

From a Failed Corner Slide to a Verified Tableau Injection

A constructive positivity proof in height two,
with exact obstructions and a Lean-checked coloring formula

Yamba AI Math Lab
Research note prepared for Christopher Yamba

23 October 2026

Abstract

An earlier research conversation proposed a universal sign-reversing tableau involution for elementary coefficients of chromatic quasisymmetric functions. The proposed corner slide fails on five vertices, even for a connected incomparability graph of a height-two poset. We give its exact validity criterion and show an infinite cardinality obstruction to its intended unrestricted shape-to-shape injection. Allowing the moved entry to pass an initial block of comparable entries repairs validity; the absence of a three-element chain makes this repair injective. This construction is the existing injection of Cho–Huh, which we independently recovered and explicitly attribute. We give a complete mathematical proof of the height-two positivity corollary, reproducible finite calculations, and a precise account of which combinatorial statements are checked by Lean. The certificate now includes complete finite tableau families, a survivor-count identity in every inversion degree, a signed-array involution, and finite alternant algebra. The all-alphabet coloring-to-survivor equality is now checked in Lean, using actual proper colorings and the original complete survivor families. We claim a formalization of the established height-two result, not a novel injection or a solution for arbitrary height.

Claim status. The original universal slide is **refuted**. The local repair and height-two injection are **proved** below. Height-two positivity is an **established result reproduced here**. Finite computations are **empirical, exact, and bounded**. H3 is **proved and Lean-checked for height two**, for every finite alphabet, including zero. No arbitrary-height solution is claimed.

1 The problem and its actual frontier

For a labeled graph G on $[n]$, a proper coloring is a map $\kappa : [n] \rightarrow \mathbb{N}_{>0}$ assigning distinct colors to adjacent vertices. Set

$$X_G(\mathbf{x}; q) = \sum_{\kappa \text{ proper}} q^{\text{asc}_G(\kappa)} \prod_{i=1}^n x_{\kappa(i)}, \quad \text{asc}_G(\kappa) = \#\{ij \in E : i < j, \kappa(i) < \kappa(j)\}.$$

For a natural unit interval graph this is symmetric. Its elementary expansion is

$$X_G = \sum_{\lambda \vdash n} c_\lambda(q) e_\lambda,$$

where $e_\lambda = \prod_i e_{\lambda_i}$. The constructive goal is to interpret every $c_\lambda(q)$ as a sum of nonnegative monomials, ideally through a uniform, explicitly invertible map.

Established context; theorem statements inspected, not full proof audits

Hikita proved ordinary Stanley–Stembridge positivity; his result gives nonnegativity at positive real values of q [3]. Such numerical nonnegativity does not by itself imply nonnegative polynomial coefficients: $q^2 - q + 1$ is positive for every real q but has a negative coefficient. Sagan–Tom supply substantial tableau and coefficient results [4]. A September 2026 preprint explicitly claims full coefficientwise positivity [5]; we inspected its statement and parts of its argument, without a full proof audit or a local formal rebuild. Thus we do not infer the current global status from an older claim that the problem is open.

The supplied notes rank Problem 06 highest and include supposed completed proofs. Those claims require checking. Our first question is precise: *Does the stated corner map define an injection on the actual signed tableaux?* A negative answer is decisive for that proposed proof, even when the desired positivity statement remains true.

2 Objects, conventions, and the signed coefficient

Let P be a finite strict poset. An English P -tableau is a bijective filling of a partition diagram by its elements, with strictly increasing chains along rows and no P -descent down a column. If u is above v in adjacent cells, the column condition is $v \not\prec_P u$.

For $G = \text{inc}(P)$, set

$$\text{inv}_G(T) = \#\{ij \in E(G) : i < j, j \text{ lies in an earlier row than } i\}.$$

No graph edge joins elements within a chain row. Consequently this statistic equals the number of graph edges whose larger endpoint precedes the smaller endpoint in the row-major word $w(T)$.

For computations we use a weakly increasing Hessenberg vector $h = (h_1, \dots, h_n)$ with $i \leq h_i \leq n$. The associated order and graph are

$$i <_P j \iff h_i < j, \quad ij \in E(G_h) \iff i < j \leq h_i \quad (i < j).$$

The Schur-tableau expansion and dual Jacobi–Trudi identity give

$$X_{G_h}(\mathbf{x}; q) = \sum_{\mu \vdash n} F_\mu(q) s_\mu(\mathbf{x}), \quad F_\mu(q) = \sum_{T \in \text{PT}_\mu(P_h)} q^{\text{inv}_{G_h}(T)}, \quad (1)$$

$$s_\mu = \det(e_{\mu'_i - i + j})_{1 \leq i, j \leq \mu_1}. \quad (2)$$

We use $e_0 = 1$ and $e_r = 0$ for $r < 0$. The imported expansion (1) is due to Shareshian–Wachs [2]. It is an explicit mathematical dependency of the positivity corollary; it is not re-proved by our Lean code.

3 Why the original slide fails

Consider shape $(2^{k+1}, 1^m)$, with final two-cell row (a, b) , followed by singleton entries $c_1, \dots, c_m$. The proposed move splits (a, b) into successive singleton rows, leaving all following entries in order.

Theorem 3.1 (Exact boundary criterion). *For any finite poset and valid source tableau, this slide is valid if and only if $m = 0$ or $c_1 \not\prec_P b$. It preserves the row-major word and graph inversions whenever its target is valid.*

Proof. The only newly introduced column comparisons are a above b and, when $m > 0$, b above c_1 . The first is valid by $a <_P b$ and asymmetry. The second is precisely $c_1 \not\prec_P b$. All other conditions are inherited, and every label still appears once. The row-major words coincide. The word interpretation of inversions gives the last assertion. $\square$

There is a stronger obstruction than the failure of one algorithm.




Figure 1: A five-vertex failure with $h = (3, 4, 4, 5, 5)$. The graph has edges 12, 13, 23, 24, 34, 45 and is connected. The poset has height two. The inversion degree is 2; unchanged weight does not repair invalidity.

Theorem 3.2 (No unrestricted two-sector injection). *For every $n \geq 5$, in the total order on $[n]$, there are more tableaux of shape $(2, 2, 1^{n-4})$ than of shape $(2, 1^{n-2})$. No injection from the entire former set into the latter exists.*

Proof. These are standard Young tableaux. Transposing reduces the counts to shapes $(n-2, 2)$ and $(n-1, 1)$. For the latter, the sole entry in the second row can be any of $2, \dots, n$, giving $n-1$ choices. For the former, choose two second-row entries from $\{2, \dots, n\}$; every pair is valid except $\{2, 3\}$. Indeed the top-left entry is 1, and the next first-row entry is smaller than the larger chosen entry unless the pair is $\{2, 3\}$. The count is $\binom{n-1}{2} - 1 = n(n-3)/2$. At $n=5$ this exceeds $n-1$ by one; the difference increases by $n-2$ when n increases to $n+1$. It is therefore positive for every $n \geq 5$. $\square$

At five vertices the missing determinant sectors are visible directly:

$$[e_{4,1}]X_G = F_{2111} - F_{221} - F_{311} + F_{32} + F_{41} - 2F_5. \quad (3)$$

For a total order every F_μ is the number of standard tableaux of shape μ , and the six terms in (3) are $4 - 5 - 6 + 5 + 4 - 2 = 0$. The graph is edgeless, so $X_G = e_1^5$. The incomplete two-sector expression would instead give -1 . This is a proof obstruction, not a counterexample to elementary positivity.

F_{2111} 4	F_{221} -5	F_{311} -6	F_{32} 5	F_{41} 4	F_5 -2
-----------------	-----------------	-----------------	---------------	---------------	-------------

The whole signed coefficient is zero. Every sector matters.

Figure 2: The complete dual Jacobi–Trudi cancellation for $e_{4,1}$ in the five-vertex total order. A wider-shape term cannot be silently dropped.

The earlier top-only arm–leg rule also fails its positivity requirement. For $h = (2, 2, 3)$, the valid hook with first row $(1, 3)$ and leg (2) has negative sign. The proposed move of 3 above 2 is invalid since $2 <_P 3$; the reverse move needs $1 <_P 2$, which is false. It is a negative fixed point. The established hook involution permits moves at other leg positions [4].

4 The correct insertion and its inverse

Known construction, independently recovered; attribution: Cho–Huh, Section 3.1

The needed change is to move b farther down when the boundary is blocked. The construction below is the injection ψ_ℓ of Cho–Huh [1]; it is reproduced, not proposed as a new discovery.

Definition 4.1 (Comparable-prefix insertion). Write the singleton suffix as CD , where $C = (c_1, \dots, c_r)$ is its longest initial block with $c_i <_P b$ for every i . Replace

$$\text{front}, (a, b), C, D \mapsto \text{front}, (a), C, (b), D.$$

Call this map B . Entries of C, D are singleton rows. The front consists of the preceding two-cell rows.




Figure 3: The repaired map on the same connected example. The exchanged labels 3, 5 are comparable, so they are not an incomparability edge and no graph inversion changes.

Theorem 4.2 (Validity and weight). *For every finite poset, B sends valid tableaux of shape $(2^{k+1}, 1^m)$ to valid tableaux of shape $(2^k, 1^{m+2})$. It preserves inv_G for the incomparability graph.*

Proof. After removing b , the original first column is unchanged. Insert b after C . The entry immediately above it is either a or the last entry of C ; both are smaller than b in P . If a following entry exists, it is not smaller than b by maximality of C . The new adjacent column comparisons are therefore valid. Other comparisons are inherited, and rows are chains. The filling remains bijective. The row relation of b with a changes from the same row to a later row, and b passes the entries of C . Each such pair is comparable, hence is not an edge of G . All edge orders are unchanged. $\square$

Theorem 4.3 (Height-two injection). *If P contains no strict chain of three elements, then B is injective. Its inverse on its image selects the first singleton entry x below a such that $a <_P x$, removes it, and restores (a, x) as the final two-cell row.*

Proof. The moved entry b satisfies $a <_P b$. Every singleton crossed before it satisfies $c_i <_P b$. If $a <_P c_i$ held, then $a <_P c_i <_P b$ would be a three-element chain. Thus no earlier singleton is eligible, and b is exactly the first eligible entry. Extracting it recovers b and the entire original suffix in order, as well as the unchanged front and a . This is a left inverse, proving injectivity. $\square$

The hypothesis is essential for this inverse. In the five-element total order, $(12, 34, 5)$ and $(12, 35, 4)$ both map to $(12, 3, 4, 5)$. The map remains valid and weight preserving, but is no longer injective. Changing insertion depth alone does not solve the unrestricted problem.

5 A complete mathematical positivity corollary in height two

Suppose P_h has height at most two. Every tableau row then has length at most two. Write

$$\mu_j = (2^j, 1^{n-2j}), \quad F_j(q) = F_{\mu_j}(q), \quad 0 \leq j \leq \lfloor n/2 \rfloor,$$

Outside that range we define $\text{PT}_{\mu_j}(P_h) = \emptyset$ and $F_j = 0$, without assigning a partition to an expression with a negative exponent. For $1 \leq j \leq \lfloor n/2 \rfloor$, the two-column determinant is

$$s_{\mu_j} = e_{n-j}e_j - e_{n-j+1}e_{j-1}, \quad s_{\mu_0} = e_n.$$

Collecting terms in (1) gives, for $0 \leq k \leq \lfloor n/2 \rfloor$,

$$c_{(n-k,k)}(q) = F_k(q) - F_{k+1}(q), \quad (4)$$

where a zero part is omitted. If $n = 2k$ the second term is zero. Coefficients indexed by partitions with more than two parts vanish.

Theorem 5.1 (Height-two positivity; Cho–Huh). *For every natural unit interval order P_h of height at most two, all elementary coefficients of $X_{G_h}(\mathbf{x}; q)$ belong to $\mathbb{N}[q]$. More precisely,*

$$c_{(n-k,k)}(q) = \sum_{U \in \text{PT}_{\mu_k}(P_h) \setminus B(\text{PT}_{\mu_{k+1}}(P_h))} q^{\text{inv}_{G_h}(U)}.$$

Proof. For μ_{k+1} the map B is precisely the valid injection of Theorems 4.2–4.3, with k two-cell front rows. It preserves every inversion degree. Subtract its image from the positive set in (4); all remaining monomials have coefficient $+1$. When the source is empty the assertion is immediate. $\square$

Equivalently, on the disjoint union of a negative copy of the source and a positive copy of the target, pair each source with its image. Fix the other positive objects. This is a sign-reversing, weight-preserving involution with positive fixed points. Both the signed sets and the injection are necessary; an isolated flag-toggling lemma does not establish their existence.

6 Lean: what was actually checked

The project uses Lean 4.34.0 and `Std`, without a Mathlib download. Its actual tableau predicate requires a distinct vertex list, a permutation of that list as the flattened filling, nonempty chain rows, a partition shape, and every adjacent column condition. This prevents the semantic gap in the earlier boundary object, which omitted the comparison below the moved entry.

6.1 Signed arrays: a checked cancellation map

The new signed-array certificate permits empty rows. Fix a weakly decreasing function $\lambda : \mathbb{N} \rightarrow \mathbb{N}$, an array with m rows, and permutation labels $\pi(0), \dots, \pi(m-1)$. Its row lengths satisfy the integer equations


$$|R_i| = \lambda(\pi(i)) - \pi(i) + i \quad (0 \leq i < m).$$

The subtraction takes place in $\mathbb{Z}$: a negative requested length admits no row. Each row is a strict P -chain, and the concatenated rows bijectively fill the given distinct vertices. These are the actual predicates `ChainArray` and `SignedArray`, rather than coefficient identities assumed as inputs.

Theorem 6.1 (Formal signed-array involution). *For a finite poset of height at most two, the bottommost-defect map on these signed arrays is an involution preserving the vertex filling, shifted shape, chain rows, and tableau graph-inversion degree. Every moved array reverses $(-1)^{\text{inv}(\pi)}$. Its fixed arrays are exactly actual tableaux followed by empty rows. Their labels are the identity permutation, their row lengths are $\lambda(0), \dots, \lambda(m-1)$, and their signs are $+1$.*

Proof. At the bottommost failing first-column boundary, apply the local toggle in Figure 4, swapping the two row labels. The lower first entry stays b , preserving the good suffix and the pivot on the return move. The shifted keys $\lambda(j) - j$ are distinct; this excludes adjacent length increases of one, including $(0, 1)$ and $(1, 2)$. Height two bounds every chain row by two cells and excludes a second-column defect between two pair rows. Consequently the only fixed signed arrays are actual tableau bodies with trailing empty rows. The shape equations then force increasing, hence identity, permutation labels. The toggle swaps adjacent comparable entries in the row-major

word, preserving graph inversions, and adjacent distinct permutation labels, reversing their inversion parity. The Lean source proves these steps as `move_involution`, `move_signedArray`, `move_tableauGraphInversions`, `signedArray_fixed_iff_padded`, `fixed_labels_identity`, `fixed_row_lengths`, `move_sign_reversing`, and `fixed_permutationSign`. $\square$



same lower first entry; same inversion degree; opposite signs

Figure 4: The checked local toggle. The dashed box denotes an empty row, not a cell of the filling. The row labels p, q are exchanged along with the change of row lengths.

This local theorem supplies the cancellation move. The Mathlib bridge below proves the complete finite summation, the canonical padding/removal bijection with the existing tableau families, and compatibility between inversion parity and alternant permutation signs. These connect the actual coloring polynomial to the survivor expansion in Theorem 6.4.

6.2 Complete families and degreewise cancellation

Let $\mathcal{T}_k = \text{PT}_{(2^k, 1^{n-2k})}(P)$ when $2k \leq n$, and let $\mathcal{T}_k = \emptyset$ otherwise. Put $\mathcal{S}_k = \mathcal{T}_k \setminus B_k(\mathcal{T}_{k+1})$, where B_k is insertion at the final two-cell row. For a finite family A , write $C_d(A) = \#\{T \in A : \text{inv}_G(T) = d\}$.

Theorem 6.2 (Formal degreewise survivor identity). *For every finite poset of height at most two and every $k, d \in \mathbb{N}$,*

$$C_d(\mathcal{T}_k) = C_d(\mathcal{T}_{k+1}) + C_d(\mathcal{S}_k).$$

The empty vertex set has one empty tableau of degree zero. A single vertex also has one tableau of degree zero. At either an odd or even maximal index, the next source is empty and every target is a survivor.

Proof. Generate every filling of the specified shape from the finite vertex list, retain exactly the actual tableaux, and remove duplicates. Membership in this family is proved equivalent to tableau validity and the exact row-length list; completeness is not an input assumption. The source shape determines the split index k . Theorem 4.3 gives a weight-preserving injection into the target family. Its image in degree d has exactly $C_d(\mathcal{T}_{k+1})$ objects. Partitioning the target into this image and its complement proves the equality using addition of natural counts. The guard $2k \leq n$ handles every out-of-range index before forming $n - 2k$. $\square$

The declarations `mem_family` and `family_nodup` establish the enumeration properties. The all-index result is `height_two_degreewise_cancellation`. For a relation defined only on the vertex set, `finiteRelation` extends it by false outside that set, and `finite_height_two_degreewise_cancellation` requires asymmetry and height only on those vertices.

6.3 Actual polynomials and checked finite alternants

The separate polynomial project uses Lean 4.34.1 and Mathlib revision `d13f23b723b8a846827a245b89c10fc7d3f11612`. Its ring is $R_m = \mathbb{Z}[q][x_1, \dots, x_m]$. The definition of $X_{h,m}$ is the finite sum over actual proper functions $[n] \rightarrow [m]$, with their actual color monomials and ascent powers. The elementary polynomial is the actual subset sum $e_b = \sum_{|S|=b} \prod_{i \in S} x_i$. Tableau and survivor coefficients use the existing complete families; no equality with $X_{h,m}$ is assumed in any definition.




Figure 5: The checked finite decomposition. The weight-preserving injection identifies the source with one part of the target; the other part consists of survivors. No subtraction or polynomial-coefficient assumption is used.

For $\gamma \in \mathbb{N}^m$, define

$$A(\gamma) = \sum_{\sigma \in S_m} \text{sgn}(\sigma) \prod_{i=1}^m x_{\sigma(i)}^{\gamma_i}, \quad \delta = (m-1, m-2, \dots, 0), \quad \Delta = A(\delta).$$

The sign is the actual permutation sign. An empty product or empty permutation has value one.

Theorem 6.3 (Formal finite alternant ingredients). *For every $m, b \in \mathbb{N}$ and $\gamma \in \mathbb{N}^m$,*

$$A(\gamma)e_b = \sum_{\substack{S \subseteq [m] \\ |S|=b}} A(\gamma + \mathbf{1}_S).$$

If two distinct coordinates of γ are equal, then $A(\gamma) = 0$. The coefficient of $\mathbf{x}^\delta$ in Δ is one. Consequently $\Delta \neq 0$ and $\Delta p = \Delta t$ if and only if $p = t$ in R_m . For $m = 0$, $\Delta = 1$.

Proof. Elementary subset sums are unchanged by permuting variables. Multiply each signed monomial by this subset sum, distribute both finite sums, and move the permutation through the elementary factor. This gives the first identity, including $b > m$, when both sides are zero. A transposition of equal exponent positions fixes the monomial and reverses its sign, so $A = -A$. The concrete domain R_m has no zero divisors and $2 \neq 0$, hence $A = 0$.

The coordinates of δ are distinct. Only the identity permutation contributes to $\mathbf{x}^\delta$, with sign one. This proves the coefficient claim and nonvanishing; cancellation follows in the same domain. For zero variables the sole empty permutation contributes one.

The checked declarations are `alternant_mul_elementary`, `alternant_eq_zero_of_eq`, `vandermonde_delta_coefficient`, `vandermonde_ne_zero`, `vandermonde_cancel` and `vandermonde_empty`, all in namespace `PartitionH3`.

□

The two-elementary specialization is now checked:

$$\Delta e_a e_b = \sum_{j=0}^b A(\delta + (2^j, 1^{a+b-2j}, 0^{m-a-b+j})), \quad b \leq a, \quad a+b \leq m.$$

Actual subsets are classified by their two prefix lengths. A rise inside one block causes equal adjacent natural exponents, so its alternant vanishes. The surviving subsets have prefix lengths j and $b-j$. The declaration `vandermonde_mul_two_elementary` includes balanced and empty cases. The declaration `alternating_reconstruction` reconstructs an alternating polynomial from its actual finite support and dominant coefficients; an expansion is not assumed.

6.4 Actual colorings, signed arrays, and the final bridge

Proper color classes are chains: if $i < j$ have the same color, properness excludes an edge, giving $h_i < j$. Height two bounds a class by two vertices. For adjacent colors, let U be their support. Flip every isolated vertex of the graph induced on U , except that if $|U| = 3$ and there is no isolated vertex, flip all of U . This rule depends only on U , which it preserves, so it is an involution. The nontrivial three-vertex case is a path; natural unit interval monotonicity puts its center between its endpoints. Each coloring has one ascent on its two edges. Outside comparisons are unchanged because the colors are adjacent. The map preserves properness and ascent weight and exchanges the two multiplicities. Summing it proves symmetry.




Figure 6: The exceptional adjacent-color switch. If the induced graph has one edge, only its isolated vertex flips. These are the actual cases of `toggleColors`, not a free choice of a matching coloring.

Place sorted color classes in *increasing* color order, including empty classes, and attach permutation label $\sigma(i)$ to row i . Their lengths satisfy the literal integer equations

$$|R_i| = \lambda_{\sigma(i)} - \sigma(i) + i.$$

The inverse coloring assigns each vertex its unique row index. Sortedness and chain validity recover every original row. The list permutation sign is proved equal to the sign of σ . Original array graph inversions count *descents* of this coloring. Color reversal and the proved symmetry identify its descent polynomial with the ascent polynomial. Thus the coefficient of $x^{\delta+\lambda}$ in $\Delta X_{h,m}$ is exactly the signed weight sum over the original complete shifted-array universe.




Figure 7: The checked bridge. Increasing color rows use descents: on one edge with colors $(2, 1)$ their word is 21 , with one inversion. Reversal plus symmetry returns the requested ascent convention.

The bottommost-defect involution cancels every nonfixed array. Fixed labels are the identity and fixed rows are precisely tableaux with trailing empty rows. Removing those rows is a proved bijection with the complete original family, preserving its inversion statistic. A decreasing partition with row lengths at most two is necessarily $(2^k, 1^{n-2k}, 0^{m-n+k})$; all other sectors vanish. Dominant-coefficient reconstruction therefore gives

$$\Delta X_{h,m} = \sum_k W(\mathcal{T}_k) A(\delta + (2^k, 1^{n-2k}, 0^{m-n+k})).$$

H2 is lifted from every degree to the actual polynomial identity $W(\mathcal{T}_k) = W(\mathcal{T}_{k+1}) + W(S_k)$. The two-elementary identity and finite telescoping give the same alternant sum for $\Delta \sum_k W(S_k) e_{n-k} e_k$. Its boundary term is zero because $\mathcal{T}_{\lfloor n/2 \rfloor + 1}$ is empty. Nonzero Vandermonde cancellation proves the equality for $m \geq n$.

Theorem 6.4 (All-alphabet Lean survivor formula). *For every natural unit interval poset P_h of height at most two on n vertices and every finite alphabet size $r \geq 0$,*

$$X_{h,r}(x; q) = \sum_{k=0}^{\lfloor n/2 \rfloor} \left(\sum_{T \in S_k} q^{\text{inv}(T)} \right) e_{n-k}(x) e_k(x).$$

Here S_k is the original full image complement of the Cho–Huh injection.

Proof. The preceding construction proves the equality for $m \geq n$. Take $m = \max(n, r)$ and set $x_r, \dots, x_{m-1}$ to zero. A coloring using one of these colors contributes zero; the remaining colorings are in bijection with actual r -colorings. Elementary subset products restrict by the same rule. This argument includes $r = 0$ and $n = 0$ without dividing by a variable or assuming a nonempty coloring type. $\square$

The exact declaration is `PartitionH3.height_two_survivor_expansion`. Its only mathematical assumption on the supplied natural unit interval poset is `HeightTwo`; it quantifies over every r . Its direct Lean check reports only `propext`, `Quot.sound`, and `Classical.choice`. No proof placeholder, custom axiom, assumed Schur expansion, or final coefficient equality is an input. This formalizes the established height-two theorem and construction of Cho–Huh [1], following the finite coloring and alternant route of Shareshian–Wachs [2]. The intermediate Schur expansion is not itself claimed as a separate Lean theorem.

The signed calculation takes place over $\mathbb{Z}[q]$. To return to $\mathbb{N}[q]$, define the coloring polynomial, elementary subset products and survivor weights by the same actual sums in the natural coefficient semiring. The coefficient embedding $\mathbb{N} \hookrightarrow \mathbb{Z}$ induces injective maps on both polynomial levels, and maps each of these sums to its integer counterpart. Pulling back the equality gives the exact natural-coefficient declaration `PartitionH3.height_two_survivor_expansion_nat`, with the same height and alphabet scope and standard axioms.

The reproducible development preserves the separate five-module, Lean 4.34.0 Std certificate. Its Mathlib bridge has twenty-one modules pinned to Lean 4.34.1 and the recorded Mathlib revision. The verification entry point checks fresh project builds, printed signatures and standard-axiom reports; semantic review and its source hashes are recorded separately in the claim-to-theorem index. Fresh project builds check the original five Std modules and all twenty-one Mathlib modules, with respectively 54 and 137 principal dependency reports limited to standard axioms. Separate integration reviews support the integer formula and natural-coefficient transport at their recorded source hashes. These checks and reviews close H3 for height two; the general-height objective remains outside this certificate.

7 A worked three-vertex example

Take $h = (2, 3, 3)$. The only comparability is $1 <_P 3$, and G_h is the path 1–2–3. Write a singleton-row tableau as its top-to-bottom word. The complete family $\mathcal{T}_0$ consists of 123, 132, 213, 321, with weights $1, q, q, q^2$. The only member of $\mathcal{T}_1$ is $[1, 3]/[2]$, with weight q . Insertion sends it to 132. Therefore

$$W(\mathcal{S}_0) = 1 + q + q^2, \quad W(\mathcal{S}_1) = q, \quad X_{h,r} = (1 + q + q^2)e_3 + qe_2e_1.$$

The last equality is an instance of the established mathematical theorem and of the all-alphabet Lean survivor formula above.

On two colors, $e_3 = 0$ and the two proper colorings give $q(x_1^2x_2 + x_1x_2^2) = qe_2e_1$. With three distinct colors, the coefficient of $x_1x_2x_3$ is $1 + 4q + q^2$. The same coefficient in e_2e_1 is 3, so subtracting its contribution leaves $1 + q + q^2$. The independent coloring oracle and the tableau enumerator both give these exact values.




Figure 8: The path has exactly two proper colorings on two colors. The sole source tableau maps to 132, leaving 123, 213, 321 as survivors in the singleton family.

8 Finite evidence from two independent calculations

The primary enumerator generates permutations of labels, filters by actual tableau conditions, computes inversion degrees, and expands the integer determinant (2). A separate oracle enumerates proper colorings with prescribed ordered color multiplicities and expands elementary products as choices of subsets. Exact rational elimination recovers elementary coefficients. The oracle uses neither tableaux nor Jacobi–Trudi.

The two methods agree for every Hessenberg vector through five vertices. The comparable-prefix map is additionally tested on all height-two vectors through seven vertices. These are finite checks of implementations, not inductive proofs. The deterministic verification script and JSON reports record the exact ranges and counts.

Number of vertices	2	3	4	5	6	7
Height-two orders	2	4	8	16	32	64
Source tableaux checked	1	5	36	282	2730	30420
Failures	0	0	0	0	0	0

For the connected five-vertex witness from Figure 1, both methods give

$$\begin{aligned}
 c_5(q) &= 1 + 3q + 4q^2 + 4q^3 + 4q^4 + 3q^5 + q^6, \\
 c_{4,1}(q) &= q + 3q^2 + 4q^3 + 3q^4 + q^5, \\
 c_{3,2}(q) &= q^2 + 2q^3 + q^4,
 \end{aligned}$$

and no other nonzero coefficients. The false top-only slide and the true positivity statement can therefore be tested on the same small object.

The checks count graph inversions exactly, not a surrogate statistic. They also test the empty graph, complete graphs, the four-vertex path, the missing suffix boundary, and collisions when the height hypothesis is removed. The largest test range is deliberately explicit; no claim about larger instances follows from the enumeration alone.

9 What the general solution still needs

The repaired map closes the height-two argument, a known result. For general natural unit interval orders, rows can be longer. The full determinant then has additional signed sectors, and the shape-to-shape strategy is impossible in the form proposed in the earlier notes (Theorem 3.2). A proof using this route would need all of the following:

1. An actual configuration space of tableaux paired with determinant permutations, including zero-index factors and every contributing shape.
2. A uniform map preserving validity and inversion degree while allowing movement between the wider signed sectors.

3. An explicit inverse or a terminating cancellation argument, with a proof of positive fixed points and no unpaired negative configurations.
4. A formal bridge from those configurations to the exact elementary coefficients. A finite lookup table cannot replace it.

The $n = 5$ total-order identity (3) is a minimal regression test for a candidate generalization. The connected height-two witness tests the local boundary separately. A future proposed map must pass both before its general proof is credible.

This investigation does not establish a new open-problem solution. It replaces an invalid purported solution with precise obstructions, a complete restricted mathematical proof, and an all-alphabet Lean formalization of that restricted result. No inference about AI architecture quality follows from this research run.

References

- [1] S. Cho and J. Huh. *On e -positivity and e -unimodality of chromatic quasisymmetric functions*. SIAM J. Discrete Math. 33 (2019), 2286–2315. Section 3.1, injection ψ_ℓ and Theorem 3.3. [arXiv:1711.07152](#).
- [2] J. Shareshian and M. L. Wachs. *Chromatic quasisymmetric functions*. Adv. Math. 295 (2016), 497–551. Definition 6.1 and Theorem 6.3. [arXiv:1405.4629](#).
- [3] T. Hikita. *A proof of the Stanley–Stembridge conjecture*. Version 2, 25 December 2025. [arXiv:2410.12758](#).
- [4] B. E. Sagan and F. Tom. *Chromatic symmetric functions and change of basis*. Algebraic Combinatorics (2026). [doi:10.5802/alco.468](#).
- [5] OpenAI. *Elementary positivity of chromatic quasisymmetric functions*. Preprint, 24 September 2026, Theorem 1.1. [Primary preprint](#). Statement and selected proof sections inspected; full proof unverified here.