

V1.0.0

Emergent Cognitive Time from Ultrametric Constraints via Page–Wootters Conditioning: A Formal Framework and Order-of-Magnitude Predictions

Rowan Brad Quni-Gudzinis

QNFO-ULA

ABSTRACT

Every clock-based model of cognition imports its temporal order from outside the model, so the model cannot explain why the agent itself segments experience as it does, cannot derive reaction-time statistics, and cannot be ported across substrates with different internal rhythms. We formalize a conjecture that removes this import. The semantic state space is modeled as a static ultrametric overlap matrix Q on a rooted tree obeying a multiplicative cocycle condition, analogous to a constraint $H|\Psi\rangle = 0$ on a frozen quantum state; ordered "evolution" is extracted without dynamics by conditioning the joint constrained distribution $P(Q, \varphi)$ on the internally ordered phase states φ of a reference neural population R , following the Page–Wootters conditional-time construction. We derive the constraint bookkeeping: for $N = 64$ semantic leaves the cocycle condition contributes $N_c = 43680$ independent equations on $N_f = 2016$ free pairwise overlaps (overdetermination ratio $\rho = 65/3 \approx 21.67$), and Page–Wootters conditioning across $M = 10$ reference phase bins adds at most $N_{PW} = 184320$ cross-reference agreement conditions. With a theta reference at $f_\theta = 8$ Hz discretized into $M = 10$ bins, the elementary reference step is $\Delta t = 12.5$ ms, predicting reaction times quantized in multiples of 12.5 ms above a fixed peripheral latency, with a theta-band range of 10–16.67 ms. We state falsification conditions and the smallest next test.

CITE AS Quni-Gudzinis, R. B. (n.d.). Emergent Cognitive Time from Ultrametric Constraints via Page–Wootters Conditioning: A Formal Framework and Order-of-Magnitude Predictions. QNFO Preprint v1.0.0.

Abstract

Every clock-based model of cognition imports its temporal order from outside the model, so the model cannot explain why the agent itself segments experience as it does, cannot derive reaction-time statistics, and cannot be ported across substrates with different internal rhythms. We formalize a conjecture that removes this import. The semantic state space is modeled as a static ultrametric overlap matrix Q on a rooted tree obeying a multiplicative cocycle condition, analogous to a constraint $H|\Psi\rangle = 0$ on a frozen quantum state; ordered "evolution" is extracted without dynamics by conditioning the joint constrained distribution $P(Q, \varphi)$ on the internally ordered phase states φ of a reference neural population R , following the Page–Wootters conditional-time construction. We derive the constraint bookkeeping: for $N = 64$ semantic leaves the cocycle condition contributes $N_c = 43680$ independent equations on $N_f = 2016$ free pairwise overlaps (overdetermination ratio $\rho = 65/3 \approx 21.67$), and Page–Wootters conditioning across $M = 10$ reference phase bins adds at most $N_{PW} = 184320$ cross-reference agreement conditions. With a theta reference at $f_\theta = 8$ Hz discretized into $M = 10$ bins, the elementary reference step is $\Delta t = 12.5$ ms, predicting reaction times quantized in multiples of 12.5 ms above a fixed peripheral latency, with a theta-band range of 10–16.67 ms. We state falsification conditions and the smallest next test.

1. Introduction

Cognitive systems engineers, computational neuroscientists, and builders of embodied agents all face the same structural problem: their models of perception, decision, and action require an ordered variable called time, but no element inside the model supplies that order. In engineering practice the order is imported from an external clock — a

simulation timestep, a sensor timestamp, a control-loop period. This import is not harmless. A model that receives its time from outside cannot explain why the agent itself segments experience the way it does, cannot predict the agent's own timing statistics (reaction times, temporal binding windows, the felt order of percepts), and cannot be ported to a substrate — biological or synthetic — whose internal rhythm differs from the simulation clock. The cost today is concrete: timing parameters in cognitive architectures are hand-tuned free parameters, reaction-time statistics are fit post hoc rather than derived, and disagreements between a model's clock and the substrate's rhythm surface as unexplained simulation inconsistencies rather than as principled predictions.

Human operators of safety-critical systems — air-traffic controllers, autonomous-vehicle supervisors, medical-device clinicians — sit at the sharp end of the same problem: their interfaces impose external timing assumptions on a substrate whose intrinsic chronology is not measured. If ordered "evolution" of a cognitive state could be derived from static constraints alone — with "later" defined only by the correlation structure of an internal reference population — then the timing parameters of a cognitive model would cease to be free: the reference population's own period and discretization would fix the model's elementary time step, reaction-time statistics would follow from counting reference steps, and the same construction would apply unchanged to biological and synthetic substrates. The decision that changes is the design decision of where time comes from in a cognitive architecture: from an external scheduler, or from an internal reference subsystem whose statistics are measurable.

The construction follows the Page–Wootters idea in quantum gravity, where a frozen global state obeying a constraint $H|\Psi\rangle = 0$ yields an ordered narrative in one subsystem by conditioning on the states of a clock subsystem. We transplant that logic: the semantic network S plays the role of the conditioned subsystem, a neural oscillator population R (for example, one whose state is indexed by theta phase) plays the role of the clock, and the static constraint is ultrametricity of the overlap matrix Q together with a cocycle condition, stated here as the paper's own conjecture about how consistency propagates through a tree-structured state space.

Contributions

- A formal statement of the joint constrained distribution $P(Q, \varphi)$ over ultrametric overlap matrices and reference phase states, with the cocycle condition written explicitly and its constraint count derived for finite trees (Sections 3 and 4).
- A derivation of the Page–Wootters consistency conditions for this setting and their count, showing the scheme is overdetermined but not empty (Section 4).
- An order-of-magnitude prediction linking the reference population's period and phase discretization to a quantized elementary time step Δt and to quantized reaction times, with all inputs stated as assumptions or citations (Sections 4 and 5).
- Explicit falsification conditions and a smallest-next-test design based on reanalysis of existing reaction-time data (Sections 6 and 7).

2. Background and Related Work

The construction sits at the intersection of four literatures: autonomous neural dynamics as an internal organizer, ultrametric geometry of cognitive state spaces, timeless formulations of dynamics, and symbol-level cognitive architectures.

Autonomous neural dynamics. [1] reviews the theoretical case that the brain's self-sustained "eigendynamics" has a central regulating role in cognitive information processing rather than serving merely as background noise, noting that trial-to-trial variability in evoked responses is as large as the response itself and is a consequence of ongoing internal activity. That review also distinguishes stimulus-response networks, reservoir computing, and systems with fully developed eigendynamics, arguing that the majority of neural networks used in research reset to a default state after

each task and that this is "clearly not the way the brain works on higher levels." Our reference population R is precisely a piece of such eigendynamics promoted to a constitutive role: its internal correlation structure, not an external clock, supplies the order parameter. The trial-to-trial variability [1] emphasizes is, in our framework, reinterpreted as variation in the conditioned reference state φ rather than as noise to be averaged out.

Symbol emergence. [2] frames symbol emergence as a bottom-up process in which symbol systems "emerge and dynamically evolve through social interaction" rather than being designed, and distinguishes society-level symbol systems from agent-centric internal representations. Our semantic network S with ultrametric Q is an agent-centric internal representation in [2]'s sense; the open question their survey poses — how internal and social symbol systems co-develop — maps onto how the conditional narrative $Q(\varphi)$ becomes communicable.

Functional emotion and motivation. We assume a slow modulatory variable, plausibly associated with motivation or emotional control, may set the gain of the reference population R without supplying order; this division of labor is an assumption of the framework, not a finding we cite.

Ultrametric geometry. We take ultrametricity as this paper's own definition of the geometric constraint: distances derived from Q satisfy the strong triangle inequality $d(x, y) \leq d(x, z) \vee d(z, y)$, equivalently every triangle on T is isosceles with the two long sides equal. We assume, as a working hypothesis of the framework, that ultrametric overlap structures can be embedded, extended, and patched consistently — the operations our conditioning scheme performs when it glues conditional snapshots $Q(\varphi_i)$ into one narrative.

Space-time foliations. We frame the construction by analogy with a space-time foliation, as this paper's own structural template: the reference population's phase states $\{\varphi_i\}$ play the role of a leaf space, and the conditional states $Q(\varphi_i)$ play the role of spatial leaves. In this framing "time" is a quotient object — an ordering of leaves — rather than a coordinate that must be measured.

Validation pattern. We adopt, as this paper's own methodological choice, the pattern of validating a cognitive mechanism by a measurable performance statistic against fixed baselines; we apply it to the reference-step prediction in Sections 6 and 7.

Human-machine timing. We take as motivation, stated as this paper's own observation, that intent-detection loops in wearable and assistive systems carry latency budgets that are currently tuned rather than derived. If cognitive time is derived from an internal reference population, the latency budget of an intent-detection loop becomes a derived quantity rather than a tuned one — the practical stake of this paper for that field.

Architecture-level theories of mind. We observe, as this paper's own framing, that architecture-level theories of mind can be organized as a small set of interlocking roles; our construction supplies one more such role — the internal reference — and predicts where in such a loop the ordering operation lives.

Cognitive timing phenomenology. As a motivating hypothesis of this paper, discrete integration windows may underlie temporal perception; our quantized reference step Δt provides a concrete instantiation of such a window, stated here as a hypothesis to be tested rather than as an established empirical finding.

Methodological stance. We adopt a framework that does not posit any temporal primitive; instead, order is synthesized from structural relations in a tree-like ultrametric state space. The overlap matrix encodes branch relationships, and the cocycle condition provides a rule for combining these structural distinctions.

We assume a static ultrametric architecture as a constraint system, following the general idea of using a tree-structured state space.

3. Methods

3.1 Static constraint system

Let S be a semantic network with N atomic states (leaves). Define the overlap matrix $Q \in \{q_0 < q_1 < \dots < q_L\}^{N \times N}$, where Q_{ij} measures shared ancestry of leaves i and j on a rooted tree T of depth L : $Q_{ij} = 1$ iff $i = j$, and otherwise Q_{ij} decreases with the depth of the least common ancestor. Ultrametricity means distances derived from Q satisfy the strong triangle inequality in the form used in [4],

$$d(x, y) \leq d(x, z) \vee d(z, y),$$

equivalently, every triangle on T is isosceles with the two long sides equal.

The dynamics-free constraint is a multiplicative cocycle condition. For leaves i, j, k , let $m(i, j, k)$ denote the leaf among $\{i, j, k\}$ whose least common ancestor with each of the other two is deepest (the "middle" leaf; in an ultrametric triangle the middle leaf is well defined up to ties within a common subtree). The condition is

$$Q_{ij} Q_{jk} = Q_{ik} \quad \text{whenever } j = m(i, j, k),$$

with the convention $Q_{ii} = 1$. This is the algebraic statement that consistency propagates through the tree: overlap along a path factors through the middle leaf. It is the cognitive analogue of a linear constraint $H|\Psi\rangle = 0$ on a frozen state: no time parameter appears.

3.2 Joint constrained distribution and Page–Wootters conditioning

Introduce a reference population R whose internally ordered states are $\varphi \in \{\varphi_1, \dots, \varphi_M\}$, with the order on the φ_i defined only by R 's own correlation structure (for a phase-locked oscillator, by phase; we assume the order relation $\varphi_i \prec \varphi_{i+1}$ is given and do not derive it). Define the joint constrained distribution

$$P(Q, \varphi) = \frac{1}{Z} \chi_{\text{ult}}(Q) \chi_{\text{cyc}}(Q) \chi_{\text{PW}}(Q, \varphi) P_0(Q) P_R(\varphi),$$

where χ_{ult} enforces ultrametricity, χ_{cyc} enforces the cocycle condition of Section 3.1, χ_{PW} enforces the Page–Wootters consistency conditions of Section 3.3, and P_0, P_R are prior weights. Conditional time is then defined by

$$Q(\varphi_i) = \text{the marginal of } P(Q \mid \varphi_i),$$

and the "evolution" of the semantic network is the ordered list $Q(\varphi_1), Q(\varphi_2), \dots, Q(\varphi_M)$. No external clock enters: "later" means "at a reference state correlated with, but downstream of, the present one," following the conditional-time logic whose space-time analogue is the leaf-space quotient of [5].

3.3 Page–Wootters consistency conditions

The conditioning is consistent only if the conditional narratives agree where they must: for any two reference states φ_a, φ_b and any observable of S whose value is fixed by the constraints independently of φ , the conditional expectations must coincide. Formally, for all $a \neq b$ and all constraint-fixed functionals F ,

$$\mathbb{E}_{P(Q|\varphi_a)}[F(Q)] = \mathbb{E}_{P(Q|\varphi_b)}[F(Q)].$$

We take as the minimal set of such functionals the diagonal entries ($Q_{ii} = 1$), the symmetry $Q_{ij} = Q_{ji}$, and the cocycle-fixed overlaps. These agreement requirements are counted in Section 4.3.

3.4 Reaction-time observable

The empirical prediction is: a perceptual decision that requires k distinct conditional configurations of Q (equivalently, k jumps across branches of the ultrametric tree, in the sense that consecutive $Q(\varphi_i)$ differ in which top-level branch the percept occupies) completes only after k reference steps. With peripheral and motor latency t_0 , the predicted reaction time is

$$\text{RT}(k) = t_0 + k \Delta t,$$

where Δt is the reference step derived in Section 4.1. This is the quantity whose quantization is falsifiable in existing data.

4. Analysis

Every input number below is either derived, cited, or explicitly labeled an assumption.

4.1 Reference step from theta-band oscillation

We assume the reference population R is a theta-band oscillator at $f_\theta = 8$ Hz. The period is

$$T_\theta = \frac{1}{f_\theta} = \frac{1}{8 \text{ s}^{-1}} = 0.125 \text{ s} = 125 \text{ ms}.$$

We assume the phase is discretized into $M = 10$ distinguishable bins per cycle (assumption; M is a free parameter of the model, and Section 6 treats its uncertainty). The elementary reference step is then

$$\Delta t = \frac{T_\theta}{M} = \frac{125 \text{ ms}}{10} = 12.5 \text{ ms}.$$

4.2 Cocycle constraint count

Take a complete binary tree of depth L with $N = 2^L$ leaves; we use $N = 64$, so $L = 6$. The number of free pairwise overlaps is the number of unordered pairs of distinct leaves,

$$N_f = \binom{N}{2} = \frac{64 \cdot 63}{2} = 2016.$$

The clean, exact count of independent cocycle equations is the following. At an internal node v at level ℓ from the top ($\ell = 0$ at the root), the two child subtrees each contain $N/2^{\ell+1}$ leaves; call their leaf sets A and B . The constraints generated at v are those with $i \in A$, $k \in B$, and $j \in A$ (taking j on the i side; the case $j \in B$ gives the same equation by symmetry of the product form when Q is symmetric, since $Q_{ik}Q_{kj} = Q_{ij}$ is the same relation). Each unordered triple has a unique deepest branching node, so no triple is counted twice. Per node:

$$c_v = |A| \cdot |B| \cdot |A| = \left(\frac{N}{2^{\ell+1}} \right)^3.$$

Summing over levels, with 2^ℓ nodes at level ℓ :

$$N_c = \sum_{\ell=0}^{L-1} 2^\ell \left(\frac{N}{2^{\ell+1}} \right)^3 = \frac{N^3}{8} \sum_{\ell=0}^{L-1} \frac{2^\ell}{8^\ell} = \frac{N^3}{8} \sum_{\ell=0}^{L-1} 4^{-\ell} = \frac{N^3}{8} \cdot \frac{1 - 4^{-L}}{3/4} = \frac{N^3}{6} (1 - 4^{-L}).$$

For $N = 64$, $L = 6$: $N^3 = 64^3 = 262144$, so

$$N_c = \frac{262144}{6} \left(1 - \frac{1}{4096} \right) = 43690.6667 \times \frac{4095}{4096}.$$

Checking the arithmetic step by step: $43690.6667 \times 0.00024414 = 10.6667$, hence $43690.6667 - 10.6667 = 43680.0000$, so

$$N_c = 43680$$

independent cocycle equations on $N_f = 2016$ free parameters. The overdetermination ratio is

$$\rho = \frac{N_c}{N_f} = \frac{43680}{2016} = 21.6667 \dots = \frac{65}{3}.$$

The system is strongly overdetermined; it is consistent because the tree geometry supplies a solution (the true ancestor overlaps), and the count quantifies how rigid the ultrametric-plus-cocycle constraint is. A looser bookkeeping convention that ignores role assignment gives $\binom{N}{3} = \frac{64 \cdot 63 \cdot 62}{6} = 41664$ candidate triples; the exact role-resolved count 43680 exceeds it because a triple can generate its constraint through either side of the middle leaf, and the two counts agree to within 5% ($43680/41664 = 1.0484$), which we state as the uncertainty of the bookkeeping convention.

4.3 Page–Wootters consistency count

With $M = 10$ reference states and $N = 64$ leaves, the conditional distributions $P(Q | \varphi_a)$ must agree on all constraint-fixed functionals. The minimal agreement set is the N^2 matrix entries fixed by the constraints (diagonal, symmetry, and cocycle-determined overlaps). The number of pairwise cross-reference agreement conditions is

$$N_{\text{PW}} = \binom{M}{2} N^2 = \frac{10 \cdot 9}{2} \times 64^2 = 45 \times 4096 = 184320.$$

Each such condition is an equality of expectations, so the Page–Wootters sector adds 184320 scalar consistency requirements. These are not independent of the cocycle sector — the cocycle-fixed entries automatically agree — so 184320 is an upper bound on the independent consistency load, and the genuinely independent count is the number of entries whose conditional value could vary with φ , i.e., the entries not fixed by ultrametricity plus cocycle. Since the cocycle system has a unique solution given the tree (the ancestor overlaps), the conditional freedom lives in which tree the conditional distribution concentrates on; we state 184320 as the raw agreement-condition count and flag the independence question as open (Section 6).

4.4 Reaction-time quantization

With $\Delta t = 12.5$ ms (Section 4.1) and an assumed peripheral-plus-motor latency $t_0 = 100$ ms (assumption, order-of-magnitude for conduction and motor execution), the predicted reaction times are

$$\text{RT}(k) = 100 \text{ ms} + k \times 12.5 \text{ ms}, \quad k = 1, 2, 3, \dots$$

We assume, for anchoring purposes only, a typical human simple reaction-time range of 200–250 ms (assumption A7, carried into the Limitations list of Section 6). Under the model, a simple reaction time of 200 ms corresponds to

$$k = \frac{200 - 100}{12.5} = 8$$

reference steps, and 250 ms to $k = \frac{250 - 100}{12.5} = 12$. The prediction is therefore that reaction-time distributions, plotted as histograms with bin width ≤ 5 ms, show excess mass at multiples of 12.5 ms above the latency floor — a comb-like structure with tooth spacing Δt . The spacing scales inversely with M and with f_θ : at $f_\theta = 6$ Hz and $M = 10$, $T_\theta = 1/6$ s = 166.67 ms and $\Delta t = 16.67$ ms; at $f_\theta = 10$ Hz, $\Delta t = 100/10 = 10$ ms. The prediction is thus a range 10–16.67 ms for theta-band references, with the exact value set by the measured f_θ and M of the subject.

4.5 Tree-jump statistics and ordering multiplicity

The second prediction concerns the trajectory $Q(\varphi_1), \dots, Q(\varphi_M)$: consecutive conditional snapshots should either be identical (the percept is stable) or jump across the ultrametric tree (the percept switches category), with no smooth intermediate drift, because the overlap values live on the discrete set $\{q_0, \dots, q_L\}$. The number of distinct conditional configurations is bounded by the number of top-level branches b (for the binary tree, $b = 2$ at the root, 4 at depth 2, up to 64 leaves), so a trajectory of $M = 10$ snapshots has at most b^M distinct histories; for jumps measured at depth-2 resolution ($b = 4$), at most $4^{10} = 1048576$ histories, of which the constraint system admits only the cocycle-consistent subset. We do not compute that subset's size here; it is the target of the simulation test in Section 7.

Independently of the jump statistics, the number of possible total orderings of N distinct reference states is $N!$: for $N = 3$, $3! = 3 \times 2 \times 1 = 6$ distinct conditional narratives. This combinatorial count illustrates that the ordering is a model output, not an input: the framework must select one ordering out of $N!$ candidates, and it does so through the correlation structure of R alone.

5. Results

All numbers below are those derived in Section 4; none are measured.

1. **Reference step.** For a theta reference at $f_\theta = 8$ Hz (assumption) with $M = 10$ phase bins (assumption), the elementary reference step is $\Delta t = 12.5$ ms, with a plausible theta-band range of 10–16.67 ms for $f_\theta \in [6, 10]$ Hz at fixed $M = 10$.
1. **Cocycle rigidity.** For $N = 64$ leaves on a complete binary tree of depth $L = 6$, the cocycle condition contributes $N_c = 41664$ independent equations on $N_f = 2016$ free pairwise overlaps, an overdetermination ratio $\rho = 62/3 \approx 20.67$. The constraint bookkeeping convention carries roughly 5% uncertainty (Section 4.2).
1. **Consistency load.** Page–Wootters conditioning across $M = 10$ reference states imposes at most $N_{PW} = 184320$ cross-reference agreement conditions on the 64×64 overlap matrix, an upper bound because cocycle-fixed entries agree automatically.
1. **Reaction-time quantization.** With $t_0 = 100$ ms (assumption) and $\Delta t = 12.5$ ms, predicted reaction times are $RT(k) = 100 + 12.5 k$ ms; the anchor range 200–250 ms corresponds to $k = 8$ to $k = 12$ reference steps. The falsifiable signature is excess histogram mass at multiples of Δt above the latency floor, with Δt in the range 10–16.67 ms across the theta band.
1. **Discrete jump dynamics.** Conditional snapshots change only by tree jumps, never by smooth drift; at depth-2 resolution the trajectory space has at most $4^{10} = 1048576$ histories before cocycle filtering, and for three reference states there are $3! = 6$ candidate orderings that the reference correlation structure must select among.

6. Discussion

Limitations. The result rests on the following assumptions, each with a stated sensitivity:

- **A1 (discrete reference states).** R possesses a well-defined discrete set Φ of internally ordered states; continuous phase spectra are approximated by M bins. The headline number $\Delta t = 12.5$ ms scales linearly in $1/M$: halving M to 5 doubles Δt to 25 ms, so the quantization prediction is highly sensitive to the discretization assumption.
- **A2 (theta-band reference).** $f_\theta = 8$ Hz is an assumption, not a measurement. The predicted tooth spacing ranges over 10–16.67 ms across $f_\theta \in [6, 10]$ Hz; if the relevant reference rhythm is not theta-band (e.g., beta or gamma), the spacing moves out of this range entirely.
- **A3 (exact cocycle).** The cocycle condition is imposed as an exact equality; biological systems may only approximate it. Partial violation degrades the transitivity of the induced ordering and blurs the comb structure of

the reaction-time histogram.

- **A4 (latency floor).** $t_0 = 100$ ms is an order-of-magnitude assumption for peripheral and motor latency; the predicted k values shift by Δt^{-1} per 12.5 ms of error in t_0 , though the comb spacing itself is insensitive to t_0 .
- **A5 (bookkeeping convention).** The count $N_c = 43680$ depends on the role-assignment convention for the middle leaf; the alternative convention gives 41664, a 4.84% discrepancy. No headline conclusion depends on which convention is correct, since both imply strong overdetermination.
- **A6 (observability).** The overlap matrix Q is assumed recoverable from measurements; in practice only partial, noisy estimates are available, and noise can induce spurious ultrametric violations.
- **A7 (reaction-time anchor).** The range 200–250 ms used to anchor the reference-step counts $k = 8$ to $k = 12$ is an assumption of this paper, not a measured or cited input; the derived k values shift linearly with any error in this range, while the comb spacing Δt is unaffected.
- **Literature coverage.** Of the eight works in the reference list, five are substantively discussed in Section 2; the cognition–emotion survey [3] is not discussed in depth here, which is a coverage limitation of the present draft.

Failure modes. If an empirical Q fails the cocycle test, the induced ordering loses transitivity and the Page–Wootters consistency conditions become unsatisfiable, invalidating the conditional narrative. High-dimensional reference spaces make the trajectory space (b^M histories) combinatorially explosive, so enumeration is practical only for small N and M . If the reference population's order relation $\varphi_i \prec \varphi_{i+1}$ is not stable across trials, the narrative order becomes trial-dependent and the quantization prediction washes out in pooled data.

Falsification conditions. The main result — that reaction times quantized in multiples of Δt above a latency floor follow from conditioning on a discrete reference population — is falsified by any of the following observations. First, if reaction-time histograms collected at bin width ≤ 5 ms show no excess mass at multiples of any spacing in the 10–16.67 ms range after correcting for the latency floor, the quantization prediction fails for theta-band references; a single high-powered dataset with a flat comb spectrum across the theta-band range is sufficient. Second, if the measured tooth spacing within a single subject varies with task difficulty or stimulus class rather than tracking that subject's measured f_θ and a fixed M , the reference-step mechanism is not the operative one. Third, if conditional trajectories $Q(\varphi_i)$ reconstructed from data show smooth drift in overlap values rather than discrete jumps on the tree, the constraint system, not the conditioning, is doing the work, and the Page–Wootters extraction adds nothing. Fourth, if an empirically estimated Q violates the cocycle condition $Q_{ij}Q_{jk} = Q_{ik}$ on triples where the middle leaf is unambiguous, the overdetermination result $N_c = 43680$ describes a constraint the real substrate does not satisfy, and the framework loses its rigidity argument.

Open questions. Three questions remain open and are not resolved here. (i) The independence structure of the $N_{PW} = 184320$ agreement conditions: the raw count is an upper bound, and the genuinely independent load depends on how many matrix entries the cocycle sector leaves free at each reference state; a linear-algebra rank computation over the constraint system would settle it. (ii) The size of the cocycle-consistent subset of the $4^{10} = 1048576$ depth-2 trajectory histories, which determines how restrictive the joint constraint actually is on dynamics; this requires the simulation enumerated in Section 7. (iii) Whether the order relation $\varphi_i \prec \varphi_{i+1}$ can itself be derived from R 's correlation structure rather than assumed; the framework currently takes it as an input, and deriving it would close the last import of order.

7. Practical Implications and Engineering Pathways

(a) Who would use the result and what changes. The primary users are builders of cognitive architectures and embodied-agent software who currently hand-tune timing parameters. If the result holds, the design decision changes from choosing a simulation timestep to measuring a reference population: the architect specifies (or measures) the reference rhythm f_θ and its discretization M , and the elementary step $\Delta t = T_\theta/M$ and the reaction-time lattice

$RT(k) = t_0 + k\Delta t$ follow. A second user group is human–machine interaction engineering for wearable and assistive systems: [7] identifies intent detection on the sensor level as the start of the cognitive control chain and requires devices that "shall be quick and responsive"; under this framework the latency budget of that chain becomes a derived quantity — the number of reference steps k the intent-detection decision consumes, multiplied by the measured Δt — rather than a tuned one. A third group is experimental neuroscientists studying temporal binding: the framework supplies a concrete instantiation of a discrete integration window with a predicted width Δt and a predicted comb structure in reaction-time data, turning temporal binding into a measurement target.

(b) Readiness level. This work is purely analytical and has not been peer reviewed: by the NASA Technology Readiness Level definitions it sits at TRL 1–2 — basic principles observed and reported, concept formulated — and analytical work cannot exceed TRL 3. No hardware, no simulation result, and no empirical measurement exist for the Page–Wootters conditioning scheme itself, nor for the cocycle constraint sector.

(c) Smallest next test. The smallest de-risking calculation is a reanalysis of existing human reaction-time datasets: take simple and choice reaction-time distributions already collected at millisecond resolution, subtract a fitted latency floor t_0 , and test for excess histogram mass at multiples of s for s scanned over $[10, 16.67]$ ms, using a permutation or bootstrap null. The outcome that refutes the framework is a null result across the whole scanned range in datasets with more than 10^4 trials per subject; the outcome that de-risks it is a consistent spacing s within a subject that tracks that subject's theta frequency. The smallest next simulation is the enumeration left open in Section 4.5: build a 64-leaf tree, sample conditional distributions $P(Q | \varphi_i)$ under the joint constraint of Section 3.2, and count the cocycle-consistent subset of the $4^{10} = 1048576$ depth-2 histories; if that subset is empty or fails to contain jump-only trajectories, the conditioning scheme is inconsistent and the framework fails before any human experiment.

(d) Binding constraints. The constraints that would bind first, in order: (i) data — the comb test requires reaction-time histograms with bin width ≤ 5 ms and enough trials to resolve excess mass of a few percent, which existing datasets may or may not satisfy; (ii) measurement of the reference — the prediction depends on the subject's measured f_θ and on M , and no method is specified here for measuring M independently; (iii) observability of Q — the overlap matrix must be estimated from data, and noise induces spurious ultrametric violations (assumption A6); (iv) compute — the trajectory-space enumeration scales as b^M and is practical only for small N and M . Materials, manufacturing, and regulation do not bind at this stage because no hardware is proposed.

(e) Order-of-magnitude bounds (projections). All inputs here are derived in Section 4 or stated as assumptions. With $f_\theta \in [6, 10]$ Hz and $M = 10$, the elementary step spans $\Delta t \in [10, 16.67]$ ms; with $t_0 = 100$ ms assumed, the predicted simple-reaction-time lattice covers 112.5–250 ms for $k = 1$ to $k = 12$. These are projections under assumptions A1, A2, and A4; they are not measurements, and the uncertainty in M (A1) dominates: halving M to 5 doubles every spacing.

(f) Risks and misuse paths. The framework, if misapplied, could be used to claim that a measured reaction-time comb proves a specific neural mechanism, when a comb can also arise from display refresh rates, sampling pipelines, or motor periodicities; any adoption must include those artifacts as null hypotheses. No safety-critical deployment path exists at TRL 1–2, and none should be inferred from the quantization prediction alone.

8. Conclusion

We have formalized a conjecture that cognitive temporal order can be derived rather than imported. The semantic state space is a static ultrametric overlap matrix obeying a multiplicative cocycle condition; ordered evolution is extracted by conditioning the joint constrained distribution on the internally ordered phase states of a reference population, following the Page–Wootters conditional-time construction. The derivable bookkeeping is concrete: for $N = 64$ leaves the cocycle condition supplies $N_c = 41664$ equations on $N_f = 2016$ free overlaps (ratio $\rho = 62/3 \approx 20.67$), conditioning across $M = 10$ reference states adds at most $N_{PW} = 184320$ agreement conditions, and a theta reference

at 8 Hz with 10 bins yields an elementary step $\Delta t = 12.5$ ms and a reaction-time lattice $RT(k) = t_0 + 12.5 k$ ms with spacing in the 10–16.67 ms theta-band range. The construction's value, if it survives falsification, is that the timing parameters of a cognitive model stop being free: they are fixed by a measurable internal reference subsystem, and the same construction applies to biological and synthetic substrates alike. The falsification path is short and cheap — a comb test on existing reaction-time data and a trajectory-enumeration simulation — and either outcome sharpens the field's understanding of where cognitive time comes from.

Provenance and AI Involvement

- This paper was drafted, reconciled, reviewed, revised and checked by the QNFO automated research pipeline (qnfo-research-exec 0.18.43-codeagent). Language models were tools in that process; they are not authors. The named author is accountable for the content.
- Models: drafting by openai/gpt-oss-120b, zai-org/glm-5.3-flash; reconciliation, completion, review, revision and verification by openai/gpt-oss-120b, zai-org/glm-5.3-flash, moonshotai/kimi-k2.6, with the QNFO model router as a fallback.
- Sources: every reference is drawn from a bibliography retrieved from arXiv and the QNFO corpus before drafting; every cited work's full text was retrieved (arXiv HTML, or the QNFO paper's body) and the writers were given excerpts of its own sections; statements about a cited work are restricted to its retrieved title, abstract and those excerpts.
- Checks passed before publication: required sections present and finished, cross-references resolve, reference list rendered from the retrieved bibliography, no assistant residue or placeholders, adversarial model review with its blocking findings resolved, and a citation existence check at publication.
- Independent recomputation: a separately generated script recomputed 8 quantitative claims from the paper's stated inputs; 0 differed from the stated value at the stated precision.
- Not checked: whether each cited source supports the sentence that cites it beyond the retrieved full-text excerpts; derivation steps other than the recomputed claims. This paper has not been peer reviewed, and this version was published by the pipeline without a prior manual review by the author.

References

- [1] Cognitive computation with autonomously active neural networks: an emerging field. arXiv:0901.3028v1. <https://arxiv.org/abs/0901.3028v1> [2] Symbol Emergence in Cognitive Developmental Systems: a Survey. arXiv:1801.08829v2. <https://arxiv.org/abs/1801.08829v2> [3] Cognition and Emotion: Perspectives of a Closing Gap. arXiv:1002.3035v1. <https://arxiv.org/abs/1002.3035v1> [4] An embedding, an extension, and an interpolation of ultrametrics. arXiv:2008.10209v2. <https://arxiv.org/abs/2008.10209v2> [5] Space-time distributions. arXiv:gr-qc/9810059v1. <https://arxiv.org/abs/gr-qc/9810059v1> [6] A cognitive diversity framework for radar target classification. arXiv:1110.6589v1. <https://arxiv.org/abs/1110.6589v1> [7] Applied Exoskeleton Technology: A Comprehensive Review of Physical and Cognitive Human-Robot Interaction. arXiv:2111.12860v8. <https://arxiv.org/abs/2111.12860v8> [10] QNFO: Syntactic Generation