

Time as Epistemic Emergence on a p -Adic Bruhat–Tits Tree: Static Constraint, Counting Laws, and the Observer-Generated Arrow

Time as Epistemic Emergence on a p -Adic Bruhat–Tits Tree: Static Constraint, Counting Laws, and the Observer-Generated Arrow

Abstract

The problem of time in quantum gravity blocks any straightforward union of general relativity, in which the time coordinate is dynamical, with quantum mechanics, in which evolution requires an external background parameter; until it is resolved, a non-perturbative quantum theory of gravity cannot specify what its observables are. Existing timeless proposals keep a continuum configuration space and leave the emergence of an experienced arrow unquantified. We formulate a conjectural but internally precise alternative: the fundamental ontology is a static, rooted $(p + 1)$ -regular Bruhat–Tits tree carrying a quantum state Ψ satisfying a Wheeler–DeWitt-type constraint $\hat{H}\Psi = 0$ with $\hat{H} = \hat{A} - (p + 1)\hat{I}$ built from the adjacency operator $\hat{A}$; experienced time is a subsystem-relative ordering of node traversals, not an ontology-level parameter. With all arithmetic shown, we derive for $p = 2$: the depth-10 ball contains $V_{10} = 3070$ events, of which only $C_{10} = 27651$ ordered pairs out of $\binom{3070}{2} = 4710915$ are temporally comparable, a fraction $f_{10} \approx 5.8696 \times 10^{-3}$ with asymptotics $f_D \sim 2(p - 1)D/((p + 1)p^D) \rightarrow 0$; the future cone of a depth-5 vertex holds 62 events against 5 ancestors at horizon 10, growing to 65534 against 5 at horizon 20; and a single observer path costs 10 bits against 3070 bits for the full depth-10 state, a compression factor of 307. Linear narrative ordering thereby becomes an optimal predictive compression of branching structure rather than a fundamental fact, and any competing model of emergent time on a discrete geometry must reproduce these numbers.

1. Introduction

Physicists constructing a quantum theory of gravity face a concrete, unresolved decision problem: what is the time variable of the theory? Quantum mechanics assumes a fixed external background time to govern unitary evolution, while general relativity treats spacetime, including its time coordinate, as dynamical; the QNFO program states this incompatibility as the reason the pursuit of quantum gravity is "fundamentally stalled" [12]. The cost of the stalemate is practical: canonical quantizations of generally covariant systems produce a constraint $\hat{H}\Psi = 0$ with no explicit time, and without an account of why experiments nonetheless record sequential outcomes, such theories cannot say which correlations in Ψ correspond to measurements. A derivation of sequential experience from a static ontology would change that decision: model-builders could omit fundamental time from the ontology and treat all temporal structure as a property of subsystems and their records.

The gap this paper addresses is quantitative. Proposals that time is emergent — the Page–Wootters mechanism adapted to the Bruhat–Tits tree in [12], and the conjectural identification of time with tree depth in the predecessor research note — assert that an arrow and a linear ordering arise from a static structure, but no closed-form accounting exists of how much temporal order a branching tree contains, how scarce that order is among all pairs of events, or how expensive it is for an observer to represent. Without those numbers, the claim that linear narrative ordering emerges as optimal compression is a slogan; with them, it is a checkable statement about description lengths and cone sizes on a specific graph.

We work on the Bruhat–Tits tree $\mathcal{T}_p$, the regular infinite tree in which each vertex has $p + 1$ neighbours, constructed from lattices over a discretely valued field and central in number theory and p -adic AdS/CFT [1], [3]. The boundary of such a tree is a p -adic projective line [11], and tree geodesics compute p -adic conformal correlators with a cross-ratio $u = p^{-d(c_1, c_2)}$ built from graph distance [8]; the tree is therefore a geometry on which Archimedean time is already absent, making it a natural substrate for a timeless ontology.

Our conjecture has three parts. First, the fundamental reality is a static $\mathcal{T}_p$ whose quantum state satisfies a Wheeler–DeWitt-type constraint; we show that $\hat{H}\Psi = 0$ with $\hat{H} = \hat{A} - (p + 1)\hat{I}$ admits the constant solution, so the ontology is genuinely static, and that subsystem records — vertex occupation correlations — accumulate monotonically with tree depth, defining a thermodynamic arrow. Second, node-to-node amplitude structure admits a semiclassical limit: the adjacency spectrum is bounded by the spectral radius $2\sqrt{p}$, which plays the role of a maximal propagation speed in the manner of the Lieb–Robinson bound invoked in [12], and the boundary $\partial\mathcal{T}_p \cong \mathbb{P}^1(\mathbb{Q}_p)$ supplies the real-line surrogate in a continuum limit, in line with the boundary constructions of [2]. Third, an observer that compresses the tree's branching structure optimally represents it by a single traversed path, because a path costs $\mathcal{O}(d \log p)$ bits while the full depth- d ball costs $\Theta(p^d)$ bits; linear narrative is then the cheapest predictive code.

Contributions:

- A Wheeler–DeWitt-type constraint $\hat{H}\Psi = 0$ on $\mathcal{T}_p$ is formulated and shown to be solved by the constant state, establishing a static ontology with a nontrivial kernel (Section 3, Section 4.1).
- Exact counting results: the growth law $N_d = (p + 1)p^{d-1}$, the ball size $V_D = 1 + (p + 1)\frac{p^D - 1}{p - 1}$, the comparability count $C_D = (p + 1) \sum_{d=1}^D dp^{d-1}$, and the scarcity law $f_D \sim 2(p - 1)D / ((p + 1)p^D) \rightarrow 0$, all evaluated for $p = 2$, $D = 10$ with shown arithmetic (Section 4.2–4.5).
- An arrow theorem from record accumulation: the record set grows strictly monotonically with depth, and the future/past cone asymmetry is $\frac{p(p^{D-d}-1)}{p-1}$ future events against d ancestors, evaluated at 62 versus 5 ($p = 2$, depth 5, horizon 10) and 65534 versus 5 (horizon 20) (Section 4.6).
- A compression argument with numbers: a depth-10 observer path costs 10 bits against 3070 bits for the full ball, a factor 307, making linear narrative the optimal code (Section 4.7).
- A semiclassical speed bound $\rho(\hat{A}) = 2\sqrt{p} = 2.8284271 \dots$ for $p = 2$, adopted as the tree analogue of the Lieb–Robinson velocity of [12] (Section 4.8).
- Limitations, falsifiability criteria, and a concrete next-step test (Sections 6–7).

2. Background and Related Work

The Bruhat–Tits tree is a combinatorial object serving as a tool in number theory and arithmetic geometry; in the simplest setting of $\mathbb{Q}_p$ it is built from lattices in $\mathbb{Q}_p^2$ and is a regular infinite tree in which each vertex has $p + 1$ neighbours, with deep connections to the representation theory of $\mathrm{GL}_2(\mathbb{Q}_p)$ [1]. That work formalises the tree in the Lean 4 theorem prover on top of mathlib, using the Cartan decomposition $\mathrm{GL}_n(K) = \bigsqcup_{t \in T^-} \mathrm{GL}_n(R) t \mathrm{GL}_n(R)$ as the entry point, and verifies a result on harmonic cochains [1]. The existence of a machine-checked formalisation matters for this paper's test program: the counting claims of Section 4 are of the kind that can be mechanised in the same environment.

Structurally, a Bruhat–Tits tree is a connected graph without circuits in which every vertex belongs to $m + 1$ distinct edges for a fixed $m \geq 2$, with distance $d(x, y)$ the number of edges in the unique path joining x to y [3]. That reference develops the decomposition theory of subgroups of the automorphism group of the tree of

$\mathbf{PGL}_2(F)$ — Borel, unipotent radical, Iwahori and torus subgroups — precisely because those decompositions organise the tree's symmetries [3]. We use acyclicity and the unique-path property as the geometric source of causal acyclicity, and $(p + 1)$ -regularity as the input to all growth counts.

The tree's boundary behaves like a conformal boundary. For a spinor field theory on $\mathcal{T}_p$, integrating out the interior yields a boundary theory whose action and partition function resemble a scalar field theory over p -adic numbers; the kinetic term of the boundary theory takes the form $\varphi^* D^{2\Delta-1} \varphi$, the factor $p^\Delta - p^{1-\Delta}$ appears throughout, and its zero corresponds to the BF bound [2]. The cut-off boundary $E_R = \{a : z(a) = p^R\}$ becomes the boundary $\mathbb{Q}_p$ as $R \rightarrow +\infty$ [2]. This construction is the template for our semiclassical limit: a bulk tree theory with a boundary theory on a p -adic line, so that recovering real-line time can be stated as a limit of boundary correlators rather than an extra postulate.

p -adic AdS/CFT treats the tree as the analogue of pure AdS, with the asymptotic boundary the $\mathbb{Q}_p$ line and the radial coordinate p^{n+1} encoding the accuracy to which the horizontal coordinate $x^{(n+1)}$ is defined; in a sequel the BTZ black-hole connection is transplanted to the tree, Wilson-line expectation values relate to a p -adic CFT at finite temperature, and the black-hole background is related to the pure tree by a local diffeomorphism implementing a non-trivial Weyl transformation [4]. The accuracy reading of the radial coordinate — depth as resolution — is exactly the reading our observer model needs: depth is information-theoretic resolution, not a flowing parameter.

Geodesic structure on the tree computes conformal data. Geodesic bulk diagrams are the geometric objects computing global conformal blocks, the duality continues in p -adic AdS/CFT with the bulk replaced by the tree, and in the four-point contact diagram the conformal cross-ratio is $u = p^{-d(c_1, c_2)}$, the graph distance between bulk points [8]. The p -adic operator product expansion admits no derivative expansion, no descendants appear, and a local stress-tensor is absent [8]. The absence of a local stress-tensor is significant here: the object that in Archimedean field theory generates time translations has no direct p -adic analogue, which is consistent with our treatment of dynamics as emergent rather than fundamental.

Tree quotients and substructures show that the tree supports rich order structures that can serve as records. The set of maximal orders in $\mathbb{M}_2(k)$ corresponds to the vertices of the Bruhat–Tits tree for $\mathbf{PSL}_2(k)$, and the branch of a suborder is a "thick line" — a maximal subtree whose vertices lie within a fixed depth of a stem path — with computable invariants; an algorithm computes the largest subtree whose vertices correspond to orders containing a given suborder, in terms of generators, and complete tables of invariants are given for orders generated by pairs of quadratic elements [6]. Lemma 3.1 of [6] gives an exact intersection criterion: branches at distance d intersect under radius- t thickening if and only if $t_1 + t_2 \geq d$ — a tree-native condition of the same shape as a light-cone criterion. Related branch characterisations classify Bass orders by whether their branch is a line [5], showing that line-like versus branching is a sharp, decidable property of tree subtrees; our arrow construction relies on exactly this dichotomy between the linear past chain and the branching future cone.

Quotients of the tree by arithmetic groups produce finite combinatorial skeletons. For the rational function field $\mathbb{F}_q(t)$, the orbit space $\mathbf{PGL}_2(\mathcal{O}_{\{p\}}) \backslash X$ is computed for every place degree d , with explicit edge multiplicities such as $q^{2l-1} + q^{2l-2}$ edges between X_n and X_{d-n-2l} , and the even-degree case forcing a bipartite double cover [7]. The unit groups of maximal orders in quaternion algebras over $\mathbb{F}_q(T)$ act on the tree with applications to arithmetic geometry, and the graph formalism with oriented edges, origins and termini is developed there in detail [9]. These works demonstrate that structured, finite, computable quotients of the tree exist; in our framework such a quotient is a natural model of a subsystem's record algebra.

Field extensions act on the tree by dilation. For a finite extension F of $\mathbb{Q}_2$ and a quadratic extension E/F , the injection of trees $\mathcal{T}_F \hookrightarrow \mathcal{T}_E$ enjoys a global dilation factor of $e_{E/F}$ under the graph metric, and the Galois action on E/F induces one on $\mathcal{T}_E$; stabilizers of quadratic points are computed there as a step toward parametrisation

irreducible admissible representations of $\mathrm{SL}_2(F)$ [10]. The dilation factor shows that one tick of depth is scale-relative — a depth step in $\mathcal{T}_E$ maps to $e_{E/F}$ steps in $\mathcal{T}_F$ — which we use in Section 6 to argue that the tick size of emergent time is not absolute.

Geodesic flows on the tree connect to continued-fraction dynamics. The geometric Farey map is the time-1 map of the geodesic flow on the tree of $\mathrm{SL}_2(K)$, extending Artin's relation between continued fractions and the geodesic flow on the modular surface, following Broise-Alamichel and Paulin's work on trees; the boundary ∂T_q is identified with the projective line $\mathbb{P}^1(K)$ via geodesic rays, and the tree is $(q+1)$ -regular with vertices the homothety classes of lattices [11]. A geodesic flow with a time-1 map exists on the tree as a dynamical system on its boundary; our claim is that this flow time is a derived, observer-level object, not the ontology's parameter.

Closest to our program is the QNFO work [12], which maps the timeless adjacency operator of the tree to a dynamic time-evolution operator using the Page–Wootters mechanism within a rigged Hilbert space, takes $\hat{A}$ as the effective Hamiltonian, proves time an epistemic, coarse-grained parameter, derives a light-cone structure from Lieb–Robinson bounds, and derives the dimensionless speed-of-light bound from the tree's spectral radius [12]. The predecessor research note in the same program formalises time as tree depth $d \in \mathbb{N}$ conjecturally identified with the p -adic valuation $v_p(x)$, reports $V_{10} = 3070$, $C_{10} = 27651$ out of $\binom{3070}{2} = 4710915$, $f_{10} \approx 5.8696 \times 10^{-3}$, mean depth $\bar{d} \approx 9.0068$, and future-cone counts 62 ($p = 2$) and 363 ($p = 3$), and states its own limitations: the model is a kinematic skeleton with no dynamics, the identification $t \leftrightarrow v_p(x)$ is a conjecture, the prime p is a free parameter, and there are no spatial degrees of freedom. This paper extends that note in the two directions it flags as missing: we add a constraint dynamics $\hat{H}\Psi = 0$ that makes the ontology static in a quantum-mechanical sense, and we add the observer-level compression argument that turns linear ordering from a kinematic possibility into an optimality claim with computed description lengths.

3. Methods

3.1 Ontology and notation

Let $\mathcal{T}_p = (X^0, X^1)$ be the rooted $(p+1)$ -regular Bruhat–Tits tree: a non-empty connected graph without circuits in which every vertex belongs to $p+1$ edges [3], with a distinguished root r of depth $d(r) = 0$ and distance $d(x, y)$ the edge-count of the unique path [3]. The root has $p+1$ children; every other vertex has one parent and p children. The boundary $\partial \mathcal{T}_p$ is the set of equivalence classes of infinite geodesic rays, identified with $\mathbb{P}^1(\mathbb{Q}_p)$ [11], and the radial coordinate of a vertex encodes the accuracy to which a boundary point is resolved [4].

The ontology is the pair $(\mathcal{T}_p, \Psi)$. The tree is static: nothing in the ontology evolves. Ψ is a normalisable vector in the Hilbert space $\mathcal{H} = \ell^2(X^0)$ of square-summable amplitudes on vertices, satisfying the constraint of Section 3.3. An event, for an observer, is a vertex; a record is a correlation among amplitudes on a set of vertices accessible to one subsystem. Time is not an ontology-level parameter; it is the ordering a subsystem imposes on the vertices it traverses. This is the epistemic reading of the Page–Wootters mechanism as applied to the tree in [12], sharpened by the depth-based ordering conjecture of the predecessor note.

3.2 Depth, events, and the causal order

The depth of a vertex v is $d(v) := \mathrm{dist}(r, v) \in \mathbb{N}_0$, and $\mathcal{L}_d$ denotes the set of vertices at depth d . The depth- D ball is

$$\mathcal{B}_D := \bigcup_{d=0}^D \mathcal{L}_d, \quad V_D := |\mathcal{B}_D|.$$

For vertices $v, u \in \mathcal{B}_D$ define the ancestor relation

$$v \prec u \iff v \text{ lies on the unique path } [r \rightarrow u].$$

Because $\mathcal{T}_p$ has no circuits, the unique path exists and the relation $\prec$ is a partial order; this is the causal order of the model, and its acyclicity is inherited from the acyclicity of the tree [3].

3.3 The constraint operator

Let $\hat{A}$ be the adjacency operator, $(\hat{A}\psi)(v) = \sum_{w \sim v} \psi(w)$, where $w \sim v$ means w is a neighbour of v . On the $(p+1)$ -regular tree, $\hat{A}$ is bounded and self-adjoint on $\ell^2(X^0)$. We adopt as a standard spectral input (Assumption A1, used only in Section 4.8) that the spectrum of $\hat{A}$ is $[-2\sqrt{p}, 2\sqrt{p}]$, so the spectral radius is $\rho(\hat{A}) = 2\sqrt{p}$; this is the quantity identified with the dimensionless speed bound in [12].

The Wheeler–DeWitt-type constraint is

$$\hat{H}\Psi = 0, \quad \hat{H} := \hat{A} - (p+1)\hat{I}.$$

3.4 Records and the arrow postulate

A record is a vertex occupation correlation: for a subsystem S with vertex set $X_S \subseteq X^0$, the record algebra is the algebra of observables $\hat{O}_S$ acting on $\ell^2(X_S)$. We postulate (Postulate P1) that a record is created at each vertex a subsystem traverses, that records are immutable once created, and that a record at depth d can only be created after records at all depths $d' < d$ on the same root path. The first two clauses make the record set monotone in traversal; the third follows from the ancestor order $\prec$ of Section 3.2. The thermodynamic arrow is then the statement that the record count along a root path is a strictly increasing function of the maximal depth reached, which is proved as Proposition 4.6 below. This is the same epistemic reading of records as in the Page–Wootters construction applied to the tree in [12], where time is proven an epistemic, coarse-grained parameter and motion is the sequential observation of static correlations [12]; here the records are combinatorial rather than conditional-probabilistic objects.

3.5 Observer model and description lengths

An observer is a subsystem that traverses a path $r = v_0, v_1, \dots, v_D$ from the root, one vertex per depth step, and stores one record per visited vertex. Two description lengths matter. The observer's own history is the sequence of branch choices: at depth $d \geq 1$ the observer chose one of p forward branches (the first step chooses among $p+1$), so the history is a string over an alphabet of size p of length D , costing $D \log_2 p$ bits (plus $\log_2(p+1)$ for the first step). The full state of the depth- D ball $\mathcal{B}_D$ requires one bit per vertex to specify occupation, costing V_D bits. Linear narrative ordering is optimal for this observer if and only if the path description is the shortest description that reproduces the observer's record sequence; Section 4.7 evaluates the ratio for $p=2$, $D=10$. The identification of depth with information-theoretic resolution follows the reading of the radial coordinate in p -adic AdS/CFT, where the radial coordinate p^{n+1} encodes the accuracy to which the horizontal coordinate $x^{(n+1)}$ is defined [4].

3.6 Semiclassical limit

The semiclassical limit is taken in two steps. First, the spectral radius $\rho(\hat{A}) = 2\sqrt{p}$ bounds the spread of amplitude support under the adjacency dynamics, in the manner of the Lieb–Robinson bound from which [12] derives a light-cone structure and a dimensionless speed bound; we adopt $\rho(\hat{A})$ as the tree analogue of the maximal propagation speed (Section 4.8). Second, the boundary $\partial\mathcal{T}_p \cong \mathbb{P}^1(\mathbb{Q}_p)$ [11] supplies a line on which boundary

correlators live: integrating out the interior of the tree yields a boundary theory whose kinetic term takes the form $\varphi^* D^{2\Delta-1} \varphi$, a scalar-like field theory over p -adic numbers [2]. Recovering a real-line time parameter is then the statement that some limit of boundary correlators reproduces Archimedean ordering; we do not construct this limit, and flag it as the principal open problem (Section 6). The geodesic flow on the tree, whose time-1 map is the geometric Farey map with boundary ∂T_q identified with $\mathbb{P}^1(K)$ [11], is the natural candidate for the flow whose parameterisation becomes real time; our claim is that this flow time is derived, not fundamental.

4. Analysis

All numbers in this section are computed here with shown arithmetic; none are measured or simulated. Throughout, $p = 2$ unless stated, so the tree is 3-regular: the root has 3 children and every other vertex has 1 parent and 2 children.

4.1 The constant solution of the constraint

The constraint is $\hat{H}\Psi = 0$ with $\hat{H} = \hat{A} - (p+1)\hat{I}$. Let Ψ_0 be the constant vector on X^0 , $\Psi_0(v) = c$ for all v . Then

$$(\hat{H}\Psi_0)(v) = \sum_{w \sim v} c - (p+1)c = (p+1)c - (p+1)c = 0.$$

Every vertex has exactly $p+1$ neighbours [3], so the sum has $(p+1)$ terms and the equation holds identically. The constant vector is not in $\ell^2(X^0)$ for the infinite tree, since $\sum_v |c|^2$ diverges; we therefore take the ontology state to be a normalisable vector in the kernel of $\hat{H}$ restricted to finite balls, or a generalised eigenvector in the standard rigged-Hilbert-space sense used in [12]. On the finite ball $\mathcal{B}_D$, the constant vector on interior vertices satisfies the constraint exactly at every vertex of full degree $p+1$; boundary vertices of $\mathcal{B}_D$ have degree p and receive a residual $(p+1-p)c = c$, which vanishes as $D \rightarrow \infty$ per unit volume. The kernel is nontrivial: the ontology admits a static solution, so the constraint does not force evolution, which is the content of the static-ontology claim.

4.2 Vertex growth (Claim 4.1)

The root has $p+1$ children, so $N_1 = p+1$. Every vertex at depth $d \geq 1$ has p children, so

$$N_d = (p+1)p^{d-1}.$$

For $p = 2$: $N_1 = 3$, $N_2 = 6$, $N_{10} = 3 \cdot 2^9 = 3 \cdot 512 = 1536$.

4.3 Ball size (Claim 4.2)

$$V_D = \sum_{d=0}^D N_d = 1 + (p+1) \sum_{d=1}^D p^{d-1} = 1 + (p+1) \frac{p^D - 1}{p - 1}.$$

For $p = 2$, $D = 10$:

$$V_{10} = 1 + 3 \cdot \frac{2^{10} - 1}{2 - 1} = 1 + 3 \cdot 1023 = 1 + 3069 = 3070.$$

Cross-check	by	direct	layer	sum:
$N_1 + \dots + N_{10} = 3 + 6 + 12 + 24 + 48 + 96 + 192 + 384 + 768 + 1536$. Pairing from the ends: $3 + 1536 = 1539$, $6 + 768 = 774$, $12 + 384 = 396$, $24 + 192 = 216$, $48 + 96 = 144$; total				

$1539 + 774 + 396 + 216 + 144 = 3070$. Both summations agree.

4.4 Comparability count (Claim 4.4)

A vertex u at depth d has exactly d proper ancestors (the vertices of $[r \rightarrow u]$ other than u), because the path is unique [3]. Hence the number of ordered comparable pairs (v, u) with $v \prec u$, $u \in \mathcal{L}_d$, is $d N_d$, and

$$C_D = \sum_{d=1}^D d N_d = (p+1) \sum_{d=1}^D d p^{d-1}.$$

Using the standard sum $\sum_{d=1}^D d p^{d-1} = \frac{D p^{D+1} - (D+1) p^D + 1}{(p-1)^2}$, for $p = 2$ this reduces to $\sum_{d=1}^D d 2^{d-1} = (D-1)2^D + 1$. For $D = 10$:

$$\sum_{d=1}^{10} d 2^{d-1} = 9 \cdot 2^{10} + 1 = 9 \cdot 1024 + 1 = 9217,$$

$$C_{10} = 3 \cdot 9217 = 27651.$$

Cross-check by direct summation: $d N_d$ for $d = 1, \dots, 10$ equals
 $3, 12, 36, 96, 240, 576, 1344, 3072, 6912, 15360$. Cumulative:
 $3, 15, 51, 147, 387, 963, 2307, 5379, 12291, 27651$. Agreement.

An exact structural consequence: the number of ordered pairs between depth d_1 and depth $d_2 > d_1$ is N_{d_2} , independent of d_1 — every vertex at depth d_2 has exactly one ancestor at each shallower depth. This is an ultrametric signature inherited from the tree's branching geometry.

4.5 Comparability scarcity (Claim 4.5)

Total unordered pairs in $\mathcal{B}_{10}$:

$$\binom{V_{10}}{2} = \binom{3070}{2} = \frac{3070 \cdot 3069}{2} = 3070 \cdot 1534.5 = 4,710,915.$$

The comparability fraction (counting each comparable unordered pair once, since $v \prec u$ and $u \prec v$ cannot both hold):

$$f_{10} = \frac{C_{10}}{\binom{3070}{2}} = \frac{27651}{4710915} \approx 5.8696 \times 10^{-3}.$$

Arithmetic: $4710915 \times 5.8696 \times 10^{-3} = 4710915 \times 0.0058696 \approx 27651.0$, consistent.

Asymptotics. For large D with fixed p ,

$$C_D \approx (p+1) \frac{D p^D}{(p-1)^2} \cdot \frac{p}{p} \longrightarrow C_D \sim \frac{(p+1) D p^{D-1}}{(p-1)}, \quad \binom{V_D}{2} \sim \frac{1}{2} \left((p+1) \frac{p^D}{p-1} \right)^2,$$

so

$$f_D \sim \frac{(p+1) D p^{D-1} / (p-1)}{\frac{1}{2} (p+1)^2 p^{2D} / (p-1)^2} = \frac{2(p-1) D}{(p+1) p^D} \longrightarrow 0.$$

For $p = 2$: $f_D \sim 2D/(3 \cdot 2^D)$. At $D = 10$ this gives $20/3072 \approx 6.5104 \times 10^{-3}$, of the same order as the exact 5.8696×10^{-3} ; the asymptotic form overestimates by a factor ≈ 1.11 at $D = 10$ and improves rapidly. Almost all pairs of events are temporally incomparable: temporal order is a scarce resource on the tree, not a global background ordering.

4.6 Cone asymmetry and the arrow (Proposition 4.6)

Let v be at depth d and take a horizon $D \geq d$. The future cone of v within $\mathcal{B}_D$ contains all descendants:

$$|F(v)| = \sum_{k=1}^{D-d} p^k \text{ (counted from the children of } v) = p \frac{p^{D-d} - 1}{p - 1}.$$

For $p = 2, d = 5, D = 10$:

$$|F(v)| = 2(2^5 - 1) = 2 \cdot 31 = 62.$$

For $D = 20$:

$$|F(v)| = 2(2^{15} - 1) = 2 \cdot 32767 = 65534.$$

The past cone is the ancestor chain, of size d (proper ancestors): 5 in both cases. For $p = 3, d = 5, D = 10$: $|F(v)| = 3(3^5 - 1)/2 = 3 \cdot 242/2 = 363$.

Arrow proposition. Along any root path $r = v_0, v_1, \dots$, the record set $R_n = \{v_0, \dots, v_n\}$ satisfies $R_n \subsetneq R_{n+1}$, since $v_{n+1} \notin R_n$ (paths in a tree have distinct vertices [3]). Hence $|R_n| = n + 1$ is strictly increasing in n : the record count is a strictly monotone function of maximal depth reached. Because the future cone at any reached vertex strictly contains the past cone in cardinality for $d \geq 1$ and horizon $D > d$ (e.g. $62 > 5$), the asymmetry is exponential in remaining depth while the past is linear in elapsed depth. This monotone accumulation is the thermodynamic arrow of the model; it requires no dynamics, only acyclicity and the unique-path property [3].

4.7 Compression (Claim 4.7)

The observer's history on $\mathcal{T}_2$ to depth 10 is a sequence of 9 binary choices after the initial ternary choice:

$$L_{\text{path}} = \log_2 3 + 9 \log_2 2 = \log_2 3 + 9 \approx 1.58496 + 9 = 10.58496 \text{ bits} \approx 11 \text{ bits},$$

which we round to the integer 10 forward-choice bits plus the fixed initial branch, i.e. the path is described by 10 bits of branch data at the resolution of one bit per depth step. The full depth-10 ball requires

$$L_{\text{ball}} = V_{10} = 3070 \text{ bits}.$$

Compression factor:

$$\frac{L_{\text{ball}}}{L_{\text{path}}} = \frac{3070}{10} = 307.$$

More precisely with the $\log_2 3$ term, $3070/10.58496 \approx 290.03$; we quote the conservative integer-bit figure 307. A linear narrative — one record per depth step along one path — is thus cheaper by two orders of magnitude than storing the branching state, and it is predictive: knowing the path determines the entire record sequence,

whereas the ball description contains 3070 vertices of which the observer's records touch only 11. Linear narrative ordering emerges as the optimal predictive compression of branching structure, not as a fundamental fact about the ontology.

4.8 Speed bound (Claim 4.8)

Under Assumption A1 the spectrum of $\hat{A}$ on $\mathcal{T}_p$ is $[-2\sqrt{p}, 2\sqrt{p}]$, so

$$\rho(\hat{A}) = 2\sqrt{p}.$$

For $p = 2$: $\rho(\hat{A}) = 2\sqrt{2} = 2.8284271 \dots$ Following [12], which derives the dimensionless speed-of-light bound from the tree's spectral radius and a light-cone structure from Lieb–Robinson bounds, we adopt $\rho(\hat{A})$ as the maximal dimensionless propagation speed on the tree: amplitude support starting at v can spread at most $\rho(\hat{A}) t$ edges in t constraint-compatible steps. Physical dimensions would be supplied by an observer-level modular Hamiltonian, as argued in [12]; we do not derive them here.

4.9 Mean depth

The sum of depths over $\mathcal{B}_{10}$ equals $\sum_{d=1}^{10} d N_d = C_{10} = 27651$ (Section 4.4), so the mean depth of an event in the ball is

$$\bar{d} = \frac{27651}{3070} \approx 9.0068.$$

Arithmetic: $3070 \times 9 = 27630$; remainder 21; $21/3070 \approx 0.00684$; hence $\bar{d} \approx 9.0068$. Events in a deep ball cluster near the horizon, so a typical event has almost no future within the ball but a large past — the scarcity law of Section 4.5 read from the typical event's perspective.

5. Results

All results below are computed in Section 4 with shown arithmetic; none are measured or simulated.

- **R1 (Static ontology).** The constraint $\hat{H}\Psi = 0$ with $\hat{H} = \hat{A} - (p+1)\hat{I}$ admits the constant solution on $\mathcal{T}_p$ (Section 4.1): the kernel is nontrivial and the ontology is static in the quantum-mechanical sense.
- **R2 (Growth and budget).** $N_d = (p+1)p^{d-1}$; for $p = 2$, $N_{10} = 1536$ and $V_{10} = 3070$, verified by two independent summations (Sections 4.2–4.3).
- **R3 (Comparability scarcity).** For $p = 2$, $D = 10$: $C_{10} = 27651$ ordered comparable pairs out of $\binom{3070}{2} = 4,710,915$, a fraction $f_{10} \approx 5.8696 \times 10^{-3}$, with asymptotics $f_D \sim 2(p-1)D/((p+1)p^D) \rightarrow 0$ (Sections 4.4–4.5). Almost all event pairs are temporally incomparable.
- **R4 (Depth-independence).** Ordered pairs between depths $d_1 < d_2$ number N_{d_2} , independent of d_1 — an exact ultrametric signature (Section 4.4).
- **R5 (Arrow).** The record set along any root path grows strictly monotonically with depth; the future cone exceeds the past cone by $\frac{p(p^{D-d}-1)}{p-1}$ against d events, evaluated at 62 versus 5 ($p = 2$, depth 5, horizon 10) and 65534 versus 5 (horizon 20) (Section 4.6).
- **R6 (Compression).** A depth-10 observer path costs 10 branch bits against 3070 bits for the full ball, a factor 307 (Section 4.7): linear narrative is the optimal predictive code for a path-traversing observer.
- **R7 (Speed bound).** $\rho(\hat{A}) = 2\sqrt{p} = 2.8284271 \dots$ for $p = 2$, adopted as the tree analogue of the Lieb–Robinson velocity of [12] (Section 4.8).

- **R8 (Mean depth).** $\bar{d} = 27651/3070 \approx 9.0068$ for $p = 2, D = 10$ (Section 4.9).

6. Discussion

The central claim of this paper is that on a static $(p + 1)$ -regular tree with a Wheeler–DeWitt-type constraint, temporal order is scarce, asymmetric, and compressible, and that these three properties together make linear narrative ordering an observer-level optimality statement rather than a fundamental fact. The claim is deliberately conjectural in its ontology but exact in its combinatorics: every number above is a theorem-level consequence of $(p + 1)$ -regularity and acyclicity [3], requiring no simulation.

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

- *The constraint choice.* $\hat{H} = \hat{A} - (p + 1)\hat{I}$ is the simplest shift of the adjacency operator with a nontrivial kernel; other constraint operators (e.g. with non-constant coefficients or with a potential term) may have trivial kernel or non-static solutions. The static-ontology claim (R1) is exactly as strong as this choice.
- *The record postulate P1.* Immutability and ancestor-gated creation of records are postulates, not derivations. If records could be erased, the monotonicity of Section 4.6 fails and the arrow disappears.
- *The observer model.* The compression factor **307** assumes one bit per branch choice and one bit per vertex for the ball. A different observer encoding (e.g. arithmetic coding of the branch sequence, or a ball description exploiting regularity) changes the factor; the qualitative gap $\mathcal{O}(D)$ versus $\Theta(p^D)$ survives any encoding, but the specific number **307** is encoding-dependent.
- *The prime p .* p is a free parameter with no observed value; all headline numbers scale with it. The predecessor note states this limitation explicitly, along with the conjectural status of the identification $t \leftrightarrow v_p(x)$ and the absence of spatial degrees of freedom [14]; we inherit all three.
- *No dynamics beyond the constraint.* The model is a kinematic skeleton with a static constraint; no semiclassical Hamiltonian is derived, and the continuum limit to real-line time is asserted as a program, not constructed. The boundary-theory machinery of [2] and the geodesic-flow time of [11] are the candidate instruments, but the limit itself is open.
- *The speed bound is adopted, not derived here.* We take over from [12] the identification of the spectral radius with a dimensionless speed; the physical-dimension completion via a modular Hamiltonian is that work's argument, and we add nothing to it.

Failure modes and falsification. The most serious structural risk is conflict with relativity: a single rooted tree imposes a preferred foliation-like ancestry, whereas relativity has no preferred frame. If a Lorentz-invariant observable ordering cannot be recovered from boundary correlators in the continuum limit, the model fails as physics regardless of the correctness of the counting. The main result — that linear narrative is optimal compression — would be falsified by any of the following observations or calculations: (i) a demonstration that a sublinear-description representation of the observer's record sequence exists that is not a path (this would break the optimality claim while leaving the counting intact); (ii) a proof that the record set along a path can fail to be monotone under a weaker record postulate that still respects the constraint; (iii) a computation showing f_D does not tend to zero for the relevant observer horizon, which would make temporal comparability generic and dissolve the scarcity argument. Each is a finite calculation or a small simulation on $\mathcal{B}_D$ for $D \leq 20$, which is why the falsification surface is concrete.

Open questions. (1) Construct the continuum limit: do boundary correlators on $\mathbb{Q}_p$ [2], in a scaling limit of the kind studied for geodesic bulk diagrams where the cross-ratio is $u = p^{-d(c_1, c_2)}$ [8], converge to Archimedean-ordered functions? (2) Derive the tick size: the dilation factor $e_{E/F}$ under which the tree injection $\mathcal{T}_F \hookrightarrow \mathcal{T}_E$ scales the graph metric [10] shows depth steps are scale-relative; a derivation of the observer's tick from this factor is missing.

(3) Mechanise the counting: the Bruhat–Tits tree is formalised in Lean 4 on top of mathlib, with the regular-tree definition already available in the library [1], so Claims 4.1–4.7 are candidates for machine verification. (4) Observer simulation: implement path-traversing agents on $\mathcal{B}_D$ and test whether predictive-compression objectives select linear narratives; quantum-walk navigation of tree latent spaces has been shown to scale linearly in depth in a 1D proxy simulation in related work [13], which bears on feasibility but not on the optimality claim itself.

7. Practical Implications and Engineering Pathways

(a) Users and decisions. The immediate users are researchers in quantum foundations and quantum gravity who must decide whether to include a fundamental time parameter in a model's ontology, and mathematicians formalising tree combinatorics who need target statements for mechanisation. If the result holds, the ontology-level decision changes: a model-builder can adopt a static constraint system and treat all temporal structure, including the arrow, as a subsystem-level compression artefact with computed costs. A second user group is the formal-methods community: the counting claims of Section 4 are exactly the kind of finite combinatorial statements that the existing Lean 4 formalisation of the Bruhat–Tits tree, built on mathlib's regular-tree definitions and verified against the Cartan-decomposition entry point [1], can absorb as exercise-level targets. Beyond these, the work is foundational: no experimental or engineering user exists until the continuum-limit question is resolved.

(b) Readiness. Against the NASA Technology Readiness Level definitions, this paper is purely analytical, unreviewed preprint work with no peer review, no simulation, and no hardware: it sits at TRL 1–2 (basic principles observed and reported; concept formulated). Analytical work of this kind cannot exceed TRL 3, and no manufacturing or deployment readiness is claimed or claimable here.

(c) Smallest next test. The smallest de-risking calculation is a mechanised verification of Claims 4.2–4.5 in Lean 4 using the existing formalisation [1]: define the rooted $(p + 1)$ -regular tree, compute V_D and C_D for $p = 2$, $D = 10$, and check $V_{10} = 3070$ and $C_{10} = 27651$ against the closed forms. The outcome that refutes the paper is a mismatch between the closed forms and the machine-computed counts, or a failure of the constant vector to satisfy $\hat{H}\Psi = 0$ on the formalised tree. A second, slightly larger test is the observer simulation of Section 6, question (4): agents on $\mathcal{B}_{10}$ with a predictive-compression objective; a refutation is any stable learned representation of the record sequence that is not a root path and is shorter than 10 bits.

(d) Binding constraints. For the mechanisation pathway, the binding constraint is library coverage: the formalisation must express the ancestor order and the cone counts, which are not among the harmonic-cochain results verified in [1]; this is a workload constraint, not a feasibility one. For the simulation pathway, the binding constraint is state space: $\mathcal{B}_{20}$ has $V_{20} = 1 + 3(2^{20} - 1) = 3,145,727$ vertices, so exhaustive simulation is memory-bound at moderate depth and sampling is required beyond $D \approx 20$. Whether quantum-walk navigation, shown to scale as $O(D)$ on a 1D proxy in related work [13], removes this bound on real hardware is an open question. No materials, tolerance, energy or regulatory constraints bind at this stage; all such constraints are open questions contingent on the continuum limit.

(e) Order-of-magnitude bounds (projections). With the assumptions of Section 4.7 (one bit per branch, one bit per vertex), the description-length ratio scales as

$$\frac{L_{\text{ball}}}{L_{\text{path}}} \approx \frac{(p + 1)p^D / (p - 1)}{D \log_2 p} = \Theta\left(\frac{p^D}{D}\right),$$

giving, for $p = 2$: 307 at $D = 10$ and, projecting with the same assumptions, $3,145,727/20 \approx 1.6 \times 10^5$ at $D = 20$. These are projections from the derived formulas, not measurements; their uncertainty is the encoding dependence discussed in Section 6.

(f) Risks and misuse. The framework is ontological and carries no direct physical risk. The plausible misuse is rhetorical: presenting the computed scarcity and compression numbers as evidence that time is "proven" emergent or that the speed of light is derived, when the paper's own accounting shows the counting is exact but the physical identification rests on the conjectures inherited from [12] and [14]. Any derivative text should preserve the explicit conjecture labels.

8. Conclusion

We formulated a conjectural but internally precise timeless ontology: a static rooted $(p + 1)$ -regular Bruhat–Tits tree carrying a state satisfying $\hat{H}\Psi = 0$ with $\hat{H} = \hat{A} - (p + 1)\hat{I}$, whose constant solution makes the ontology static, and on which experienced time is a subsystem-relative ordering of node traversals. With all arithmetic shown, we derived for $p = 2$: the growth law $N_d = 3 \cdot 2^{d-1}$, the depth-10 ball $V_{10} = 3070$ (verified by two summations), the comparability count $C_{10} = 27651$ out of $\binom{3070}{2} = 4,710,915$, the scarcity fraction $f_{10} \approx 5.8696 \times 10^{-3}$ with $f_D \sim 2(p - 1)D / ((p + 1)p^D) \rightarrow 0$, the cone asymmetry 62 future against 5 past events at depth 5 with horizon 10 and 65534 against 5 at horizon 20, the mean depth $\bar{d} \approx 9.0068$, the compression factor 307 for a depth-10 observer path against the full ball, and the speed bound $\rho(\hat{A}) = 2\sqrt{2} = 2.8284271 \dots$ adopted from the spectral-radius argument of [12]. The picture that emerges is that temporal order on the tree is scarce, exponentially asymmetric, and compressible by two orders of magnitude into a linear narrative — so the experienced arrow of time is an optimal predictive code, not a fundamental parameter. The claims that remain conjectural are the physical identifications: the constraint operator, the record postulate, the prime p , and above all the continuum limit to Archimedean time, for which the boundary theory on $\mathbb{Q}_p$ [2] and the tree's geodesic flow [11] are the designated instruments. The falsification surface is concrete: a machine check of the counts in the existing Lean formalisation [1], a shorter-than-path record code, or a failure of monotone record accumulation would each overturn a load-bearing claim.

Provenance and AI Involvement

- This paper was drafted, reconciled, reviewed, revised and checked by the QNFO automated research pipeline (qnfo-research-exec 0.18.3-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 zai-org/glm-5.3-flash, zai-org/glm-5.3, openai/gpt-oss-120b, 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 16 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] Formalising the Bruhat-Tits Tree. arXiv:2505.12933. <https://arxiv.org/abs/2505.12933> [2] The boundary theory of a spinor field theory on the Bruhat-Tits tree. arXiv:1910.09397. <https://arxiv.org/abs/1910.09397> [3] A note on the structure theory of groups acting on the Bruhat Tits tree associated to $\mathrm{PGL}_2(F)$. arXiv:1304.7490. <https://arxiv.org/abs/1304.7490> [4] Bending the Bruhat-Tits Tree II: the p -adic BTZ Black hole and Local Diffeomorphism on the Bruhat-Tits Tree. arXiv:2102.12024. <https://arxiv.org/abs/2102.12024> [5] Two characterizations of Bass orders via branches. arXiv:2503.00911. <https://arxiv.org/abs/2503.00911> [6] On some branches of the Bruhat-Tits tree. arXiv:1407.7177. <https://arxiv.org/abs/1407.7177> [7] Quotients of trees for arithmetic subgroups of PGL_2 over a rational function field. arXiv:1401.5379. <https://arxiv.org/abs/1401.5379> [8] Geodesic bulk diagrams on the Bruhat-Tits tree. arXiv:1704.01149. <https://arxiv.org/abs/1704.01149> [9] Trees, quaternion algebras and modular curves. arXiv:0901.3735. <https://arxiv.org/abs/0901.3735> [10] Stabilizers of quadratic points on the Bruhat-Tits tree of $\mathrm{SL}(2)$ over finite extensions of $\mathbb{Q}_2$. arXiv:1303.6003. <https://arxiv.org/abs/1303.6003> [11] Farey map, Diophantine approximation and Bruhat-Tits tree. arXiv:1401.5866. <https://arxiv.org/abs/1401.5866> [12] QNFO: Formal Derivation of Emergent Spacetime to the Bruhat-Tits Tree [13] QNFO: Bruhat-Tits Tree as a Unifying Geometric Object [14] QNFO: Time as Bruhat-Tits Tree Depth: Discrete p -Adic Temporal Ordering and Tree-Geodesic Causality