

V1.0.0

Superposition as a Projection Artifact of Discrete Tree Geometry

Rowan Brad Quni-Gudzinis

QNFO-ULA

ABSTRACT

Quantum superposition underlies virtually every quantum technology, yet its ontological status remains contested. The present work addresses the gap between the continuous Hilbert-space formalism and a hidden-variable ontology by proposing that superposition arises from projecting a single, definite trajectory through a discrete, depth-organized tree of configurations onto a resolution-limited measurement basis. We construct an explicit binary-branching tree model, derive the coarse-graining that yields the standard wavefunction, and show that the Born rule is consistent with branch-counting combined with a renormalisation-group scaling of branch volumes. A concrete numerical example with depth $d = 5$ demonstrates that the derived probability $P_A = 0.625$ matches $|a|^2 = 0.625$ for amplitudes $a = \sqrt{0.625}$ and $b = \sqrt{0.375}$. The result suggests that the model can reproduce diagonal Born probabilities without invoking ontological multiplicity, though a full resolution of the measurement problem remains open, and that experimental tests can be designed by probing the discreteness scale of the underlying tree.

CITE AS Quni-Gudzinis, R. B. (n.d.). Superposition as a Projection Artifact of Discrete Tree Geometry. QNFO Preprint v1.0.0.

Abstract

Quantum superposition underlies virtually every quantum technology, yet its ontological status remains contested. The present work addresses the gap between the continuous Hilbert-space formalism and a hidden-variable ontology by proposing that superposition arises from projecting a single, definite trajectory through a discrete, depth-organized tree of configurations onto a resolution-limited measurement basis. We construct an explicit binary-branching tree model, derive the coarse-graining that yields the standard wavefunction, and show that the Born rule is consistent with branch-counting combined with a renormalisation-group scaling of branch volumes. A concrete numerical example with depth $d = 5$ demonstrates that the derived probability $P_A = 0.625$ matches $|a|^2 = 0.625$ for amplitudes $a = \sqrt{0.625}$ and $b = \sqrt{0.375}$. The result suggests that the model can reproduce diagonal Born probabilities without invoking ontological multiplicity, though a full resolution of the measurement problem remains open, and that experimental tests can be designed by probing the discreteness scale of the underlying tree.

1. Introduction

Laboratories building quantum computers, interferometers, and macroscopic-superposition experiments face a persistent conceptual obstacle: the measurement problem. Engineers must design error-correction protocols, yet the underlying theory offers no clear account of why a single outcome is observed when the formalism predicts a superposition of outcomes. If the superposition were merely an artifact of a hidden discrete structure, then the design of fault-tolerant hardware could be guided by the scale at which the tree discretisation becomes relevant, potentially reducing overheads.

The present paper shows that the apparent superposition observed in double-slit interference, qubit rotations, and large-scale Bose-Einstein-condensate experiments can be reproduced by a single deterministic path through a binary tree whose depth encodes the resolution of the measurement apparatus. Collapse then corresponds to Bayesian updating

of the observer's knowledge about which leaf of the tree was traversed. The Born rule emerges from the relative count of leaves associated with each coarse-grained branch.

Contributions

- **C1.** Definition of a discrete binary-branching tree model for quantum configurations and derivation of the coarse-graining map to a wavefunction (Section 3).
- **C2.** Analytical proof that branch-counting yields the Born rule $P = |a|^2$ for any two-path superposition (Section 4).
- **C3.** Numerical illustration with depth $d = 5$ showing $P_A = 0.625$ matches the squared amplitude (Section 5).
- **C4.** Discussion of experimental signatures and a minimal pilot test based on macroscopic superposition of a Bose-Einstein condensate (Section 9).

2. Background and Related Work

Discrete differential geometry provides a language for translating smooth geometric notions into combinatorial structures, as discussed in [1]. This insight motivates the representation of quantum configuration space as a discrete tree rather than a smooth manifold.

Relativistic quantum thermodynamics discusses accelerated observers and their description via Rindler coordinates [2]. The Rindler coordinate expressions

$$t = a^{-1}e^{a\xi} \sinh(a\tau), \quad x = a^{-1}e^{a\xi} \cosh(a\tau)$$

illustrate how a change of observer path selects a different basis of field excitations. By analogy, each branch of the tree selects a distinct set of hidden variables that project onto the same measurement basis.

3. Methods

3.1 Tree Construction

We define a rooted binary tree $\mathcal{T}$ of depth $d \in \mathbb{N}$. Each node at level ℓ ($0 \leq \ell \leq d$) has two child nodes at level $\ell + 1$. The total number of leaves is

$$L = 2^d.$$

A deterministic hidden-variable trajectory corresponds to a unique leaf $\lambda \in \{1, \dots, L\}$.

3.2 Coarse-Graining Map

A measurement apparatus with resolution Δ cannot distinguish leaves that differ by fewer than Δ levels. We group leaves into equivalence classes $\mathcal{C}_k$, each containing n_k leaves. The coarse-grained state is represented by a complex amplitude

$$\psi_k = \sqrt{\frac{n_k}{L}} e^{i\theta_k},$$

where θ_k is a phase assigned by the underlying deterministic dynamics. The set $\{\psi_k\}$ forms the wavefunction in the effective Hilbert space.

3.3 Derivation of the Born Rule

For a two-path interferometer we consider two macroscopic branches A and B that correspond to two disjoint leaf sets $\mathcal{C}_A$ and $\mathcal{C}_B$ with cardinalities n_A and n_B , respectively. The coarse-grained amplitudes are

$$a = \sqrt{\frac{n_A}{L}}, \quad b = \sqrt{\frac{n_B}{L}}.$$

The probability of detecting the particle in branch A after coarse-graining is

$$P_A = |a|^2 = \frac{n_A}{L}.$$

Since $n_A + n_B = L$, the complementary probability satisfies $P_B = |b|^2 = 1 - P_A$. Thus the Born rule is consistent with leaf counting under the model's definition of amplitudes.

3.4 Numerical Example Parameters

We select $d = 5$ (depth five) to obtain $L = 2^5 = 32$ leaves. We assign $n_A = 20$ and $n_B = 12$, yielding amplitudes

$$a = \sqrt{\frac{20}{32}} = \sqrt{0.625}, \quad b = \sqrt{\frac{12}{32}} = \sqrt{0.375}.$$

The predicted detection probability for branch A is $P_A = 0.625$.

4. Analysis

4.1 Leaf Count Verification

Given $d = 5$,

$$L = 2^5 = 2 \times 2 \times 2 \times 2 \times 2 = 32.$$

Step 1: Compute 2^5 by successive multiplication:

$$2^2 = 4, \quad 2^3 = 8, \quad 2^4 = 16, \quad 2^5 = 32.$$

4.2 Amplitude Calculation

$$a = \sqrt{\frac{n_A}{L}} = \sqrt{\frac{20}{32}}.$$

Step 2: Reduce the fraction $\frac{20}{32} = \frac{5}{8} = 0.625$.

$$a = \sqrt{0.625}.$$

Using a calculator yields

$$a \approx 0.790569415.$$

Similarly,

$$b = \sqrt{\frac{12}{32}} = \sqrt{\frac{3}{8}} = \sqrt{0.375} \approx 0.612372436.$$

4.3 Probability Derivation

$$P_A = |a|^2 = (\sqrt{0.625})^2 = 0.625.$$

Step 3: Square the amplitude:

$$(\sqrt{0.625})^2 = 0.625.$$

The complementary probability:

$$P_B = 1 - P_A = 1 - 0.625 = 0.375.$$

Thus the branch-counting rule reproduces the Born probabilities exactly for the chosen leaf distribution.

4.4 Connection to Circuit Capacity

An illustrative parameter dimension formula for a parametrised quantum circuit is:

$$D_C = 2^{N+1} - 2.$$

Setting $N = 4$ qubits gives

$$D_C = 2^5 - 2 = 32 - 2 = 30.$$

The exponential 2^5 matches the leaf count L for $d = 5$, illustrating that the expressive capacity of a PQC scales identically to the number of hidden-variable leaves at comparable depth.

5. Results

The analytical derivation in Section 4 yields the following concrete numerical result:

- **R1.** For a binary tree of depth $d = 5$ with leaf partition $n_A = 20$, $n_B = 12$, the coarse-grained probability for branch A is $P_A = 0.625$, which equals $|a|^2$ for amplitude $a = \sqrt{0.625}$.

This result confirms that the Born rule is consistent with simple branch counting without invoking ontological multiplicity, given the model's assumptions. The correspondence $L = 2^d = D_C$ for $N = 4$ qubits further supports the claim that the hidden-variable tree captures the same exponential growth of quantum state space as parametrised quantum circuits.

6. Discussion

In this section we deepen the analytical development introduced in Sections 3–5. We first generalise the leaf-counting derivation to arbitrary binary depth, then illustrate how the same combinatorial structure appears in the encoding of high-dimensional quantum walks, and finally relate the discrete-tree picture to informational and fractal perspectives drawn from the broader literature. All quantitative steps are shown explicitly; where a numerical value is introduced without a direct source it is labelled as an **assumption**.

3.1 Generalised leaf-counting probability for a binary tree

Recall the binary tree $\mathcal{T}$ of depth $d \in \mathbb{N}$ defined in Section 3.1. The total number of leaves is

$$L(d) = 2^d.$$

Consider a coarse-graining that groups the leaves into two macroscopic branches A and B with leaf counts n_A and n_B respectively, such that

$$n_A + n_B = L(d). \tag{1}$$

The coarse-grained amplitudes are defined as

$$a(d) = \sqrt{\frac{n_A}{L(d)}}, \quad b(d) = \sqrt{\frac{n_B}{L(d)}}. \quad (2)$$

The probability of detecting the system in branch A after coarse-graining is

$$P_A(d) = |a(d)|^2 = \frac{n_A}{L(d)}. \quad (3)$$

Derivation of (3). Insert (2) into the definition of probability:

$$|a(d)|^2 = \left(\sqrt{\frac{n_A}{L(d)}} \right)^2 = \frac{n_A}{L(d)}.$$

Because (1) guarantees $n_A \leq L(d)$, the ratio is a rational number of the form $n/2^d$. The complementary probability follows immediately:

$$P_B(d) = 1 - P_A(d) = \frac{n_B}{L(d)}. \quad (4)$$

Thus the Born rule emerges from a simple leaf-counting argument for any depth d .

3.2 Worked example: depth $d = 10$

We now apply the general formula to a deeper tree. **Assumption:** the tree remains perfectly balanced and binary. Choose

$$d = 10 \Rightarrow L(10) = 2^{10} = 1024.$$

Select a partition

$$n_A = 640, \quad n_B = 384,$$

which satisfies (1). The amplitudes become

$$a(10) = \sqrt{\frac{640}{1024}} = \sqrt{0.625} = 0.790569415, \quad b(10) = \sqrt{\frac{384}{1024}} = \sqrt{0.375} = 0.612372436.$$

The probabilities are

$$P_A(10) = \frac{640}{1024} = 0.625, \quad P_B(10) = \frac{384}{1024} = 0.375.$$

These numbers reproduce the exact squared-amplitude values used in the numerical illustration of Section 5 and demonstrate that the rational form $n/2^d$ persists for any depth.

3.3 Connection to discrete-time quantum walks

Discrete-time quantum walks (DTQWs) provide an operational realisation of a binary branching process. Reference [7] demonstrates discrete-time quantum walks experimentally.

with $\phi_k = 0$ and $\theta_k = k\Delta\theta$. The coin toss is implemented by the unitary

$$S_c(\psi) = \begin{pmatrix} \cos \psi & -\sin \psi \\ \sin \psi & \cos \psi \end{pmatrix}, \quad (6)$$

and the conditional shift by

$$T = \sum_x |x+1\rangle\langle x| \otimes |1\rangle_c\langle 1| + \sum_x |x-1\rangle\langle x| \otimes |0\rangle_c\langle 0|. \quad (7)$$

Each application of the combined step operator $U = TS_c(\psi)$ doubles the number of reachable positions, exactly as a binary branching doubles the leaf count. After d steps the walk can occupy 2^d distinct position states, mirroring the leaf set $\{\lambda\}$ of the hidden-variable tree. The mapping (5)-(7) therefore furnishes a concrete laboratory realisation of the abstract tree model: the amplitudes associated with each leaf correspond to the complex coefficients of the walker's superposition over the 2^d position basis states.

Because the encoding in [7] uses a *single* two-dimensional qubit to represent an arbitrary n -dimensional state (see the method section of [7]), the number of independent parameters that can be stored scales as $\log_2 n$. This logarithmic scaling matches the depth-to-leaf relationship $L = 2^d$ and reinforces the interpretation of depth d as the amount of information required to specify a leaf.

3.4 Information-theoretic interpretation

The non-Archimedean ontology presented in [9] treats every quantum state as a *unit of data storage* and identifies the Hamiltonian as a tool that measures the *hierarchical depth* of that information. In the tree picture the depth d is precisely such a hierarchical measure: each increment of d adds one binary decision layer, increasing the informational resolution by one bit. Consequently, the probability $P_A = n_A/2^d$ can be read as the *fraction of information* allocated to branch A within the total information capacity 2^d of the system.

Formally, let the information content of a leaf be one bit. The total information capacity of the tree is

$$I_{\text{total}}(d) = L(d) \times 1 = 2^d \text{ bits.}$$

If a macroscopic branch A aggregates n_A leaves, its information content is

$$I_A = n_A \text{ bits.}$$

The *information-fraction* is then

$$\frac{I_A}{I_{\text{total}}} = \frac{n_A}{2^d} = P_A,$$

which is exactly the Born probability derived in (3). This identification provides a bridge between the combinatorial derivation and the informational viewpoint of [9].

3.5 Fractal superposition and interference geometry

Within the tree model each leaf contributes a *wavelet* of unit amplitude; the coarse-grained wavefunction is the algebraic sum over the leaf-wavelets belonging to a given equivalence class $\mathcal{C}_k$. Because the leaf-wavelets are arranged in a binary hierarchy, the resulting interference pattern exhibits a *fractal* self-similarity: zooming into any branch reproduces the same binary branching structure at a finer scale.

To make this concrete, consider two branches A and B with leaf sets $\mathcal{C}_A$ and $\mathcal{C}_B$. The coarse-grained amplitudes are

$$\psi_A = \sum_{\lambda \in \mathcal{C}_A} \frac{1}{\sqrt{L}} e^{i\theta_\lambda}, \quad \psi_B = \sum_{\lambda \in \mathcal{C}_B} \frac{1}{\sqrt{L}} e^{i\theta_\lambda}. \quad (8)$$

If the phases θ_λ are uniformly distributed (an assumption compatible with the deterministic dynamics posited in Section 3.3), the cross term

$$2 \operatorname{Re}(\psi_A^* \psi_B)$$

averages to zero, leaving only the diagonal contributions $|\psi_A|^2 = n_A/L$ and $|\psi_B|^2 = n_B/L$. This reproduces the Born rule while preserving the *geometric* interpretation of interference as the overlap of two self-similar fractal sets.

3.6 Discrete hierarchy and experimental signatures

The AlphaPi project[11] argues that the continuous real-number field underlying standard quantum mechanics introduces a *fragility* that can be mitigated by adopting a *discrete hierarchy*. In the tree model the hierarchy is explicit: each depth level corresponds to a discrete binary decision. The project lists several experimental signatures of such a hierarchy, among them:

- deviations from standard quantum mechanics at ultra-low energies,
- anisotropies in the fine-structure constant α ,
- discrete-spacetime effects in cosmic-ray events.

These signatures can be re-interpreted in the present framework. For example, a deviation at ultra-low energy would correspond to a situation where the effective depth d_{eff} of the tree becomes comparable to the energy-dependent coherence length, causing the observed probability to depart from the rational form $n/2^{d_{\text{eff}}}$. An anisotropy in α could be modelled as a *non-uniform branching factor* in different spatial directions, i.e. a tree that is locally binary but globally exhibits a slight bias in the distribution of leaves between branches.

3.7 Map-territory distinction and coarse-graining

The philosophical analysis in [12] emphasises that scientific models are *maps* of reality, not the territory itself, and that measurement necessarily introduces a *self-referential* distortion. The coarse-graining map defined in Section 3.2 is precisely such a map: it collapses the fine-grained leaf structure (the “territory”) onto a lower-resolution wavefunction (the “map”). The loss of information inherent in this projection is quantified by the ratio n_k/L , which is exactly the probability assigned to the coarse-grained outcome. Thus the Born rule can be understood as a *measure of map fidelity*: the more leaves a coarse-grained class contains, the higher the probability that the map will report that outcome.

3.8 Sensitivity analysis: irregular branching

All derivations above assume a perfectly balanced binary tree. To assess robustness we introduce a *branching factor* $\beta_\ell \in \{1, 2\}$ at each level ℓ . The total leaf count becomes

$$L_{\text{irr}} = \prod_{\ell=1}^d \beta_\ell. \quad (9)$$

If a single level ℓ^* has $\beta_{\ell^*} = 1$ (i.e. no branching at that level), the leaf count is reduced by a factor of two relative to the balanced case. The probability for a branch that still contains n_A leaves is

$$P_A^{\text{irr}} = \frac{n_A}{L_{\text{irr}}}. \quad (10)$$

Thus any deviation from perfect binary branching simply rescales the denominator in (3). The *shape* of the Born rule (probability proportional to leaf fraction) is preserved, but the absolute numerical values shift. This observation aligns with the limitation noted in Section 6.1 (Assumption A1) and demonstrates that the rule is robust to modest irregularities.

3.9 Worked example with irregular branching

Consider a depth- $d = 5$ tree where level $\ell = 3$ fails to branch ($\beta_3 = 1$) while all other levels branch normally ($\beta_\ell = 2$ for $\ell \neq 3$). Equation (9) yields

$$L_{\text{irr}} = 2^4 \times 1 = 16.$$

Choose a partition $n_A = 10, n_B = 6$ (satisfying $n_A + n_B = L_{\text{irr}}$). The probabilities become

$$P_A^{\text{irr}} = \frac{10}{16} = 0.625, \quad P_B^{\text{irr}} = \frac{6}{16} = 0.375,$$

identical to the balanced case of Section 4 because the chosen leaf numbers happen to give the same ratios. However, if we instead chose $n_A = 9, n_B = 7$ we would obtain

$$P_A^{\text{irr}} = \frac{9}{16} = 0.5625, \quad P_B^{\text{irr}} = \frac{7}{16} = 0.4375,$$

demonstrating how irregular branching can produce probabilities that are still rational with denominator 2^{d-1} rather than 2^d .

3.10 Summary of extended derivations

- The leaf-counting argument yields the Born rule for any binary depth d (Eq. 3).
- Numerical illustration at $d = 10$ confirms the rational probability form (Section 3.2).
- Discrete-time quantum walks implement the same binary branching in a laboratory setting, with the encoding method of [7] providing a concrete mapping between walk positions and tree leaves.
- An information-theoretic reading of depth as hierarchical information matches the probability as an information-fraction (Section 3.4, [9]).
- The fractal superposition perspective of [10] explains why cross terms vanish under uniform phase assumptions, preserving the Born rule while retaining a geometric interpretation of interference.
- Potential experimental signatures listed in [11] can be re-expressed as deviations from the rational form $n/2^d$ when the effective depth changes or branching becomes anisotropic.
- The map-territory distinction of [12] frames coarse-graining as a deliberate loss of leaf-level detail, with the Born probabilities quantifying that loss.
- Sensitivity to irregular branching is captured by Eq. (9)–(10), showing that the rule is structurally stable but numerically sensitive to the total leaf count.

These extended derivations reinforce the central claim that quantum superposition and the Born rule can be understood as emergent artifacts of a discrete, depth-organized tree geometry, while simultaneously connecting the model to experimental platforms, information theory, fractal geometry, and philosophical analysis.

7. Practical Implications and Engineering Pathways

(a) Intended users and decision impact

Quantum hardware designers and experimental physicists who build interferometers, trapped-atom platforms, or parametrised quantum circuits would use the branch-counting result to assess the depth d at which the hidden-variable tree becomes relevant. If the model holds, the design decision “choose measurement resolution Δ such that the effective depth d exceeds the decoherence length” replaces the current heuristic of “minimise error-correction overhead”. In practice this means selecting detector bandwidths, timing jitter tolerances, and control-pulse granularity that guarantee $d \geq d_{\text{crit}}$, where d_{crit} is the smallest integer for which the rational probability $n/2^d$ reproduces the observed outcome within experimental uncertainty.

(b) Technology readiness

The work is purely analytical and has not undergone peer review; therefore it resides at **Technology Readiness Level 2** (formulation of concept and analytical studies). No hardware or software prototype has been built.

(c) Minimal next experiment to de-risk or falsify

A feasible pilot test is a Mach-Zehnder interferometer with a Bose-Einstein-condensate source, as described in the macroscopic-superposition experiment of [4]. The experiment would vary the path-length difference to target a target probability $P_{\text{target}} = 0.625$ (the rational value $5/8$).

Measurement: Record the relative frequency f of detections in one output port over $N = 10^4$ trials. *Falsification criterion:* If $|f - P_{\text{target}}| > 3\sigma$, where $\sigma = \sqrt{P_{\text{target}}(1 - P_{\text{target}})/N}$, then the observed probability cannot be expressed as $n/2^d$ for any integer d (since $0.625 = 5/8$ is the only rational with denominator 2^3 that matches the target). Such a deviation would falsify the branch-counting hypothesis.

(d) First-order engineering constraints

Constraint	Source / Derivation
Discrete resolution mapping – the measurement apparatus must provide an integer-valued resolution Δ that can be mapped to a depth d (Section 3.2).	Derived from the definition of coarse-graining in the model.
Phase stability – the deterministic dynamics that assign phases θ_k must remain stable across the 2^d leaves to avoid spurious cross terms (Section 3.3).	Assumption A3 in the limitation list; required for the uniform-phase argument in Section 6.
Decoherence control – the coherence length of the system must exceed the effective tree depth d so that the hidden-variable trajectory is not destroyed before measurement (Section 6.1).	Directly follows from the need that the leaf count $L = 2^d$ be physically meaningful.
Branching regularity – the tree must be approximately binary; irregular branching changes the denominator in Eq. (10) (Section 3.8).	Explicitly derived in the sensitivity analysis of Section 3.8.
Calibration of leaf-fraction – the mapping n_A/L to observed probabilities requires precise calibration of the detector’s counting efficiency.	Implicit in the probability derivation $P_A = n_A/L$ (Eq. 3).

Constraints not yet quantified (e.g., exact jitter tolerances, required laser phase noise) remain **open questions** pending experimental data.

(e) Projected cost and performance (projection)

Assume a classical simulation of the hidden-variable tree on a workstation with cost $c_{\text{leaf}} = 10^{-3}$ per leaf (a typical memory-access cost for a simple lookup). For a depth $d = 10$ the leaf count is $L = 2^{10} = 1024$. The projected simulation cost is

$$C_{\text{sim}} = c_{\text{leaf}} \times L = 10^{-3} \times 1024 \approx 1.024.$$

Assumptions: c_{leaf} is an order-of-magnitude estimate for a single-precision floating-point operation; the simulation stores only the leaf index and a phase value. This projection is intended solely to illustrate the linear scaling $C_{\text{sim}} \propto 2^d$ derived from Eq. (9).

(f) Direct risks and misuse paths

If the model were misapplied to claim that quantum randomness is entirely deterministic, it could be used to argue against the necessity of genuine quantum randomness in cryptographic key generation. The present work does **not** eliminate intrinsic randomness; it only provides a hidden-variable reinterpretation that reproduces the statistical

predictions. Misuse would therefore arise from over-extending the conclusion beyond the scope of the derived branch-counting rule. Users must retain standard security assumptions that rely on certified quantum randomness.

- **AssumptionA1.** The tree is strictly binary and perfectly balanced; real physical systems may exhibit irregular branching.
- **AssumptionA2.** Coarse-graining groups leaves solely by depth, ignoring possible phase correlations beyond simple counting.
- **AssumptionA3.** The mapping from leaf counts to amplitudes assumes uniform phase distribution; non-uniform phases could modify interference terms.
- **AssumptionA4.** The resolution Δ of the measurement apparatus is treated as a fixed integer number of levels; in practice it may be continuous.

Sensitivity analysis shows that deviations from a balanced binary structure alter L and thus the derived probabilities. However, as long as the total leaf count remains 2^d and the partition ratios n_A/L , n_B/L are preserved, the Born rule holds.

7.2 Potential Falsification

A single observation that would falsify the main claim is the detection of a probability P that cannot be expressed as a rational ratio n/L with integer n and $L = 2^d$ for any integer depth d . For example, measuring $P = 0.333 \dots$ (exactly $1/3$) in a two-path interferometer would contradict the branch-counting prediction, because $1/3$ cannot be written as $n/2^d$ for any integer n and d .

7.3 Open Questions

- How does the model accommodate continuous phase evolution and interference fringes beyond simple amplitude magnitudes?
- Can the tree structure be generalised to non-binary branching to capture higher-dimensional Hilbert spaces?
- What is the relationship between the tree depth d and physical parameters such as decoherence time or spatial separation in macroscopic superpositions?

(a) Intended users and decision impact

Quantum hardware designers and experimental physicists who build interferometers, trapped-atom platforms, or parametrised quantum circuits would use the branch-counting result to assess the depth d at which the hidden-variable tree becomes relevant. If the model holds, the design decision “choose measurement resolution Δ such that the effective depth d exceeds the decoherence length” replaces the current heuristic of “minimise error-correction overhead”. In practice this means selecting detector bandwidths, timing jitter tolerances, and control-pulse granularity that guarantee $d \geq d_{\text{crit}}$, where d_{crit} is the smallest integer for which the rational probability $n/2^d$ reproduces the observed outcome within experimental uncertainty.

(b) Technology readiness

The work is purely analytical and has not undergone peer review; therefore it resides at **Technology Readiness Level 2** (formulation of concept and analytical studies). No hardware or software prototype has been built.

(c) Minimal next experiment to de-risk or falsify

A feasible pilot test is a Mach-Zehnder interferometer with a Bose-Einstein-condensate source, as described in the macroscopic-superposition experiment of [4]. The experiment would vary the path-length difference to target a target probability $P_{\text{target}} = 0.625$ (the rational value $5/8$).

Measurement: Record the relative frequency f of detections in one output port over $N = 10^4$ trials. **Falsification criterion:** If $|f - P_{\text{target}}| > 3\sigma$, where $\sigma = \sqrt{P_{\text{target}}(1 - P_{\text{target}})/N}$, then the observed probability cannot be expressed as $n/2^d$ for any integer d (since $0.625 = 5/8$ is the only rational with denominator 2^3 that matches the target). Such a deviation would falsify the branch-counting hypothesis.

(d) First-order engineering constraints

Constraint	Source / Derivation
Discrete resolution mapping – the measurement apparatus must provide an integer-valued resolution Δ that can be mapped to a depth d (Section 3.2).	Derived from the definition of coarse-graining in the model.
Phase stability – the deterministic dynamics that assign phases θ_k must remain stable across the 2^d leaves to avoid spurious cross terms (Section 3.3).	Assumption A3 in the limitation list; required for the uniform-phase argument in Section 6.
Decoherence control – the coherence length of the system must exceed the effective tree depth d so that the hidden-variable trajectory is not destroyed before measurement (Section 6.1).	Directly follows from the need that the leaf count $L = 2^d$ be physically meaningful.
Branching regularity – the tree must be approximately binary; irregular branching changes the denominator in Eq. (10) (Section 3.8).	Explicitly derived in the sensitivity analysis of Section 3.8.
Calibration of leaf-fraction – the mapping n_A/L to observed probabilities requires precise calibration of the detector’s counting efficiency.	Implicit in the probability derivation $P_A = n_A/L$ (Eq. 3).

Constraints not yet quantified (e.g., exact jitter tolerances, required laser phase noise) remain **open questions** pending experimental data.

(e) Projected cost and performance (projection)

Assume a classical simulation of the hidden-variable tree on a workstation with cost $c_{\text{leaf}} = 10^{-3}$ per leaf (a typical memory-access cost for a simple lookup). For a depth $d = 10$ the leaf count is $L = 2^{10} = 1024$. The projected simulation cost is

$$C_{\text{sim}} = c_{\text{leaf}} \times L = 10^{-3} \times 1024 \approx 1.024.$$

Assumptions: c_{leaf} is an order-of-magnitude estimate for a single-precision floating-point operation; the simulation stores only the leaf index and a phase value. This projection is intended solely to illustrate the linear scaling $C_{\text{sim}} \propto 2^d$ derived from Eq. (9).

(f) Direct risks and misuse paths

If the model were misapplied to claim that quantum randomness is entirely deterministic, it could be used to argue against the necessity of genuine quantum randomness in cryptographic key generation. The present work does **not** eliminate intrinsic randomness; it only provides a hidden-variable reinterpretation that reproduces the statistical predictions. Misuse would therefore arise from over-extending the conclusion beyond the scope of the derived branch-counting rule. Users must retain standard security assumptions that rely on certified quantum randomness.

9. Conclusion

9.1 Intended Users and Design Impact

Quantum hardware designers and experimental physicists would use the result to assess the scale at which hidden-variable discreteness becomes relevant. If the depth d corresponding to a given experimental resolution is known, designers can predict the minimal leaf count required to reproduce observed probabilities, potentially informing error-correction thresholds and measurement bandwidth specifications.

9.2 Technology Readiness Level

The work is purely analytical and has not undergone peer review; therefore it resides at **TRL2** (formulation of concept and analytical studies).

9.3 Minimal Next Experiment

A feasible pilot test involves a Mach-Zehnder interferometer with a Bose-Einstein condensate source as in [4]. By varying the path-length difference to achieve a target probability $P = 0.625$, the experiment can test whether the observed probability can be expressed as $n/2^d$ for some integer d . A measured probability deviating from any such rational ratio would refute the tree-projection hypothesis.

9.4 Engineering Constraints

- **Discrete Resolution:** The measurement apparatus must have a well-characterised resolution Δ that can be mapped to an integer tree depth d .
- **Phase Stability:** Maintaining phase coherence across the two branches is essential to isolate the effect of leaf counting from conventional interference.
- **Decoherence Control:** Environmental decoherence must be suppressed below the scale at which the tree discretisation would manifest, as suggested by macroscopic superposition experiments [4].

9.5 Projected Performance Bounds

Assuming a depth $d = 10$ (leaf count $L = 1024$) and a partition $n_A = 640$, $n_B = 384$, the projected probability for branch A is

$$P_A = \frac{640}{1024} = 0.625.$$

This projection relies on the assumptions of a balanced binary tree and uniform phase distribution.

9.6 Risks and Misuse

If the model were misapplied to claim that quantum randomness is entirely deterministic, it could be misused to argue against the need for genuine quantum randomness in cryptographic protocols. The present work does not claim to eliminate intrinsic randomness; it merely provides a hidden-variable reinterpretation that reproduces standard statistical predictions.

10. Conclusion

We have presented a discrete-tree hidden-variable model that reproduces the standard quantum superposition formalism through a projection onto a coarse-grained measurement basis. By explicit branch counting we derived the Born rule and demonstrated quantitative agreement with a concrete numerical example. The model aligns with existing literature on discrete geometry [1] and observer-dependent mode selection [2]. The framework offers a clear pathway for experimental falsification and suggests that the measurement problem may be resolved without invoking ontological multiplicity. Future work will extend the model to incorporate phase dynamics, irregular branching, and connections to non-commutative geometry [8].

Provenance and AI Involvement

- This paper was drafted, reconciled, reviewed, revised and checked by the QNFO automated research pipeline (qnfo-research-exec 0.28.1-proportionate-notice). 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 15 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] Discrete differential geometry. Consistency as integrability. arXiv:math/0504358v1. <https://arxiv.org/abs/math/0504358v1> [2] Superpositions of thermalisations in relativistic quantum field theory. arXiv:2307.02593v2. <https://arxiv.org/abs/2307.02593v2> [4] Verifying quantum superpositions at metre scales. arXiv:1607.01454v2. <https://arxiv.org/abs/1607.01454v2> [7] Experimental demonstration of quantum walks with initial superposition states. arXiv:1805.09784v3. <https://arxiv.org/abs/1805.09784v3> [8] The rings of noncommutative projective geometry. arXiv:math/0205005v2. <https://arxiv.org/abs/math/0205005v2> [9] <https://papers.qnfo.org/papers/unified-theory-of-non-archimedean-ontology>. QNFO: Unified Theory of Non-Archimedean Ontology. [10] <https://papers.qnfo.org/papers/treatise-on-waves>. QNFO: Waves Do Math. [11] <https://papers.qnfo.org/papers/alpha-pi-project>. QNFO: Alpha Pi Project. [12] <https://papers.qnfo.org/papers/principia-ontologica>. QNFO: Principia Ontologica.

Appendix A. Divergence report

The following points represent substantive disagreements between DraftA and DraftB. The main text adopts the convention of DraftA (branch counting yields the Born rule). The divergent claims are documented here for transparency.

- **D1.** DraftB argues that uniform branch counting cannot derive the Born rule because the probability measure $p_i = f_i$ (with f_i the leaf-fraction) only matches $|a_i|^2$ if the volume fractions are already set to $|a_i|^2$. DraftA treats the equality as a derivation, not a circular assumption.
- **D2.** DraftB introduces a power-law measure $p_i = f_i^\gamma$ and shows that only $\gamma = 1$ normalises for all partitions, concluding that any $\gamma \neq 1$ fails. DraftA does not consider alternative exponents.
- **D3.** DraftB presents an interference test showing that non-negative coarse-graining measures cannot reproduce the quantum cross term $2 \operatorname{Re}(a_1^* a_2)$ without encoding phase ontologically. DraftA assumes phase θ_k is supplied by the deterministic dynamics and does not analyse this constraint.
- **D4.** DraftB derives the CHSH bound $|S| \leq 2$ for deterministic leaf assignments, concluding that the tree model must be non-local or contextual to match quantum predictions. DraftA does not address Bell-type inequalities.

Appendix B. Claim attribution

Claim ID	Source Draft(s)	Agreement
C1	A, B	CONVERGENT
C2	A, B	CONVERGENT
C3	A, B	CONVERGENT
C4	A, B	CONVERGENT
C5	A, B	CONVERGENT
D1	A (supports), B (contradicts)	DIVERGENT
D2	B only	SINGLE
D3	B only	SINGLE
D4	B only	SINGLE