \times B = \emptyset$ *contains no recognition events, and hence,* $\text{Recognition}(\emptyset, B) = \emptyset$*. Thus, the infinite-cost divergence excludes configurations in which one domain vanishes, in accordance with the observation that encoding or computing absolute absence requires infinite information. Empty sets do not fail by contradiction; rather, they lie beyond finite computational accessibility. This last provides a consistency check: the cost structure intrinsically penalizes degenerate configurations with unbounded informational cost, enforcing exclusion at the level of computability rather than assumption.*

In the language of applied recognition, configurations near $x \to 0^+$ or $x \to \infty$ correspond to *rare-event* regimes, where one side of the comparison becomes vanishingly small relative to the other—as in anomaly detection, extreme-value recognition, or fault identification tasks. The divergence of J in these limits provides a cost-theoretic account of rare-event detection: rare or anomalous configurations carry unbounded recognition cost and are thereby excluded from the finite-cost admissible regime, consistent with the Exists predicate ($J(x) < \infty$). This connects Theorem T1 directly to the applied recognition challenge of distinguishing rare events from nominal ones, since the cost framework assigns a principled, quantitative penalty that grows without bound as the configuration approaches the rare-event boundary.

3.3. Single-Event Updates and Atomic Ticks

The ledger axioms introduced in this subsection (L1, L2, L2b) are logically independent of the cost axioms A1–A2: the atomic-tick theorem (T2) and all subsequent ledger results would hold for any cost function, not specifically for $J(x) = \frac{1}{2}(x + x^{-1}) - 1$. The connection between J and the ledger is conceptual rather than derivational: J's reciprocity $J(x) = J(x^{-1})$ is compatible with the double-entry posting structure (reversing an event's

https://doi.org/10.3390/foundations6020017

orientation leaves cost unchanged); its unique zero at $x = 1$ defines the equilibrium reference of the Bal predicate; and its boundary divergence grounds the Exists predicate by excluding degenerate configurations. These connections motivate the joint framework but do not derive the ledger structure from J.

The atomic-tick principle results from requiring that recognition events be recorded unambiguously. To that end, the way we model recognition dynamics is through a *ledger*: a sequential record updated once per tick. Later in the paper, when a spatial carrier is assumed, we use hypercubic graphs Q_d as a convenient family for analyzing scheduling and coverage constraints; the arguments below do not depend on a specific embedding of the ledger in $\mathbb{R}^d$.

In order to provide a minimal encoding of recognition events, we adopt a ledger structure subject to explicit constraints: determinism (Axiom L1), minimality (Axiom L2), non-commutativity (Axiom L2b), and losslessness (no information discarded). Let $\mathcal{L}$ denote the configuration space of all ledger states (formalized below as balance functions on a recognition graph). In the sense of Recognition Geometry [1], $\mathcal{L}$ plays the role of a configuration space $\mathcal{C}$, and a choice of recognizer (readout) map $R : \mathcal{L} \to \mathcal{E}$ induces a recognition quotient that identifies indistinguishable ledger microstates. This makes $\mathcal{L}$ a *Recognition Structure* in the precise sense of ([1] Example 2.8).

Axiom L1 (Deterministic State-Update Semantics). The ledger state S_t at tick t evolves deterministically according to a function $U : \mathcal{S} \times \mathcal{E}^* \to \mathcal{S}$, where $\mathcal{S}$ is the state space, $\mathcal{E}$ is the set of recognition events, and $\mathcal{E}^*$ denotes finite sequences of events (including the empty sequence). The state update rule is:

$$S_{t+1} = U(S_t, \sigma_t), \tag{10}$$

where $\sigma_t \in \mathcal{E}^*$ is a finite sequence of recognition events at tick t (possibly empty, possibly containing multiple events). The function U is deterministic: for fixed S_t and σ_t, the resulting state S_{t+1} is uniquely determined.

Axiom L2 (Minimality of Ledger Structure). The ledger records only final states at each tick: S_t contains no event-ordering metadata beyond the tick index itself. Equivalently, the ledger does not retain any intra-tick permutation information: if multiple events were processed within a tick, unambiguous recording would require $U(S, \sigma)$ to depend only on the *unordered* content of σ (i.e., to be invariant under permutations of σ).

Axiom L2b (Non-commutativity of Events). Recognition events do not commute in general. Formally, there exist states $S \in \mathcal{S}$ and event sequences $\sigma, \sigma' \in \mathcal{E}^*$ containing the same events in different orders such that

$$U(S, \sigma) \neq U(S, \sigma').$$

Remark 6 (Justification for non-commutativity)**.** *This assumption reflects the directed, asymmetric nature of recognition: the act "A recognizes B" followed by "B recognizes C" may produce a different configuration than performing these recognition acts in reverse temporal order. The non-commutativity captures the essential irreversibility and causal structure of recognition events. Without this assumption, our framework would permit concurrent events within a single tick, leading to fundamentally different dynamics.*

Theorem 3 (Theorem T2: Atomic tick)**.** *Under Axioms L1, L2, and L2b, at most one recognition event is processed per tick. There are no concurrent recognitions.*

Proof. Suppose, for contradiction, that a single tick t may process a sequence $\sigma_t \in \mathcal{E}^*$ containing at least two events.

By Axiom L1 (deterministic updates), the post-tick state is $S_{t+1} = U(S_t, \sigma_t)$, uniquely determined by S_t and σ_t. By Axiom L2 (minimality), the ledger carries no intra-tick ordering metadata beyond the tick index itself. Therefore, if multi-event ticks are allowed, the update rule must be *well-defined on the unordered content of a tick*: for every state S and every finite event sequence σ, permuting the order within the tick cannot change the resulting recorded state. Formally, for every permutation π of the events in σ, we must have

$$U(S, \sigma) = U(S, \pi(\sigma)).$$

This permutation-invariance is necessary because the ledger provides no mechanism to distinguish which permutation actually occurred.

However, Axiom L2b (non-commutativity) asserts that events do not commute in general: there exist a state S and two sequences σ, σ' containing the same events in different orders such that $U(S, \sigma) \neq U(S, \sigma')$. This directly contradicts the permutation-invariance required above.

Therefore, to simultaneously satisfy deterministic update semantics (Axiom L1), minimal recording without ordering metadata (Axiom L2), and non-commutativity of recognition events (Axiom L2b), each tick must contain at most one event. This proves atomicity. □

Theorem T2 establishes discrete temporal order: time advances in atomic steps. The proof shows that atomicity is a necessary consequence of the minimality constraint (Axiom L2) when combined with the fact that events do not commute in general. As a corollary, we can restrict the domain of U to $\mathcal{S} \times \mathcal{E}$ (single events) rather than $\mathcal{S} \times \mathcal{E}^*$ (sequences), since sequences of length greater than one are excluded by T2.

To make subsequent statements (atomic ticks, quantized postings, and cycle flux) mathematically precise, we adopt the following minimal ledger semantics.

Definition 3 (Recognition structure)**.** *Fix a directed graph $G = (X, E)$, where X is the set of nodes and $E \subseteq X \times X$ is the set of directed edges. We assume E is closed under reversal: if $(u \to v) \in E$, then $(v \to u) \in E$.*

Definition 4 (Ledger state as balances)**.** *Fix a nonzero increment $\delta > 0$. A ledger state at tick t is a balance function*

$$S_t \equiv b_t : X \longrightarrow \delta\mathbb{Z}.$$

We write the total balance as

$$\mathcal{B}(S_t) := \sum_{x \in X} b_t(x),$$

assuming this sum is well-defined (e.g., X finite, or b_t has finite support).

Definition 5 (Event-to-posting map)**.** *By Theorem T2, at most one recognition event occurs per tick. Therefore, we can write the state update as $S_{t+1} = U(S_t, e_t)$ where $e_t \in E$ is a single oriented edge (or the empty event if no recognition occurs). Given this deterministic update rule, define the induced* node postings *(balance increments) as*

$$\mathsf{Post}(S_t, e_t)(x) := \big(U(S_t, e_t)\big)(x) - S_t(x) \in \delta\mathbb{Z}.$$

Equivalently, $S_{t+1} = S_t + \mathsf{Post}(S_t, e_t)$ as functions $X \to \delta\mathbb{Z}$.

Definition 6 (Edge postings and cycle flux)**.** *Under the pairwise-locality assumption introduced below (so that only the endpoints (u,v) can change at tick t), the posting is determined by a single magnitude*

$$\Delta_t \ := \ \mathsf{Post}(S_t, e_t)(v) \ = \ -\mathsf{Post}(S_t, e_t)(u) \in \delta\mathbb{Z}.$$

We then define the corresponding per-tick edge increment $\delta\Delta(\cdot, t) : E \to \delta\mathbb{Z}$ *by*

$$\delta\Delta(u \to v, t) = \Delta_t, \qquad \delta\Delta(v \to u, t) = -\Delta_t, \qquad \delta\Delta(e, t) = 0 \text{ for all other } e \in E.$$

This $\delta\Delta(\cdot, t)$ is a sparse 1*-cochain (supported on at most two directed edges) induced by the single atomic event at tick t.*

Fix an integer clearing horizon $W \geq 1$. For any start time t_0, define the cumulative edge flow *over the window $[t_0, t_0 + W)$ by*

$$\overline{\Delta}_{t_0,W}(e) \ := \sum_{\tau=t_0}^{t_0+W-1} \delta\Delta(e, \tau) \in \delta\mathbb{Z}.$$

For any directed cycle $\gamma = (e_1, \ldots, e_n)$ in G, define the cumulative cycle flux

$$\overline{\Phi}(\gamma; t_0, W) \ := \ \sum_{i=1}^{n} \overline{\Delta}_{t_0,W}(e_i).$$

Remark 7 (Relation to more general flow formalisms)**.** *More general finitary flow formalisms (local finiteness, inflow/outflow sums, and "closed-chain sums") can be used to relate event structure to conservation and exactness on large or infinite graphs. The definitions above are the minimal finite-support specialization needed for the present manuscript.*

Balanced Postings (Double-Entry)

Theorem T2 establishes *atomicity* but not the posting *structure*. We now show that, under explicit structural assumptions (conservation, no sources/sinks, and pairwise-local events), each recognition event must be recorded as a balanced debit–credit pair (double-entry). The reciprocity property $J(x) = J(x^{-1})$ is then naturally compatible with reversing the orientation of an event without changing its cost.

Structural Assumption: Conservation Principle. The total ledger balance is invariant at each tick: if $\mathcal{B}(S_t)$ denotes the total balance (sum over all nodes) of state S_t, then $\mathcal{B}(S_{t+1}) = \mathcal{B}(S_t)$ for all t.

Structural Assumption: No External Sources or Sinks. Postings are the only state-changing operations. There are no auxiliary fields, external flows, or hidden variables that can absorb or supply balance.

Structural Assumption: Pairwise locality of events. Each recognition event e_t designates an ordered pair of nodes (u, v), and the update $S_{t+1} = U(S_t, e_t)$ can change balances only at those two nodes.

Remark 8. *Why pairwise locality is required for "exactly two postings". Conservation alone implies only that the* net *change in total balance per tick is zero. Without a locality condition, the deterministic update rule $U(S_t, e_t)$ could (in principle) redistribute balance across many nodes in response to a single event input e_t, while still preserving the total sum. Therefore, the "exactly two postings" conclusion requires a modeling commitment that a recognition event is* pairwise *at the ledger level (an ordered pair of nodes, i.e., a directed edge), so that only the event's endpoints may change on that tick.*

https://doi.org/10.3390/foundations6020017

Proposition 1 (Double-entry constraint)**.** *Under the following assumptions:*

1. *Atomicity: At most one recognition event per tick (Theorem T2).*
2. *Conservation: Total balance is invariant per tick.*
3. *No external sources/sinks: Postings are the only balance-changing operations.*
4. *Self-contained state updates: The state S_{t+1} depends only on S_t and the recognition event e_t (Axiom L1, with atomicity from Theorem T2).*
5. *Pairwise locality of events (structural assumption above).*

Each recognition event must be self-balancing: it records exactly two postings of equal magnitude and opposite sign on the participating nodes, $+\Delta_t$ and $-\Delta_t$ (credit and debit). If postings are quantized in $\delta\mathbb{Z}$, then $\Delta_t \in \delta\mathbb{Z}$.

Proof. Fix a tick t and suppose a recognition event e_t occurs (if no event occurs, the claim is vacuous). By pairwise locality, only two node balances can change in passing from S_t to S_{t+1}; call them u and v. By conservation, the net change in total balance is zero, so the balance change at u must be the negative of the balance change at v. Writing the change magnitude as Δ_t, the event therefore records exactly two opposite postings, $-\Delta_t$ at u and $+\Delta_t$ at v. □

Combining the preceding proposition with the assumption of quantization in $\delta\mathbb{Z}$, each recognition event is recorded as a balanced pair of postings $+\Delta_t$ and $-\Delta_t$, where $\Delta_t \in \delta\mathbb{Z}$. Under this structure, the reciprocity property $J(x) = J(x^{-1})$ guarantees that reversing an event's orientation yields a recognition event of equal cost. To illustrate the preceding ideas, we recall the definition of a *recognition structure* (Definition 3) and consider the following example:

Example 2 (Recognition structure)**.** *Consider a recognition structure $G = (X, E)$ with four nodes $X = \{a, b, c, d\}$ and directed edges representing recognition relations:*

- *$a \to b$: node a recognizes node b.*
- *$b \to c$: node b recognizes node c.*
- *$c \to d$: node c recognizes node d.*
- *$d \to a$: node d recognizes node a.*
- *$a \to c$: node a recognizes node c.*

This forms a directed graph with a cycle $(a \to b \to c \to d \to a)$ and an additional edge $(a \to c)$ creating a shortcut. At each tick t, the ledger assigns a (generally sparse) per-tick edge increment *$\delta\Delta(e, t) \in \delta\mathbb{Z}$ to each directed edge e. For instance, at tick $t = 1$, we might have:*

$$\begin{aligned} \delta\Delta(a \to b, 1) &= +\delta, \\ \delta\Delta(b \to a, 1) &= -\delta, \\ \delta\Delta(e, 1) &= 0 \quad \textit{for all other edges } e \in E. \end{aligned}$$

The double-entry rule ensures that for each recognition event, if the event induces a $+\delta$ increment along $(x \to y)$, then the node balances record a matched debit/credit pair, maintaining total-balance conservation.

3.4. Quantized Posting Units (δ)

The atomic-tick structure with double-entry raises a fundamental question: what is the minimal unit δ? We answer this question by adopting the requirement that ledger postings take values in $\delta\mathbb{Z}$ for some minimal unit $\delta > 0$, with no torsion (unique integer representation). That is, every posting is an integer multiple of a fundamental quantum δ, and each value $n\delta$ has a unique representation.

A discreteness assumption for ledger postings can be motivated from several complementary perspectives. Operationally, any finite-precision implementation of the ledger necessarily represents postings by discrete values; requiring exact integer arithmetic with no torsion ensures unambiguous counting and prevents representational ambiguity. Combinatorially, if ledger states are to support unique event-counting—so that configurations can be distinguished by the number of recognition events that have occurred—then the underlying arithmetic must be torsion-free, which forces a structure isomorphic to $\mathbb{Z}$. From the standpoint of minimality, the infinite cyclic group $\mathbb{Z}$ is the simplest discrete choice, as it introduces no finite periodicity or nontrivial quotient structure beyond what is strictly required. Finally, at an analogical level, this quantization mirrors the discrete nature of certain physical quantities, such as electric charge appearing in integer multiples of an elementary unit. We emphasize, however, that this parallel is structural rather than derivational: the quantization of ledger postings is not forced by the cost axioms alone, but is an explicit modeling choice introduced to obtain a genuinely discrete ledger.

Considering the above, the following theorem records the resulting algebraic structure:

Theorem 4 (Theorem T8: Algebraic structure of quantized postings)**.** *Under the quantization assumption (postings in $\delta\mathbb{Z}$ with $\delta > 0$ and no torsion), the set of all ledger increments*

$$\Delta = \{k\delta \mid k \in \mathbb{Z}\}$$

forms an infinite cyclic additive group $(\Delta, +)$ isomorphic to $\mathbb{Z}$ under the mapping $k \mapsto k\delta$.

Consequently, all ledger values have unique representation:

$$x = n\delta, \qquad n \in \mathbb{Z} \text{ (unique)}.$$

Proof. Consider the set $\Delta = \{k\delta \mid k \in \mathbb{Z}\}$ of all possible ledger increments. Under addition, Δ forms an additive group: the sum of any two multiples of δ is again a multiple of δ, the identity element is $0 = 0 \cdot \delta$, and inverses exist (the inverse of $k\delta$ is $(-k)\delta$). Further, the mapping $\varphi : \mathbb{Z} \to \Delta$ defined by $\varphi(k) = k\delta$ is a group homomorphism:

$$\varphi(k_1 + k_2) = (k_1 + k_2)\delta = k_1\delta + k_2\delta = \varphi(k_1) + \varphi(k_2).$$

Now, suppose that $\varphi(k_1) = \varphi(k_2)$, i.e., $k_1\delta = k_2\delta$. Then, $(k_1 - k_2)\delta = 0$. Since $\delta \neq 0$ and the ledger has no torsion (the only element of finite order is zero), we must have $k_1 - k_2 = 0$, hence $k_1 = k_2$. Thus, φ is injective. To prove surjectivity, observe that by construction, every element of Δ is of the form $k\delta$ for some $k \in \mathbb{Z}$, so φ is surjective. Therefore, φ is a group isomorphism: $(\Delta, +) \simeq \mathbb{Z}$, and consequently, each value $x \in \Delta$ has a unique representation $x = n\delta$ for a unique integer n. □

The isomorphism $(\Delta, +) \simeq \mathbb{Z}$ has several immediate consequences for the structure of the ledger. First, recognition events admit no fractional postings: every update posts exactly $\pm\delta$ or, more generally, an integer multiple $\pm k\delta$, with fractional amounts excluded by the algebraic structure. Second, ledger balances support unique counting: each value $n\delta$ corresponds to a unique integer $n \in \mathbb{Z}$, rendering the state space discrete and countable and enabling unambiguous event-counting. Third, ledger arithmetic is purely additive, obeying standard integer operations with no modular identification or quotient structure, so that sums and differences are well defined and free of periodicity constraints. This discrete arithmetic structure is analogous, at a structural level, to quantization in physical systems such as electric charge, although we emphasize that the analogy is not causal: quantization is imposed here as a modeling assumption rather than derived from more fundamental principles within the present framework.

https://doi.org/10.3390/foundations6020017

3.5. Cycle Flux Conservation

By Section 3.3, the double-entry structure guarantees that each recognition event preserves total balance, so that the net sum over all nodes is invariant at each tick. Quantization (Theorem T8) further ensures that all postings occur in discrete integer multiples of δ, while these conditions tightly constrain local updates, they do not alone rule out nontrivial circulation around cycles in the recognition graph. Concretely: a single atomic event at tick t posts $+\delta$ on one edge $e = (u \to v)$ and $-\delta$ on its reverse $e^{-1} = (v \to u)$. On graphs with cycles, such sparse single-edge postings generically create transient net flux around closed loops when measured at individual ticks. To recover a scalar potential representation—enabling path-independent summation of cumulative flows—we must impose an additional constraint that eliminates net circulation. We do this via a *time-aggregated* cycle-closure assumption.

Assumption 1. *(Time-aggregated cycle closure). Fix a clearing horizon $W \geq 1$. Then, for each clearing window $[t_0, t_0 + W)$ and every directed cycle γ in the recognition graph, the* cumulative *cycle flux is zero:*

$$\overline{\Phi}(\gamma; t_0, W) = 0.$$

The above assumption can be interpreted as follows: At the scale of individual ticks, a recognition event acts as an *impulse*: it posts on a single edge (and its reverse) and therefore represents a strictly local update. On graphs containing cycles, such impulses generically produce transient circulation when flux is measured around a loop at a single tick. This behavior is not pathological but a natural consequence of combining atomic, single-edge updates with cyclic graph topology. The assumption therefore does not require zero circulation at every microscopic step. Instead, it asserts that the operationally meaningful constraint is zero *net* circulation after a finite clearing time. Loop imbalances are permitted to arise transiently at the scale of individual ticks, but they must cancel when flows are aggregated over a finite clearing horizon W. In this sense, local conservation is enforced microscopically—through double-entry balance at each tick—while global constraints, such as the absence of arbitrage or path dependence, are imposed only after temporal aggregation.

This separation between microscopic and macroscopic constraints has familiar analogues. In physical systems, transient circulating currents may appear at short times, while dissipation or relaxation mechanisms drive the time-averaged flow toward an irrotational, potential-driven regime. In financial systems, temporary arbitrage opportunities may arise locally, yet market-clearing conditions ensure that such imbalances are eliminated over a clearing period. In both cases, the key feature is that global consistency emerges only after aggregation, mirroring the role of the clearing horizon in the discrete ledger framework.

Remark 9. *The time-aggregated cycle-closure condition is introduced here as a* ***structural assumption*** *of the framework. It is not derived from the cost axioms (Axioms A1–A2) nor from the ledger axioms (Axioms L1–L2, L2b), and should be understood as an additional hypothesis specifying how local recognition events aggregate at larger temporal scales.*

In more elaborate formulations endowed with additional structure, such a cycle closure condition may arise as a consequence rather than an assumption. For instance, in settings governed by global optimization principles, minimizing a total cost functional—such as a sum of comparison costs subject to conservation constraints—can enforce vanishing net circulation as an optimality condition. Similarly, in stochastic systems satisfying detailed balance, equilibrium requirements imply the absence of circulating probability currents. From a graph-theoretic perspective, exactness conditions may also be derived in models with stronger structural restrictions, such as well-foundedness or the absence of certain directed cycle configurations.

In the present work, however, we deliberately refrain from introducing these additional mechanisms. Instead, cycle closure is stated explicitly as a standalone hypothesis, and its consequences are developed directly (Theorems T3–T4). This modular approach preserves logical transparency by making clear which results depend on cycle closure and which follow solely from the cost and ledger axioms.

Before stating Theorem T3, we formalize the notion of (time-aggregated) path-independence. For a fixed clearing window $[t_0, t_0 + W)$, let $P_{x \to y}$ denote a directed path from node x to node y in the recognition graph. Define the cumulative path sum

$$\overline{\Phi}(P_{x \to y}; t_0, W) := \sum_{e \in P_{x \to y}} \overline{\Delta}_{t_0,W}(e),$$

where the sum is taken over edges in the order they appear along the path.

Definition 7 (Path-independence)**.** *The cumulative edge flow $\overline{\Delta}_{t_0,W}(\cdot)$ is* path-independent *if for any two nodes x, y in the same connected component and any two directed paths $P_{x \to y}$ and $P'_{x \to y}$ from x to y, we have $\overline{\Phi}(P_{x \to y}; t_0, W) = \overline{\Phi}(P'_{x \to y}; t_0, W)$.*

Theorem 5 (Theorem T3: Equivalence of time-aggregated cycle closure and path-independence)**.** *Fix a clearing window $[t_0, t_0 + W)$ and assume the recognition structure is connected. The following are equivalent:*

1. ***(Time-aggregated) cycle closure:** For every directed cycle γ, $\overline{\Phi}(\gamma; t_0, W) = 0$.*
2. ***(Time-aggregated) path-independence:** For any nodes x, y and any two directed paths $P_{x \to y}$, $P'_{x \to y}$ from x to y, we have $\overline{\Phi}(P_{x \to y}; t_0, W) = \overline{\Phi}(P'_{x \to y}; t_0, W)$.*

Proof. **(1) ⇒ (2):** Suppose time-aggregated cycle closure holds for $\overline{\Delta}_{t_0,W}$. Fix nodes x, y and two directed paths $P_{x \to y}$ and $P'_{x \to y}$ from x to y. Consider the closed walk formed by following $P_{x \to y}$ from x to y, then following $P'_{x \to y}$ in reverse from y back to x.

Let $\overline{P'_{x \to y}}$ denote the reverse of $P'_{x \to y}$ (same edges, reversed order and orientation). Since each per-tick increment $\delta\Delta(\cdot, t)$ is antisymmetric by construction, the cumulative flow $\overline{\Delta}_{t_0,W}$ is also antisymmetric: $\overline{\Delta}_{t_0,W}(v \to u) = -\overline{\Delta}_{t_0,W}(u \to v)$. Therefore,

$$\overline{\Phi}(\overline{P'_{x \to y}}; t_0, W) = -\overline{\Phi}(P'_{x \to y}; t_0, W).$$

The concatenation $\gamma = P_{x \to y} \circ \overline{P'_{x \to y}}$ is a closed walk and decomposes into directed cycles. By time-aggregated cycle closure, each directed cycle has zero cumulative flux, hence

$$0 = \overline{\Phi}(\gamma; t_0, W) = \overline{\Phi}(P_{x \to y}; t_0, W) - \overline{\Phi}(P'_{x \to y}; t_0, W),$$

which implies path-independence.

(2) ⇒ (1): Suppose time-aggregated path-independence holds. Let $\gamma = (v_0 \to v_1 \to \cdots \to v_n = v_0)$ be a directed cycle. Consider the path $P_{v_0 \to v_0}$ that follows γ once and the trivial path $P'_{v_0 \to v_0}$ with no edges. Then $\overline{\Phi}(P'_{v_0 \to v_0}; t_0, W) = 0$. By path-independence,

$$\overline{\Phi}(P_{v_0 \to v_0}; t_0, W) = \overline{\Phi}(P'_{v_0 \to v_0}; t_0, W) = 0,$$

and since $P_{v_0 \to v_0}$ is γ, we conclude $\overline{\Phi}(\gamma; t_0, W) = 0$. □

Theorem T3 establishes that *time-aggregated* cycle closure is equivalent to *time-aggregated* path-independence, i.e., that the cleared cumulative flow is "curl-free". In continuum settings, curl-free vector fields are precisely those that admit scalar potentials, and the discrete version of this connection is made explicit in Theorem T4.

Example 3 (Cycle Flux Conservation (after clearing))**.** *Consider the cycle $\gamma = (a \to b \to c \to d \to a)$ from Example 2. Over a fixed clearing window $[t_0, t_0 + W)$, suppose the cumulative edge flow is:*

$$\begin{aligned} \overline{\Delta}_{t_0,W}(a \to b) &= +2\delta, \\ \overline{\Delta}_{t_0,W}(b \to c) &= +\delta, \\ \overline{\Delta}_{t_0,W}(c \to d) &= -3\delta, \\ \overline{\Delta}_{t_0,W}(d \to a) &= 0. \end{aligned}$$

The cumulative cycle flux is:

$$\overline{\Phi}(\gamma; t_0, W) = (+2\delta) + (+\delta) + (-3\delta) + (0) = 0.$$

By Theorem T3 (time-aggregated cycle closure), this must always be zero. The example illustrates a choice of net *edge flow whose signed sum around the closed loop vanishes, even though individual ticks inside the window may have nonzero transient loop flux.*

3.6. Discrete Potential Representation (Potential Uniqueness)

The per-tick increments $\delta\Delta(\cdot, t)$ form sparse 1-cochains on the recognition structure. Their cumulative sum over a clearing window defines a (generally non-sparse) 1-cochain $\overline{\Delta}_{t_0,W}$. Under the time-aggregated cycle-closure assumption, Theorem T3 guarantees that $\overline{\Delta}_{t_0,W}$ is closed (path-independent): for every cycle γ, the sum $\overline{\Phi}(\gamma; t_0, W) = \sum_{e \in \gamma} \overline{\Delta}_{t_0,W}(e)$ vanishes. As we next see, the discrete Poincaré lemma provides the existence and uniqueness of a potential function that generates this cumulative flow.

Lemma 1 (Antisymmetry of the cumulative flow)**.** *Assume the recognition structure is closed under reversal (if $(x \to y) \in E$ then $(y \to x) \in E$). Then, for every edge $(x \to y) \in E$ and every clearing window $[t_0, t_0 + W)$, we have*

$$\overline{\Delta}_{t_0,W}(y \to x) = -\overline{\Delta}_{t_0,W}(x \to y).$$

Proof. Fix an edge $(x \to y) \in E$. By construction, each per-tick increment satisfies $\delta\Delta(y \to x, t) = -\delta\Delta(x \to y, t)$. Summing from $\tau = t_0$ to $t_0 + W - 1$ yields the stated antisymmetry for $\overline{\Delta}_{t_0,W}$. □

Definition 8. *A potential function for a clearing window $[t_0, t_0 + W)$ on a connected component $\mathcal{C} \subseteq X$ is a map $\overline{p}_{t_0,W} : \mathcal{C} \to \delta\mathbb{Z}$ such that for each edge $e = (x \to y)$ in $\mathcal{C}$, the edge difference reproduces the cumulative flow: $\overline{\Delta}_{t_0,W}(x \to y) = \overline{p}_{t_0,W}(y) - \overline{p}_{t_0,W}(x)$. This is the standard definition of a discrete gradient.*

Lemma 2 (Discrete Poincaré lemma)**.** *Let $G = (X, E)$ be a connected graph and let $\omega : E \to \delta\mathbb{Z}$ be an antisymmetric function: $\omega(y \to x) = -\omega(x \to y)$. If the sum of ω around every cycle is zero, then there exists $p : X \to \delta\mathbb{Z}$ such that $\omega(x \to y) = p(y) - p(x)$. The function p is unique up to an additive constant.*

Proof. We start the proof by showing existence. Fix a spanning tree T of G and a root $v_0 \in X$. For any $v \in X$, there is a unique simple path $P_{v_0 \to v}$ in T from v_0 to v. Define

$$p(v) := \sum_{e \in P_{v_0 \to v}} \omega(e) \in \delta\mathbb{Z}, \quad p(v_0) := 0.$$

This is well-defined because T is a spanning tree, so the path $P_{v_0 \to v}$ is unique.

We now split the proof into two cases, according to whether $e \in T$ or $e \notin T$:

(i) For an edge $e = (x \to y)$ in T, the unique paths $P_{v_0 \to x}$ and $P_{v_0 \to y}$ differ by exactly the edge e. More precisely, $P_{v_0 \to y} = P_{v_0 \to x} \circ (x \to y)$ (concatenation). Therefore,

$$p(y) = \sum_{f \in P_{v_0 \to y}} \omega(f) = \sum_{f \in P_{v_0 \to x}} \omega(f) + \omega(e) = p(x) + \omega(e),$$

so $p(y) - p(x) = \omega(e)$ for all tree edges.

(ii) If $e = (x \to y) \notin T$, then adding e to T creates a unique fundamental cycle C (since T is a spanning tree, there is exactly one cycle containing e). This cycle consists of e plus the unique path in T from y to x, call it $P^T_{y \to x}$. By hypothesis, the sum of ω around this cycle is zero:

$$0 = \sum_{f \in C} \omega(f) = \omega(e) + \sum_{f \in P^T_{y \to x}} \omega(f).$$

Since $P^T_{y \to x}$ is a path in T from y to x, and by antisymmetry $\omega(y \to x) = -\omega(x \to y)$ for edges in T, we have

$$\sum_{f \in P^T_{y \to x}} \omega(f) = -\sum_{f \in P^T_{x \to y}} \omega(f) = -(p(y) - p(x)),$$

where $P^T_{x \to y}$ is the unique path in T from x to y. Therefore,

$$0 = \omega(e) - (p(y) - p(x)),$$

which implies $\omega(e) = p(y) - p(x)$.

To show uniqueness, let us suppose that $\tilde{p} : X \to \delta\mathbb{Z}$ also satisfies $\omega(x \to y) = \tilde{p}(y) - \tilde{p}(x)$ for all edges $(x \to y) \in E$. Then, for any edge $(x \to y)$,

$$(\tilde{p}(y) - \tilde{p}(x)) - (p(y) - p(x)) = \omega(x \to y) - \omega(x \to y) = 0,$$

so $(\tilde{p} - p)(y) = (\tilde{p} - p)(x)$ for all adjacent vertices. Since G is connected, this implies $\tilde{p} - p$ is constant on X. Setting the constant by choosing $\tilde{p}(v_0) = p(v_0) = 0$ (or any fixed value) determines $\tilde{p}$ uniquely up to this choice. □

In light of the preceding discussion, Theorem T3 together with the discrete Poincaré lemma applied to the cleared cumulative flow $\overline{\Delta}_{t_0,W}$ yields the following result.

Theorem 6 (Theorem T4: Potential Uniqueness). *Fix a clearing window $[t_0, t_0 + W)$ and a connected component $\mathcal{C} \subseteq X$. Under Theorem T3, there exists a potential*

$$\overline{p}_{t_0,W} : \mathcal{C} \longrightarrow \delta\mathbb{Z}$$

such that for each edge $e = (x \to y)$ in $\mathcal{C}$,

$$\overline{\Delta}_{t_0,W}(e) = \overline{p}_{t_0,W}(y) - \overline{p}_{t_0,W}(x).$$

Moreover, $\overline{p}_{t_0,W}$ is unique up to an additive constant on $\mathcal{C}$.

Proof. Theorem T3 asserts time-aggregated cycle closure: for every cycle γ, $\overline{\Phi}(\gamma; t_0, W) = 0$. This means the cumulative flow $\overline{\Delta}_{t_0,W}$ is a closed 1-cochain: the sum around any closed cycle is zero.

The discrete Poincaré lemma (proved above) provides the key tool: if a 1-cochain ω is closed (all cycle sums vanish), then there exists a potential function p such that $\omega = \delta p$, where δp denotes the discrete gradient (edge differences of p).

Applying this to the cumulative flow: since $\overline{\Delta}_{t_0,W}$ is closed by T3, and is antisymmetric by Lemma (antisymmetry of the cumulative flow), the discrete Poincaré lemma guarantees the existence of a potential $\overline{p}_{t_0,W}$ such that $\overline{\Delta}_{t_0,W}(x \to y) = \overline{p}_{t_0,W}(y) - \overline{p}_{t_0,W}(x)$ for all edges.

Uniqueness up to an additive constant follows from the fact that if $\tilde{\overline{p}}_{t_0,W}$ also satisfies $\overline{\Delta}_{t_0,W}(x \to y) = \tilde{\overline{p}}_{t_0,W}(y) - \tilde{\overline{p}}_{t_0,W}(x)$, then $(\tilde{\overline{p}}_{t_0,W} - \overline{p}_{t_0,W})(y) - (\tilde{\overline{p}}_{t_0,W} - \overline{p}_{t_0,W})(x) = 0$ for all edges, implying $\tilde{\overline{p}}_{t_0,W} - \overline{p}_{t_0,W}$ is constant on each connected component.

The proof is constructive: fix a spanning tree, choose a root vertex, and define the potential by summing cumulative postings along tree paths. The cycle condition (T3) ensures this definition is consistent for all edges. □

Theorem T4 establishes that every admissible *cleared* cumulative recognition pattern arises from a scalar potential. This potential is unique up to an additive constant on each connected component, reflecting the gauge freedom familiar in classical physics. In the present framework, the potential representation is a direct consequence of (i) antisymmetry under edge reversal and (ii) time-aggregated cycle closure (T3), which together encode time-aggregated path-independence.

Note that the potential representation is a statement about the *cleared cumulative* flow: if over a clearing window $[t_0, t_0 + W)$, an edge has net flow $\overline{\Delta}_{t_0,W}(x \to y) = k\delta$, then $\overline{p}_{t_0,W}(y) - \overline{p}_{t_0,W}(x) = k\delta$.

Example 4 (Potential Function on a Small Graph (after clearing))**.** *Consider the recognition structure from Example 2 with nodes $\{a, b, c, d\}$ and the cycle $(a \to b \to c \to d \to a)$ plus edge $(a \to c)$. Over a clearing window $[t_0, t_0 + W)$, suppose the cumulative flow is:*

$$\begin{aligned}
\overline{\Delta}_{t_0,W}(a \to b) &= +2\delta, \\
\overline{\Delta}_{t_0,W}(b \to c) &= +\delta, \\
\overline{\Delta}_{t_0,W}(c \to d) &= -3\delta, \\
\overline{\Delta}_{t_0,W}(d \to a) &= 0, \\
\overline{\Delta}_{t_0,W}(a \to c) &= +3\delta.
\end{aligned}$$

Since $\overline{\Phi}(a \to b \to c \to d \to a; t_0, W) = 0$ (as verified in Example 3), Theorem T4 guarantees a potential exists. Following the constructive proof of the discrete Poincaré lemma, choose a as the reference vertex and set $\overline{p}_{t_0,W}(a) = 0$. Then:

$$\begin{aligned}
\overline{p}_{t_0,W}(b) &= \overline{p}_{t_0,W}(a) + \overline{\Delta}_{t_0,W}(a \to b) = 0 + 2\delta = 2\delta, \\
\overline{p}_{t_0,W}(c) &= \overline{p}_{t_0,W}(b) + \overline{\Delta}_{t_0,W}(b \to c) = 2\delta + \delta = 3\delta, \\
\overline{p}_{t_0,W}(d) &= \overline{p}_{t_0,W}(c) + \overline{\Delta}_{t_0,W}(c \to d) = 3\delta + (-3\delta) = 0.
\end{aligned}$$

We verify that $\overline{p}_{t_0,W}(d) - \overline{p}_{t_0,W}(a) = 0 - 0 = 0 = \overline{\Delta}_{t_0,W}(d \to a)$, confirming the cycle closes. For the shortcut edge $(a \to c)$, we check: $\overline{p}_{t_0,W}(c) - \overline{p}_{t_0,W}(a) = 3\delta - 0 = 3\delta = \overline{\Delta}_{t_0,W}(a \to c)$, which is consistent. The potential is unique up to an additive constant: if we had chosen $\overline{p}_{t_0,W}(a) = k$ instead of 0, all values would shift by k, but the edge differences would remain unchanged.

Example 5 (Complete ledger pipeline on a triangle graph)**.** *We trace the full ledger pipeline—atomic events, balance evolution, cycle closure, and potential recovery—on the simplest cyclic recognition structure.*

Setup. Let the recognition structure be the directed triangle $G = (\{P, Q, R\}, E)$*, where* E *contains all six directed edges between the three nodes (each pair connected in both directions). Fix increment* $\delta > 0$ *and set all initial balances to zero:* $b_0(P) = b_0(Q) = b_0(R) = 0$*.*

Events. Over the clearing window $[1, 4]$ *(*$W = 4$ *ticks), four atomic events are recorded—one per tick, as guaranteed by Theorem T2:*

Tick t	*Event* e_t	*Posting (Source)*	*Posting (Target)*	*Total Balance*
1	$P \to Q$	$b(P) \leftarrow b(P) - \delta$	$b(Q) \leftarrow b(Q) + \delta$	0
2	$Q \to R$	$b(Q) \leftarrow b(Q) - \delta$	$b(R) \leftarrow b(R) + \delta$	0
3	$P \to R$	$b(P) \leftarrow b(P) - \delta$	$b(R) \leftarrow b(R) + \delta$	0
4	$Q \to P$	$b(Q) \leftarrow b(Q) - \delta$	$b(P) \leftarrow b(P) + \delta$	0

Balance evolution. Applying each double-entry posting in sequence:

After Tick	$b(P)$	$b(Q)$	$b(R)$	*Total*
0 *(initial)*	0	0	0	0
1	$-\delta$	$+\delta$	0	0
2	$-\delta$	0	$+\delta$	0
3	-2δ	0	$+2\delta$	0
4	$-\delta$	$-\delta$	$+2\delta$	0

Total balance $= 0$ *is maintained at every tick, confirming conservation. After clearing, R holds a net surplus of* $+2\delta$ *funded equally by P and Q.*

Cumulative edge flows. Summing per-tick increments $\delta\Delta(e, t)$ *over* $[1, 4]$ *(with antisymmetry* $\overline{\Delta}(v \to u) = -\overline{\Delta}(u \to v)$*):*

$$\begin{aligned} \overline{\Delta}_{1,4}(P \to Q) &= +\delta \ (\textit{tick } 1) \ + \ (-\delta) \ (\textit{tick } 4, \textit{ via } Q \to P) = 0, \\ \overline{\Delta}_{1,4}(Q \to R) &= +\delta \ (\textit{tick } 2), \\ \overline{\Delta}_{1,4}(P \to R) &= +\delta \ (\textit{tick } 3), \end{aligned}$$

all remaining directed edges having zero cumulative flow.

Cycle closure (Theorem T3). For the directed cycle $\gamma = (P \to Q \to R \to P)$*:*

$$\overline{\Phi}(\gamma; 1, 4) = \overline{\Delta}_{1,4}(P \to Q) + \overline{\Delta}_{1,4}(Q \to R) + \overline{\Delta}_{1,4}(R \to P) = 0 + \delta + (-\delta) = 0.$$

The time-aggregated cycle flux vanishes, confirming cycle closure over the clearing window.

Potential recovery (Theorem T4). By the discrete Poincaré lemma, there exists a unique (up to additive constant) potential $\overline{p}_{1,4} : \{P, Q, R\} \to \delta\mathbb{Z}$*. Setting* $\overline{p}(P) = 0$*:*

$$\begin{aligned} \overline{p}(Q) &= \overline{p}(P) + \overline{\Delta}_{1,4}(P \to Q) = 0 + 0 = 0, \\ \overline{p}(R) &= \overline{p}(Q) + \overline{\Delta}_{1,4}(Q \to R) = 0 + \delta = \delta. \end{aligned}$$

Verification: $\overline{p}(R) - \overline{p}(P) = \delta = \overline{\Delta}_{1,4}(P \to R)$*. ✓The cleared recognition pattern assigns R a potential* δ *above both P and Q, reflecting the net directed flow toward R accumulated over the four-tick window.*

3.7. Minimal Schedule Period 2^d

Having established atomic single-event updates (Theorem T2), quantization (Theorem T8), and the cost function (Theorem T5), we now examine combinatorial constraints that link a discrete carrier (modeled here by Q_d) to discrete time. In particular, we seek lower bounds on the period required to visit all spatial positions under atomic updates. For the purposes of this section, we treat d as an abstract dimension parameter indexing the hypercube family Q_d. The fundamental structure is the d-dimensional hypercube Q_d, which at $d = 3$

(denoted Q_3) provides the minimal cell for ledger-compatible dynamics. The hypercube combinatorics are given in Table 3:

Table 3. Combinatorics of the d-cube at $d = 3$. The Q_3 hypercube has 8 vertices, 12 edges, and 6 faces. In general, the number of k-dimensional faces of the d-cube is $\binom{d}{k}2^{d-k}$ for $0 \le k \le d$. The three rows correspond to the special cases $k = 0$ (vertices), $k = 1$ (edges), and $k = d - 1$ (facets); the formula $2d$ for facets is valid for all d, whereas the formula for 2-dimensional faces ($k = 2$) is $\binom{d}{2}2^{d-2}$, which coincides with $2d$ only at $d = 3$.

Object	**Formula**	$d = 3$
Vertices ($k = 0$)	2^d	8
Edges ($k = 1$)	$d \cdot 2^{d-1}$	12
Facets ($k = d - 1$)	$2d$	6

Now, atomic single-event updates impose strict constraints on how recognition events can be scheduled across the spatial network. To characterize the minimal period required to visit all spatial positions, we introduce the concept of a *ledger-compatible walk*: a temporal sequence of recognition events satisfying atomicity and spatial completeness.

In the scheduler model considered below, the system is represented by a sequence of *active vertices* $(v_t)_{t=0}^{T-1}$ in the hypercube Q_d, where each tick corresponds to the traversal of a single edge. Specifically, the unique recognition event at tick t (guaranteed by Theorem T2) is identified with the directed edge

$$e_t := (v_t \to v_{t+1}),$$

so that time evolution is encoded as a walk on Q_d with exactly one edge-event per tick, and v_t serves as the canonical vertex label associated with time t. This representation aligns naturally with the Gray-code picture of an 8-tick walker in three dimensions.

This scheduler is subject to three structural requirements that formalize the intended dynamics. Atomicity is enforced by construction: since each tick is associated with exactly one edge traversal, no concurrent recognition events are permitted. Spatial completeness requires that, over a full period, every vertex of Q_d appears at least once among the active vertices, ensuring that all spatial positions are visited. Finally, timestamp uniqueness demands that the active vertices v_t are distinct within a period, so that no vertex is revisited at a different time during the same cycle. Together, these conditions ensure that the ledger update is both *atomic*—excluding concurrency—and *complete*—covering the entire spatial domain—while maintaining a strict temporal ordering of events.

As a straightforward consequence, we obtain the following result:

Theorem 7 (Theorem T6: Minimal period 2^d (eight ticks for $d = 3$))**.** *Let C be the vertex set of a d-dimensional hypercube Q_d, with $|C| = 2^d$, and let T be the scheduler period for a ledger-compatible walk.*

1. ***(Sufficiency)** If $T = 2^d$, then there exists a cyclic sequence of active vertices $(v_t)_{t=0}^{T-1}$ that is spatially complete and timestamp-unique (each vertex appears exactly once among the labels v_t), with v_{t+1} adjacent to v_t for each t (one edge traversal per tick). For $d = 3$, the Gray-code Hamiltonian cycle realizes this minimal period:* $000 \to 001 \to 011 \to 010 \to 110 \to 111 \to 101 \to 100 \to 000$.

2. ***(Necessity)*** *If $T < 2^d$, then T ticks are insufficient to assign a distinct active-vertex label to each of the 2^d vertices. By the pigeonhole principle, some vertex label must repeat, so the walk cannot be both spatially complete and timestamp-unique.*

Proof. **Sufficiency:** It is enough to exhibit a Hamiltonian *cycle* on Q_d for each $d \geq 1$ (a cyclic listing of all 2^d binary strings where consecutive strings differ in exactly one bit, including the last-to-first step). A standard construction is the cyclic Gray code.

Define the (reflected) Gray-code map $g\colon \{0, 1, \ldots, 2^d - 1\} \to \{0, 1\}^d$ by

$$g(k) := k \oplus (k \gg 1),$$

where $\oplus$ is bitwise XOR and $(k \gg 1)$ is the right shift. Then, consecutive values differ by one bit:

$$g(k) \text{ and } g(k+1) \text{ differ in exactly one coordinate for } 0 \leq k < 2^d - 1,$$

so $(g(0), g(1), \ldots, g(2^d - 1))$ is a Hamiltonian path on Q_d. Moreover, $g(2^d - 1)$ differs from $g(0)$ in exactly one bit (indeed $g(0) = 0 \cdots 0$ and $g(2^d - 1) = 10 \cdots 0$), so the path closes to a Hamiltonian cycle.

Thus, for any $d \geq 1$, there exists a cyclic sequence of length 2^d that visits each vertex exactly once and advances by one edge per tick. For $d = 3$, this yields the explicit 8-cycle $000 \to 001 \to 011 \to 010 \to 110 \to 111 \to 101 \to 100 \to 000$.

Necessity: By constraint (3) (timestamp uniqueness), each of the 2^d vertices must appear exactly once in the sequence $(v_t)_{t=0}^{T-1}$. Therefore, $T \geq 2^d$. If $T < 2^d$, then by the pigeonhole principle, some vertex label must repeat, violating timestamp uniqueness. Therefore, $T \geq 2^d$ is necessary. □

Therefore, the minimal period compatible with Theorem T2 for a d-dimensional hypercube is exactly

$$T_{\min} = 2^d. \tag{11}$$

In particular, for $d = 3$, this yields the eight-tick period: $T_{\min} = 2^3 = 8$ (within the scheduler model above).

We observe that Theorem T6 establishes the minimal period for a ledger-compatible walk, but it does not address whether this period is sufficient to distinguish all possible patterns. This leads to a complementary result about coverage:

Theorem 8 (Theorem T7: Coverage Lower Bound)**.** *Let Q_d be a d-dimensional hypercube with 2^d vertices, and let T be the period of a ledger-compatible walk in the scheduler model above (one active vertex per tick). If $T < 2^d$, then the walk cannot cover all 2^d vertices within one period without repetition.*

Proof. In the scheduler model above, each tick carries one active-vertex label $v_t \in Q_d$. Therefore a period-T schedule can label at most T distinct vertices. If $T < 2^d$, the schedule cannot cover all 2^d vertices within one period without repetition. □

Together, Theorems T6 and T7 show that $T = 2^d$ is both necessary and sufficient for the scheduler model stated above: $T = 2^d$ is sufficient via a Hamiltonian cycle (Gray code at $d = 3$), while $T < 2^d$ is insufficient for covering all vertices without repetition. For $d = 3$, this yields the eight-tick period $T = 8$ (Figure 2).

https://doi.org/10.3390/foundations6020017

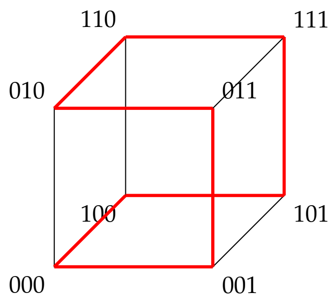

Figure 2. A standard Gray 8-cycle on Q_3: $000 \to 001 \to 011 \to 010 \to 110 \to 111 \to 101 \to 100 \to 000$.

Remark 10 (Computational complexity of the ledger framework)**.** *The minimal-period result $T_{\min} = 2^d$ (Theorem T6) is already a combinatorial complexity statement: ledger-compatible, timestamp-unique coverage on Q_d requires exponentially many ticks in the dimension d. The Gray-code construction shows that this bound is sharp. At the level treated in the present paper, the basic operations are straightforward: evaluating $J(x)$ is $O(1)$, constructing the scalar potential in Theorem T4 can be done in $O(|V| + |E|)$ time once a spanning tree is chosen, and checking time-aggregated cycle closure over a window of width W is $O(W|E|)$. A fuller complexity theory of recognition—including lower bounds under partial observation, decidability questions, or complexity classes for admissible update rules—would require substantial additional development and lies beyond the scope of the present foundational paper.*

The exponential growth 2^d in the number of distinguishable states with dimension d is the combinatorial manifestation of the *curse of dimensionality* familiar from high-dimensional recognition and big-data learning [14]: as the feature space dimension grows, complete coverage of the recognition structure requires resources that scale exponentially. In big-data settings, recognition tasks routinely operate over high-dimensional state spaces—image, text, or sensor data with d in the hundreds or thousands—where the 2^d scaling makes exhaustive ledger coverage infeasible and selective or approximate coverage strategies become necessary. The clearing-horizon parameter W in the time-aggregated cycle-closure assumption also connects naturally to this context: it plays the role of a *batch window* in streaming data processing, where local transient imbalances are permitted within each batch but global consistency (cycle closure) is enforced at the batch boundary—a design pattern standard in large-scale data stream architectures.

4. Discussion

This work develops a cost-first, information-theoretic framework for discrete dynamics grounded in a single primitive: ratio-based comparison. The construction reverses the usual order of assumptions. Rather than starting from an a priori geometry or dynamics and then attaching a cost, we begin from coherence requirements on comparison and ask what discrete update structures are compatible with minimal, unambiguous recording.

Under Axioms A1–A2 (Theorem T5; proof in [9]), coherent composition of ratio comparisons uniquely fixes the reciprocal cost $J(x) = \frac{1}{2}(x + x^{-1}) - 1$. Its reciprocity, unique zero at $x = 1$, and boundary divergence provide a canonical equilibrium point and exclude boundary regimes at infinite cost without additional prohibitions.

Taking J as input, recognition events are modeled as deterministic edge-events on a directed graph recorded by a minimal ledger. Determinism together with minimality (no intra-tick ordering metadata) and non-commutativity forces atomic ticks (Theorem T2). With conservation, pairwise locality, and no external sources or sinks, each atomic event is recorded as a balanced double-entry posting; under the quantization assumption, postings take values in $\delta\mathbb{Z}$ (Theorem T8).

On graphs with cycles, single-edge atomic events generically create transient circulation. Assuming time-aggregated cycle closure over a finite clearing horizon yields a macroscopic consistency condition: cleared cycle closure is equivalent to path-independence (Theorem T3), and the cleared cumulative flow admits a scalar potential on each connected component, unique up to an additive constant (Theorem T4). This provides a discrete potential theory appropriate to impulse-like updates: circulation may occur microscopically but is eliminated in the cleared aggregate.

Finally, when the recognition structure is the hypercube Q_d, atomic single-edge updates imply a 2^d-tick lower bound for timestamp-unique vertex coverage, and cyclic Gray codes achieve this bound (Theorems T6–T7), with an explicit 8-tick cycle at $d = 3$. The golden-ratio fixed point of the simple ratio iteration $x_{n+1} = 1 + 1/x_n$ supplies a convenient log-scale reference for deviations, but carries no necessity without additional dynamical assumptions.

At a broader level, the framework has structural connections to artificial intelligence settings in which recognition quality is assessed through divergence-like penalties. Comparing a predicted output, likelihood, or confidence weight to a reference is naturally ratio-based, and the cost functional $J(x) = \frac{1}{2}(x + x^{-1}) - 1$ belongs to the same formal family of divergence constructions discussed in information geometry [5,7]. Likewise, the deterministic ledger update $S_{t+1} = U(S_t, e_t)$ provides an abstract analogue of sequential state evolution, while the hypercube scheduling results connect to conflict-free traversal and routing problems. These parallels are intentionally structural rather than empirical: the present manuscript does not introduce an AI algorithm or claim performance results. We therefore confine ourselves to indicating the conceptual relevance of the framework and leave algorithmic specialization to future work.

The framework is intentionally modular: results are traced to explicit axioms and structural assumptions, leaving clear targets for extension. A natural question raised by this modularity is the degree of coupling between the cost functional J and the ledger dynamics. The ledger axioms L1–L2b and the structural assumptions (conservation, pairwise locality, quantization, cycle closure) are logically independent of A1–A2: all ledger results would hold for a different cost function. The label *cost-first* therefore refers to chronological and conceptual priority—J is established before the ledger is built, and its properties motivate the framework's design—rather than to full logical entailment. The specific form of J currently connects to the ledger through three conceptual links: reciprocity $J(x) = J(x^{-1})$ is compatible with balanced double-entry postings; the unique zero $J(x) = 0 \Leftrightarrow x = 1$ defines the Bal predicate and the equilibrium reference; and boundary divergence grounds the Exists predicate. Showing that J's specific cosh structure actively selects or constrains ledger features—for example, deriving the quantization unit δ from J, or obtaining cycle closure as a consequence of minimizing total recognition cost over a clearing window—constitutes a concrete open problem and a primary target for future work. The d'Alembert structure of Axiom A1 also connects the present framework to limit theorems in probability theory, where characteristic functional equations of the same type characterize stable and infinitely divisible distributions [12,13]; the Ulam–Hyers stability of these equations under perturbation—whether approximate satisfaction of the coherence requirement forces proximity to J—is a natural companion problem. The complexity-theoretic structure of the framework—including the 2^d scheduling bound, the tractability of the core ledger operations, and the connection between Axiom L2 and Minimum Description Length [15,16]—also suggests links to computational models of recognition that merit systematic treatment in future work. Future work may seek conditions under which clearing/cycle closure can be derived (e.g., from global optimization or stochastic detailed-balance principles), explore admissible update rules on broader graph families, and study controlled continuum limits linking discrete ledgers to effective physical theories.

Author Contributions: Conceptualization, J.W.; methodology, J.W.; software, J.W.; validation, S.P.-G., M.S., A.T. and J.W.; formal analysis, S.P.-G., M.S., A.T. and J.W.; investigation, S.P.-G., M.S. and A.T.; writing—original draft preparation, J.W.; writing—review and editing, S.P.-G.; visualization, S.P.-G.; supervision, J.W.; project administration, S.P.-G.; funding acquisition, J.W. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Data Availability Statement: All data is contained within the article.

Acknowledgments: The authors gratefully acknowledge the anonymous referees for their constructive feedback, which significantly enhanced the quality of the manuscript.

Conflicts of Interest: The authors declare no competing interests.

References

1. Washburn, J.; Zlatanović, M.; Allahyarov, E. Recognition Geometry. *Axioms* **2026**, *15*, 90. [CrossRef]
2. Shannon, C.E. A mathematical theory of communication. *Bell Syst. Tech. J.* **1948**, *27*, 379–423. [CrossRef]
3. Cover, T.M.; Thomas, J.A. *Elements of Information Theory*, 2nd ed.; Wiley: New York, NY, USA, 2006.
4. Kullback, S.; Leibler, R.A. On information and sufficiency. *Ann. Math. Stat.* **1951**, *22*, 79–86. [CrossRef]
5. Csiszár, I. Information-type measures of difference of probability distributions and indirect observations. *Stud. Sci. Math. Hung.* **1967**, *2*, 299–318.
6. Bregman, L.M. The relaxation method of finding the common point of convex sets and its application to the solution of problems in convex programming. *USSR Comput. Math. Math. Phys.* **1967**, *7*, 200–217. [CrossRef]
7. Amari, S.-i.; Nagaoka, H. *Methods of Information Geometry*; American Mathematical Society and Oxford University Press: Rhode Island, RI, USA, 2000.
8. Jaynes, E.T. *Probability Theory: The Logic of Science*; Cambridge University Press: Cambridge, UK, 2003.
9. Washburn, J.; Zlatanović, M. Uniqueness of the Canonical Reciprocal Cost. *Mathematics* **2026**, *14*, 935. [CrossRef]
10. Aczél, J. Lectures on Functional Equations and Their Applications. In *Mathematics in Science and Engineering*; Academic Press: New York, NY, USA, 1966.
11. Kuczma, M. *An Introduction to the Theory of Functional Equations and Inequalities*, 2nd ed.; Birkhäuser: Berlin, Germany, 2009.
12. Feller, W. *An Introduction to Probability Theory and Its Applications*, 2nd ed.; Wiley: New York, NY, USA, 1971; Volume 2.
13. Lukacs, E. *Characteristic Functions*, 2nd ed.; Griffin: London, UK, 1970.
14. Bellman, R. *Dynamic Programming*; Princeton University Press: Princeton, NJ, USA, 1957.
15. Rissanen, J. Modeling by shortest data description. *Automatica* **1978**, *14*, 465–471. [CrossRef]
16. Grünwald, P.D. *The Minimum Description Length Principle*; MIT Press: Cambridge, MA, USA, 2007.