

A Machine-Verified Bijective Proof of the Rooted Child-Factorial Catalan Identity over Spanning Trees of the Complete Graph

Lluis Eriksson

July 2, 2026

Abstract

Let K_{n+1} be the complete graph on the vertex set $\{0, 1, \dots, n\}$, and for a spanning tree T of K_{n+1} , rooted at 0, let $c_T(v)$ denote the number of children of the vertex v . We prove the exact identity

$$\sum_{T \in \text{ST}(K_{n+1})} \prod_{v=0}^n c_T(v)! = n! C_n,$$

where C_n is the n -th Catalan number; equivalently, the normalized sum $(n+1) \left((n+1)! \right)^{-1} \sum_T \prod_v c_T(v)!$ equals C_n exactly. The proof is bijective: pairs consisting of a spanning tree together with a linear ordering of every child set are placed in explicit bijection with vertex-labeled plane trees on $n+1$ nodes whose root carries the label 0. The identity arises as the exact “second-Ursell” normalization constant in the author’s audit-first programme on four-dimensional $SU(N)$ Yang–Mills existence and mass gap, where it had been isolated as an explicitly open, named proposition in a public repository; the present paper is self-contained combinatorics and makes no claim about that programme. The entire proof has been formalized in Lean 4 against a pinned Mathlib snapshot: the headline declarations compile with no `sorry`, and the kernel’s axiom oracle reports exactly `[propext, Classical.choice, Quot.sound]`. All artifacts, including a pinned continuous-integration replay of the full verification, are public.

1 Introduction

Cluster and tree-graph expansions bound logarithms of partition functions by sums over trees; this line of reasoning goes back to Penrose’s tree-graph identity [2] and was given its modern abstract-polymer form by Kotecký and Preiss [3]. In quantitative applications one repeatedly meets weighted sums over the spanning trees of a complete graph, where the weight of a tree is a product of local factors attached to its vertices. The simplest nontrivial member of this family assigns to a rooted spanning tree the product of the factorials of its child counts.

In the author’s programme on the construction of four-dimensional $SU(N)$ Yang–Mills theory [7, 8], the terminal Kotecký–Preiss estimate requires the *exact* evaluation of this factorially weighted spanning-tree sum: the “second-Ursell” normalization constant of the expansion is precisely the ratio of the weighted sum to $(n+1)!/(n+1) = n!$, and convergence at the stated radius needs this ratio to be the Catalan number C_n on the nose, not merely of Catalan order. The statement was therefore isolated, under the name `RootedChildFactorialCatalanIdentity`,

as an explicitly open, named proposition in a public repository, so that downstream theorems could assume it honestly until an unconditional, machine-checked proof against the repository's own definitions was supplied. This paper supplies that proof.

This paper does three things.

1. It states and proves the identity (Theorem 2.1) by an explicit bijection between pairs (*spanning tree, orderings of all child sets*) and vertex-labeled plane trees with root label 0 (Section 3), combined with the enumeration of the latter objects (Section 4).
2. It reports a complete formalization of the proof in Lean 4 [5] over the Mathlib library [6], carried out against the pinned toolchain of the repository and checked by the Lean kernel with a clean axiom footprint (Section 5).
3. It documents the reproducibility artifacts: the repository, the closing pull request, the exact commit, and a pinned continuous-integration replay of the entire verification (Section 6).

The mathematics below is entirely self-contained and elementary; no result from, and no claim about, the Yang–Mills programme is used or made.

2 Statement and numerical audit

Fix $n \geq 1$ and let $V = \{0, 1, \dots, n\}$, so that $|V| = n + 1$. We write $\text{ST}(K_{n+1})$ for the set of spanning trees of the complete graph on V , encoded as sets of unordered edges. Every spanning tree T is rooted at the distinguished vertex 0: each $w \neq 0$ has a unique neighbor lying on the path from w to 0, its *parent* $p_T(w)$, and

$$\text{ch}_T(v) = \{w \neq 0 : p_T(w) = v\}, \quad c_T(v) = |\text{ch}_T(v)|$$

are the *children* and the *child count* of v . (In the formalization the parent is defined for an arbitrary edge set as the minimal neighbor one breadth-first-search level closer to the root; for a tree this minimum is over a singleton and coincides with the graph-theoretic parent. All statements below are proved against that general definition.)

Let $C_n = \frac{1}{n+1} \binom{2n}{n}$ denote the Catalan numbers, so $C_1, C_2, C_3, \dots = 1, 2, 5, 14, 42, 132, \dots$

Theorem 2.1 (rooted child-factorial Catalan identity). *For every $n \geq 1$,*

$$\sum_{T \in \text{ST}(K_{n+1})} \prod_{v \in V} c_T(v)! = n! C_n.$$

Corollary 2.2 (normalized form). *For every $n \geq 1$,*

$$(n+1) \left((n+1)! \right)^{-1} \sum_{T \in \text{ST}(K_{n+1})} \prod_{v \in V} c_T(v)! = C_n$$

as an identity of real numbers. This is the form `RootedChildFactorialCatalanIdentity` of the repository.

Remark 2.3. The unweighted analogue of the left-hand side is Cayley's formula $|\text{ST}(K_{n+1})| = (n+1)^{n-1}$ [1]. Since every weight $\prod_v c_T(v)!$ is a positive integer, Theorem 2.1 in particular interpolates between the two classical sequences: for $n = 4$, for instance, the 125 spanning trees of K_5 carry total weight $336 = 4! \cdot 14$.

Before turning to the proof we record a brute-force audit. Enumerating all spanning trees of K_{n+1} through Prüfer sequences and computing the weighted sum directly gives the following table, in exact agreement with Theorem 2.1.

n	$ \text{ST}(K_{n+1}) = (n+1)^{n-1}$	$\sum_T \prod_v c_T(v)!$	$n! C_n$
1	1	1	1
2	3	4	4
3	16	30	30
4	125	336	336
5	1296	5040	5040
6	16807	95040	95040

3 The bijection

The left-hand side of Theorem 2.1 counts, with multiplicity, the pairs

$$(T, L), \quad T \in \text{ST}(K_{n+1}), \quad L = (\ell_v)_{v \in V},$$

where ℓ_v is a duplicate-free list enumerating the child set $\text{ch}_T(v)$; for fixed T there are exactly $\prod_v c_T(v)!$ such *child listings*. We call the pairs (T, L) *ordered rooted spanning trees*.

The right-hand side counts *labeled plane trees*. A plane tree (rose tree) with labels in V is either a node carrying a label $a \in V$ together with a finite *ordered* list of labeled plane subtrees. Write $\text{lab}(u)$ for the list of labels of u in root-first order, and let

$$\mathcal{L}_n = \{u : \text{lab}(u) \text{ is a permutation of } (0, 1, \dots, n) \text{ and the root label of } u \text{ is } 0\}.$$

Section 4 shows $|\mathcal{L}_n| = n! C_n$. Theorem 2.1 therefore follows from:

Theorem 3.1 (bijection). *The map*

$$\Phi : (T, L) \longmapsto \Phi_{T,L}(0),$$

where $\Phi_{T,L}(v)$ is the labeled plane tree with root label v whose ordered subtrees are $\Phi_{T,L}(w)$ for w running through ℓ_v in list order, is a bijection from the set of ordered rooted spanning trees onto $\mathcal{L}_n$.

The recursion defining $\Phi_{T,L}$ is well founded because a child sits one breadth-first-search level below its parent, so the quantity $(n+1) - \text{level}_T(v)$ strictly decreases along the recursion. The proof of Theorem 3.1 splits into the three usual claims, whose formalized statements are quoted in Section 5.

3.1 Well-definedness

Proposition 3.2. *For every ordered rooted spanning tree (T, L) one has $\Phi_{T,L}(0) \in \mathcal{L}_n$.*

Proof sketch. The label multiset of the subtree $\Phi_{T,L}(v)$ is characterized by parent iterates:

$$w \in \text{lab}(\Phi_{T,L}(v)) \iff \exists k \geq 0 : p_T^k(w) = v \text{ and } \text{level}_T(w) = \text{level}_T(v) + k.$$

(The proof is by the same well-founded recursion as the definition of Φ .) Two consequences follow. First, subtrees rooted at distinct children of the same node have disjoint label sets: a common label w would have the two children as parent iterates at the *same* offset $k = \text{level}(w) - \text{level}(\text{child})$, hence the children would coincide. This makes $\text{lab}(\Phi_{T,L}(v))$ duplicate-free, by induction. Second, at $v = 0$ every vertex w occurs, with $k = \text{level}_T(w)$, because iterating the parent map from w reaches the root in exactly $\text{level}_T(w)$ steps. A duplicate-free list of $n + 1$ elements of V is a permutation of $(0, \dots, n)$, and the root label of $\Phi_{T,L}(0)$ is 0 by construction. $\square$

3.2 Injectivity

For a labeled plane tree u with duplicate-free labels and $v \in \text{lab}(u)$, let $\text{childrenOf}(u, v)$ denote the list of root labels of the ordered subtrees of the unique node of u labeled v .

Proposition 3.3. *Φ is injective. More precisely, if $u = \Phi_{T,L}(0)$ then $\ell_v = \text{childrenOf}(u, v)$ for every v , and T is recovered as the set of parent edges $\{\{v, w\} : w \in \ell_v\}$.*

Proof sketch. The extraction $\ell_v = \text{childrenOf}(u, v)$ is proved by the same well-founded recursion, using the pairwise disjointness of sibling label sets to show that the search for the node labeled v descends into the correct subtree. Thus L is determined by u . Next, the listings determine the parent map ($w \in \ell_v$ forces $p_T(w) = v$), and a spanning tree is determined by its parent map: its edge set *equals* its set of parent edges. The latter recovery statement is a lemma of the upstream artifact (in the terminology of the repository, $\text{penroseTree}(T) = T$ for every spanning tree T), and injectivity follows. $\square$

3.3 Surjectivity

Proposition 3.4. *Φ is surjective onto $\mathcal{L}_n$.*

Proof sketch. Given $u \in \mathcal{L}_n$, define

$$T_u = \{\{v, w\} : w \in \text{childrenOf}(u, v)\}, \quad L_u = (\text{childrenOf}(u, v))_{v \in V}.$$

Three facts are needed.

(a) *T_u is a spanning tree.* Connectivity: by structural induction on u , every label reaches the root label along parent–child edges. Edge count: the edges of T_u are in bijection with the disjoint union $\bigsqcup_v \text{childrenOf}(u, v)$ (distinct parents contribute distinct edges, by antisymmetry of the structural parent–child relation on a duplicate-free tree), and that disjoint union is in bijection with the set of non-root labels (every non-root label has a structural parent, and the parent is unique). Hence $|T_u| = n$, and a connected graph on $n + 1$ vertices with n edges is a tree.

(b) *The rooted children of T_u are the structural children:* $\text{ch}_{T_u}(v) = \text{childrenOf}(u, v)$ as sets. The key claim is that for every $w \neq 0$,

$$w \in \text{childrenOf}(u, p_{T_u}(w)),$$

proved by strong induction on the breadth-first-search level of w in T_u . Indeed $p_{T_u}(w)$ is adjacent to w , and adjacency in T_u means exactly “structural parent or structural child”. If $b := p_{T_u}(w)$ were a structural *child* of w , then $b \neq 0$ (the root label 0 is never a child), the inductive hypothesis applies to b (its level is smaller by the parent equation), and by uniqueness of structural parents $p_{T_u}(b) = w$. The two level equations $\text{level}(b) + 1 = \text{level}(w)$ and $\text{level}(w) + 1 = \text{level}(b)$ then

contradict each other. Notably, this argument needs neither a depth function on u nor an explicit parent function: the two instances of the parent specification collide by themselves.

(c) *Reconstruction:* $\Phi_{T_u, L_u}(0) = u$. This is an instance of the general lemma that $\Phi_{T, L}(\text{root}(u)) = u$ whenever the listings L agree with $\text{childrenOf}(u, \cdot)$ on the labels of u — proved by structural induction on u , the agreement hypothesis propagating to each subtree because childrenOf restricted to the labels of a subtree coincides with the subtree’s own childrenOf . For (T_u, L_u) the agreement holds by definition (L_u is $\text{childrenOf}(u, \cdot)$), and $\text{root}(u) = 0$. $\square$

4 Counting labeled plane trees

Proposition 4.1. $|\mathcal{L}_n| = n! C_n$ for every $n \geq 1$.

Proof sketch. Unlabeled plane trees with $m + 1$ nodes are counted by C_m : deleting the root’s first subtree decomposes a plane tree with $m + 1$ nodes into an ordered pair of strictly smaller plane trees, which yields the recurrence $P_{m+1} = \sum_{i+j=m} P_i P_j$ with $P_0 = 1$ (nodes counted so that P_m is the number of plane trees with $m + 1$ nodes), i.e. exactly the defining recurrence of the Catalan numbers, as implemented in Mathlib’s `Combinatorics.Enumerative.Catalan`. A labeled plane tree in $\mathcal{L}_n$ is equivalently a pair (shape, labeling): writing the labels in root-first order identifies the labelings of a fixed shape with the permutations of $(0, \dots, n)$ whose first entry is 0, of which there are $n!$, and the pair (shape, label word) determines the labeled tree uniquely (rigidity of the decoration). Hence $|\mathcal{L}_n| = n! \cdot C_n$. $\square$

Proof of Theorem 2.1 and Corollary 2.2. By Theorem 3.1 and Proposition 4.1,

$$\sum_T \prod_v c_T(v)! = \#\{(T, L)\} = |\mathcal{L}_n| = n! C_n.$$

For the corollary, multiply by $(n+1)((n+1)!)^{-1} = (n!)^{-1}$ and cancel: $(n+1)((n+1)!)^{-1} n! C_n = C_n$, using $(n+1)! = (n+1)n!$ and $n! \neq 0$. $\square$

5 The Lean formalization

Every statement above has been formalized in Lean 4 [5] over Mathlib [6], against the repository’s pinned toolchain (Lean `leanprover/lean4:v4.29.0-rc6`, Mathlib commit `07642720480157414db592fa85b626dafb71355b`) and against the repository’s *own* definitions of spanning trees, breadth-first-search parents, and rooted child counts. The headline declarations are:

```
theorem sum_prod_rootedChildCount_factorial_eq (n : ℕ) :
  (∑ T ∈ spanningTrees (T : SimpleGraph (Fin (n + 1))),
    ∏ v, (rootedChildCount T v)!) = n ! * catalan n

theorem rootedChildFactorialCatalanIdentity_holds (n : ℕ) :
  RootedChildFactorialCatalanIdentity n
```

where the second statement unfolds, per the repository’s definitions, to the real-number identity of Corollary 2.2:

```

noncomputable def rootedChildFactorialTreeSum (n : ℕ) : ℝ :=
  ∑ T ∈ spanningTrees (T : SimpleGraph (Fin (n + 1))),
    ∏ v : Fin (n + 1), ((rootedChildCount T v).factorial : ℝ)

noncomputable def rootedChildFactorialTreeSumNormalized (n : ℕ) : ℝ :=
  ((n : ℝ) + 1) * (((n + 1).factorial : ℝ))-1 *
    rootedChildFactorialTreeSum n

```

The proof development comprises five modules (about 2,200 lines in total):

Module	Lines	Role
RootedCatalanWord	234	word-side algebra; normalized-form arithmetic
PlaneTreeCatalan	214	plane trees with $m + 1$ nodes are counted by C_m
LabeledPlaneTreeCatalan	382	decoration rigidity; $ \mathcal{L}_n = n! C_n$
RootedCatalanBridge	1288	the bijection: Props. 3.2–3.4, Thm. 3.1, Thm. 2.1
RootedCatalanIdentityProof	61	Corollary 2.2 against the repository’s definitions

A few engineering points may be of independent interest to formalizers.

- *Well-founded tree building.* The map Φ is defined by well-founded recursion on $(n + 1) - \text{level}_T(v)$; the label characterization of Proposition 3.2 and the extraction of Proposition 3.3 are proved by the same measure.
- *Nested induction on rose trees.* Labeled plane trees are a nested inductive type (a node carries a *list* of subtrees). All structural lemmas — uniqueness and existence of structural parents, antisymmetry of the parent–child relation, duplicate-freeness of child listings, reachability of the root — are proved by recursion on `sizeOf`, descending into list members via `List.sizeOf_lt_of_mem`.
- *The level-clash argument.* Step (b) of Proposition 3.4 is formalized exactly as sketched: a strong induction on the breadth-first-search level in which the “bad” branch is refuted by instantiating the parent specification twice and letting the two resulting level equations contradict each other by linear arithmetic. No depth function on the plane tree is ever introduced.
- *Counting by double disjoint union.* The edge count $|T_u| = n$ is two applications of `Finset.card_biUnion`: edges $\leftrightarrow$ child slots $\leftrightarrow$ non-root labels, with the required disjointness supplied by antisymmetry and by uniqueness of parents, respectively.
- *Axiom footprint.* The development contains no `sorry` and no additional axioms; for every exported declaration, including both headline theorems, `#print axioms` reports exactly

[`propext`, `Classical.choice`, `Quot.sound`]

i.e. the standard classical axioms of Mathlib.

6 Reproducibility and artifacts

All artifacts are public in the repository

<https://github.com/lluiseriksson/rooted-tree-catalan-closure>

The closure is submitted as pull request #5 (branch `codex/catalan-rigidity-lean-modules`, head commit `1eff8fc651ed3338800b8ff67471e6091ce8ef33`), which adds the five proof modules under `lean-patch/YangMills/KP/`, updates the theorem manifest and claim-boundary documents, and archives the verification logs.

The verification is replayed from scratch by a pinned GitHub Actions workflow; the run

[https:](https://github.com/lluiseriksson/rooted-tree-catalan-closure/actions/runs/28610948452)

[//github.com/lluiseriksson/rooted-tree-catalan-closure/actions/runs/28610948452](https://github.com/lluiseriksson/rooted-tree-catalan-closure/actions/runs/28610948452)

(i) checks out the pinned upstream artifact and Mathlib snapshot, (ii) applies the documented import-narrowing patch — which leaves every upstream proof body byte-identical and only trims module import blocks so the relevant dependency chain builds in isolation — (iii) compiles the upstream adapter chain and the five new modules, (iv) runs the axiom-oracle checks and prints `#print axioms` for the closing declarations, and (v) validates the resulting logs against the archived evidence with the repository’s replay validators. The archived logs are `catalan-bridge-modules-replay.log`, `catalan-bridge-full-compile.log`, and `catalan-identity-proof-oracle.log` under `lean-patch/verification/`.

Independently of Lean, the brute-force audit of Section 2 (a 30-line Python script enumerating spanning trees via Prüfer sequences) confirms the identity for $1 \leq n \leq 6$.

7 Related work and outlook

The identity sits at the intersection of two classical themes. On the enumerative side, plane trees and their labelings belong to the standard Catalan circle of ideas [4]; the specific weighted sum over spanning trees of K_{n+1} , with weight $\prod_v c_T(v)!$, appears to be less standard, and we are not aware of a prior published bijective treatment, although the result itself is elementary once the correct bijection is identified. On the analytic side, factorially weighted tree sums are the combinatorial core of tree-graph and cluster expansions [2, 3]; within the author’s programme, Theorem 2.1 supplies the exact constant needed by the terminal Kotecký–Preiss bound [8], and the machine-checked closure documented here continues the programme’s mechanical-audit methodology [9].

Two natural directions remain. First, the same bijection should evaluate refinements of the sum, e.g. with weights $\prod_v x_{c_T(v)}$ for formal variables x_k , connecting to the exponential-generating-function calculus of rooted labeled trees. Second, on the formal side, the bridge module develops a small reusable theory of labeled rose trees (structural parents, listings, extraction, reconstruction) that may be worth contributing to Mathlib in generic form.

Acknowledgments

The Lean formalization was developed interactively with Anthropic’s Claude; every theorem stated here was checked by the Lean 4 kernel against the pinned Mathlib snapshot, with the axiom oracle output archived in the repository.

References

- [1] A. Cayley, *A theorem on trees*, Quart. J. Pure Appl. Math. **23** (1889), 376–378.
- [2] O. Penrose, *Convergence of fugacity expansions for classical systems*, in: T. A. Bak (ed.), *Statistical Mechanics: Foundations and Applications*, Benjamin, New York, 1967, 101–109.
- [3] R. Kotecký and D. Preiss, *Cluster expansion for abstract polymer models*, Comm. Math. Phys. **103** (1986), 491–498.
- [4] R. P. Stanley, *Catalan Numbers*, Cambridge University Press, 2015.
- [5] L. de Moura and S. Ullrich, *The Lean 4 theorem prover and programming language*, in: Automated Deduction – CADE 28, Lecture Notes in Computer Science 12699, Springer, 2021, 625–635.
- [6] The mathlib Community, *The Lean mathematical library*, in: Proceedings of the 9th ACM SIGPLAN International Conference on Certified Programs and Proofs (CPP 2020), ACM, 2020, 367–381.
- [7] L. Eriksson, *The Master Map: An Audit-First, Attack-Resistant Navigation Guide to the Unconditional Solution of the 4D $SU(N)$ Yang–Mills Existence and Mass Gap Problem*, ai.viXra:2602.0096, 2026.
- [8] L. Eriksson, *Closing the Last Gap in the 4D $SU(N)$ Yang–Mills Construction: A Verified Terminal KP Bound and an Explicit Clay Checklist*, ai.viXra:2602.0091, 2026.
- [9] L. Eriksson, *Mechanical Audit Experiments and Reproducibility Appendix for a Companion-Paper Programme on 4D $SU(N)$ Yang–Mills Existence and Mass Gap*, ai.viXra:2602.0117, 2026.
- [10] L. Eriksson, *rooted-tree-catalan-closure* (artifact repository), <https://github.com/lluiseiriksson/rooted-tree-catalan-closure>, pull request #5, commit 1eff8fc6, CI run 28610948452, 2026.