

Matroid log-concavity in Lean.

An explanatory companion to the strongest Mason inequality and the Heron–Rota–Welsh theorem.

About this writeup

This is a mathematical companion to two accepted Lean formalizations: the strongest Mason inequality for independent sets, and the Heron–Rota–Welsh theorem for characteristic polynomials. It explains the statements and gives a guided route through the arguments used by the accepted code. The hyperlinks point to named declarations in immutable source snapshots.

The exposition was prepared after the formal proofs were accepted. It is a guide to those artifacts, rather than the blueprint from which they were originally produced. The original mathematical theorems are credited below; no claim of a new theorem is made.

Independence and log-concavity

A **matroid** consists of a ground set E and a collection of subsets called independent sets. The empty set is independent; every subset of an independent set is independent; and if A and B are independent with $|A| < |B|$, some element of $B \setminus A$ can be added to A while preserving independence. The last property is the exchange axiom.

Linear independence is an example. So is acyclicity: take the edges of a graph as the ground set, and its forests as independent sets. A matroid’s rank function $r_M(S)$ records the largest size of an independent subset of S .

A nonnegative sequence $(a_0, \dots, a_d)$ is **log-concave** if

$$a_k^2 \geq a_{k-1}a_{k+1} \quad (1 \leq k < d).$$

Where the entries are positive, this means the ratios a_k/a_{k-1} do not increase with k . It is a constraint on how the sequence grows and falls. It does not mean that the entries themselves decrease.

The two theorems concern different sequences attached to the same finite matroid. Mason counts independent sets. Heron–Rota–Welsh studies coefficients of a polynomial formed from the rank function. Both conclusions are log-concavity, but Mason’s binomial normalization makes its assertion stronger than ordinary log-concavity of the counts.

A small example

The uniform matroid $U_{3,5}$ has five elements, with precisely the sets of size at most three declared independent. The independent-set counts in sizes $0, \dots, 5$ are

$$(I_0, I_1, I_2, I_3, I_4, I_5) = (1, 5, 10, 10, 0, 0).$$

Dividing each count by the number of all subsets of that size gives

$$\left(\frac{I_k}{\binom{5}{k}} \right)_{k=0}^5 = (1, 1, 1, 1, 0, 0).$$

This is the sequence in the strongest Mason inequality. Its entries are the probabilities that uniformly chosen subsets of each size are independent.

The same matroid has characteristic polynomial

$$\chi_{U_{3,5}}(t) = t^3 - 5t^2 + 10t - 6.$$

Its unsigned coefficients in ascending degree are $(6, 10, 5, 1)$. The two nontrivial Heron–Rota–Welsh inequalities read $10^2 \geq 6 \cdot 5$ and $5^2 \geq 10 \cdot 1$. These are a different sequence and a different theorem.

Two proof routes

The accepted Mason formalization uses a **combinatorial atlas**: a family of quadratic forms constructed from counts of ordered independent extensions. A local-to-global argument transports a reverse Cauchy–Schwarz inequality through the family. Evaluating the root form gives the normalized counting inequality.

The accepted Heron–Rota–Welsh formalization first rewrites coefficients as sums of Möbius weights on flats. It constructs an explicit quadratic form from incidences between adjacent ranks and proves the required Hodge-index inequality using localization and a weighted graph Laplacian. A coloop reduction transfers the result from the reduced characteristic polynomial to the original one.

The common linear-algebra theme is a symmetric bilinear form Q with at most one positive direction. For a vector h with $Q(h, h) > 0$, nonpositivity on the hyperplane $Q(h, x) = 0$ gives

$$Q(h, h)Q(v, v) \leq Q(h, v)^2.$$

Indeed, apply nonpositivity to $v - Q(h, v)h/Q(h, h)$ and expand. The difficult work is constructing the right form, proving its sign property, and identifying these three pairings with consecutive terms of the desired sequence.

Sources and scope

Brändén and Huh’s [Lorentzian polynomials](#) contains the strongest Mason inequality as Theorem 4.14. Anari, Liu, Oveis Gharan, and Vintzant proved it independently in [Log-Concave Polynomials III](#). The accepted code here uses the kind of combinatorial atlas method described by Chan and Pak in [Introduction to the combinatorial atlas](#).

Adiprasito, Huh, and Katz established the Heron–Rota–Welsh theorem for arbitrary matroids in [Hodge theory for combinatorial geometries](#). These results form part of the broader geometric approach to combinatorial log-concavity highlighted in [June Huh’s 2022 Fields Medal citation](#).

The artifacts establish the two specified matroid inequalities. They do not constitute formalizations of all tropical geometry, all Lorentzian polynomial theory, or the entire Hodge–Riemann theory for matroid Chow rings.

Both submissions passed the official and wide comparator lanes in Axiom Proofs. The [verification page](#) records the accepted versions and their provenance. Source repositories are private to AxiomMath, so the Lean and workflow links require repository access; this explanatory site and PDF are publicly readable.

Strong Mason

The theorem

Let M be a matroid on a finite ground set E of cardinality n . Write

$$I_j = \#\{S \subseteq E : |S| = j \text{ and } S \text{ is independent in } M\}.$$

The strongest form of Mason's conjecture, now a theorem, says that the sequence $I_j / \binom{n}{j}$, for $0 \leq j \leq n$, is log-concave. Thus, whenever $1 \leq k < n$,

$$\left(\frac{I_k}{\binom{n}{k}} \right)^2 \geq \frac{I_{k-1}}{\binom{n}{k-1}} \frac{I_{k+1}}{\binom{n}{k+1}}.$$

The normalization matters: $I_k / \binom{n}{k}$ is the probability that a uniformly chosen k -element subset is independent. The theorem controls the shape of these probabilities, giving a stronger conclusion than ordinary log-concavity of the counts I_k .

The accepted Lean statement clears the binomial denominators:

$$\binom{n}{k-1} \binom{n}{k+1} I_k^2 \geq \binom{n}{k}^2 I_{k-1} I_{k+1} \quad (k > 0).$$

This formulation remains meaningful beyond the interior of the sequence. No representability, simplicity, looplessness, or connectedness assumption is imposed. The matroid may have loops, parallel elements, or rank zero. The exported result is [branden_huh_strongest_mason](#).

Brändén and Huh prove this result as Theorem 4.14 of *Lorentzian polynomials*; Anari, Liu, Oveis Gharan, and Vintzant independently proved the strongest inequality. The formal development presented here takes a combinatorial-atlas route, closely related to the method explained by Chan and Pak. Its central objects are finite matrices of extension counts and a reverse Cauchy–Schwarz inequality.

The counting convention

The benchmark represents E by a finite type and explicitly assumes that the matroid's ground set is the entire type. Consequently, n counts actual ground elements. Its independent-set count filters the finite collection of all k -element subsets by the predicate of matroid independence.

The atlas development uses a second, set-based cardinality expression. [independentSetCount_eq_challenge](#) proves exact equality between these definitions. This is a substantive interface check: the final theorem counts the subsets specified by the benchmark, with the same cardinality and independence predicate.

For the main argument, fix $0 < k < n$. If $I_{k+1} = 0$, the desired inequality is immediate. Otherwise, a positive count supplies an independent set J of cardinality $k + 1$. This witness will justify the positivity assumptions needed by the matrix induction. See [strong_mason_finite](#).

1. Replace subsets by ordered continuations

For a finite prefix S , let $C(S, r)$ count ordered strings of r distinct new elements whose union with S is independent. The recursion is

$$C(S, 0) = \begin{cases} 1 & S \text{ independent,} \\ 0 & \text{otherwise,} \end{cases} \quad C(S, r+1) = \sum_{x \in E \setminus S} C(S \cup \{x\}, r).$$

Repetition is excluded at each insertion. Let $\mathcal{E}(S, r)$ be the collection of independent supersets of S with cardinality $|S| + r$. Double counting a distinguished newly added element proves

$$C(S, r) = r! |\mathcal{E}(S, r)|, \quad C(\emptyset, r) = r! I_r.$$

The factorial records the possible orders of the new elements. The recursive definition makes choosing the first letter an exact identity, which later becomes an identity between matrix entries. These facts are formalized in [continuation](#) and [continuation_eq_factorial_mul_card_extensionLevel](#).

Matroid augmentation also proves that $C(S, r) > 0$ whenever S is independent and $|S| + r \leq |J|$. Here J need not contain S : augmentation supplies a sufficiently large independent extension of S .

2. Put the strongest normalization into the weights

Define a coefficient depending on the final cardinality:

$$c(j) = \begin{cases} \frac{n-k+1}{n-k} & j = k+1, \\ 1 & j \neq k+1, \end{cases} \quad W(S, r) = c(|S| + r) C(S, r).$$

The denominator is positive because $k < n$. Adding the first element increases the prefix size by one and decreases the remaining length by one. Their sum stays fixed, so the weighted counts satisfy the same first-letter recursion.

At the empty prefix, the three relevant values are

$$W(\emptyset, k-1) = (k-1)! I_{k-1}, \quad W(\emptyset, k) = k! I_k,$$

$$W(\emptyset, k+1) = \frac{n-k+1}{n-k} (k+1)! I_{k+1}.$$

This single exceptional weight supplies the extra ground-set factor in the strongest Mason inequality. See [masonCoeff](#) and [weightedContinuation](#).

3. Arrange the counts into an atlas of matrices

Introduce an extra letter $*$ which leaves a prefix unchanged. A ground letter x inserts x , unless it is already present. For depth $m = r + 1$, form a matrix $A(S, m)$ indexed by $E \cup \{*\}$: its (i, j) entry is $W(T, r)$ if reading i, j from S validly produces T , and zero otherwise. Dependent extensions also give zero through the continuation count.

The matrices are symmetric and entrywise nonnegative. Reading one additional letter gives a child matrix. Commutativity of insertion and cyclic symmetry of three-letter insertion produce the inheritance identities used below. The construction is [atlasMatrix](#).

For any such matrix, write $B_A(v, w) = v^T A w$ and $Q_A(v) = B_A(v, v)$. The formal predicate [Hyperbolic](#) means

$$Q_A(v)Q_A(w) \leq B_A(v, w)^2 \quad \text{whenever } Q_A(v), Q_A(w) \geq 0.$$

This reverse Cauchy–Schwarz property is the target of the matrix induction. It is defined explicitly and proved for the atlas matrices.

4. Establish hyperbolicity at depth one

Suppose S is independent and $|S| \leq k - 1$. Its feasible one-element extensions split into parallel classes: two distinct feasible elements belong to the same class when adjoining both makes the prefix dependent. Matroid augmentation proves transitivity of this relation.

The depth-one matrix has star–star entry 1, star–feasible entries 1, zero entries within a parallel class, and a common coefficient c between different classes. Here $c = 1$ unless $|S| = k - 1$, when $c = (n - k + 1)/(n - k)$. Infeasible rows vanish.

For a vector v , let b_q be the sum of its coordinates in parallel class q , and set $s = \sum_q b_q$. Expanding gives

$$Q_A(v) = (v_* + s)^2 + (c - 1)s^2 - c \sum_q b_q^2.$$

On the hyperplane $B_A(v, e_*) = v_* + s = 0$, ordinary Cauchy–Schwarz gives $s^2 \leq \ell \sum_q b_q^2$, where ℓ is the number of classes. If $c = 1$, nonpositivity follows immediately. In the exceptional case,

$$\ell \leq n - |S| = n - k + 1, \quad \ell(c - 1) \leq c,$$

which again proves $Q_A(v) \leq 0$ on that hyperplane. A quadratic-form lemma converts this into hyperbolicity. Dependent prefixes and empty feasible sets are handled separately. See [atlasMatrix_one_hyperbolic](#).

5. Pass from children to parents

The linear-algebra engine considers a symmetric nonnegative parent A , a positive vector h , and hyperbolic children B_i satisfying

$$(Av)_i = B_i(v, h), \quad \sum_i h_i Q_{B_i}(v) = Q_A(v).$$

Child reverse Cauchy–Schwarz yields

$$Q_A(v) \leq \sum_i h_i \frac{(Av)_i^2}{(Ah)_i}.$$

Positive diagonal normalization turns this into $Q_T(x) \leq \|Tx\|^2$ for a symmetric nonnegative matrix T having a positive fixed vector p . A maximum-ratio argument bounds its eigenvalues by $|\lambda| \leq 1$; the energy inequality gives $\lambda \leq \lambda^2$. Thus every positive eigenvalue is 1. A strictly positive hub column forces that eigenspace to be one-dimensional. The spectral theorem then makes Q_T nonpositive on $p^\perp$, establishing hyperbolicity and transferring it back to A .

This argument is [hyperbolic_of_local_quadratic_data](#).

For the atlas, apply it to

$$A_t = tA(S, m+1) + (1-t)A(S, m), \quad 0 < t < 1,$$

with ground-letter weights t and star weight $1-t$. Remove infeasible coordinates, whose rows and columns are identically zero. The independent witness J ensures positivity of the needed continuations, including the star hub. The insertion identities provide inheritance and the weighted quadratic identity exactly.

Finally, hyperbolicity passes to the endpoint as $t \rightarrow 1$. Induction from depth one reaches $A(\varnothing, k)$. See [hyperbolic_right_endpoint](#) and [atlasMatrix_hyperbolic_iter](#).

6. Evaluate the root and clear the factors

Let u be one on ground coordinates and zero at the star, and let e_* select the star. At the root $A = A(\varnothing, k)$, expanding finite sums gives

$$Q_A(e_*) = W(\varnothing, k-1), \quad B_A(u, e_*) = W(\varnothing, k), \quad Q_A(u) = W(\varnothing, k+1).$$

All these counts are nonnegative, so root hyperbolicity yields

$$(k!I_k)^2 \geq (k-1)!I_{k-1} \frac{n-k+1}{n-k} (k+1)!I_{k+1}.$$

Canceling positive factorial and denominator factors gives

$$(k+1)(n-k+1)I_{k-1}I_{k+1} \leq k(n-k)I_k^2.$$

The evaluation is [normalized_inequality_of_root_hyperbolic](#). Binomial recurrences convert this into the benchmark’s exact statement. If $k \geq n$, both $\binom{n}{k+1}$ and I_{k+1} vanish, so both sides are zero. The final adapter proves this boundary case explicitly: [branden_huh_strongest_mason](#).

Scope and provenance

The accepted development formalizes this matroid inequality through combinatorial atlas matrices. It does not formalize the full theory of Lorentzian polynomials or tropical geometry. Mathlib supplies general infrastructure, including matroids, finite cardinalities, real symmetric spectral theory, and limits; the submitted modules assemble the missing argument.

[Submission issue #15](#) passed both Axiom Proofs comparator lanes on September 1, 2026. The frozen submission is commit `f0629e6b9c80a8aedef6a923892150735f56570a` ; its archived proof comprises 13 Lean source files and uses Lean 4.33.0. The verification record reports only the standard foundational dependencies `propext` , `Classical.choice` , and `Quot.sound` .

The recorded final worker used Codexiom with `openai/gpt-5.6-sol-pro` . It received a preloaded, audited AxiomMath atlas development, checked the dependency chain and counting conventions, completed the exact benchmark adapter, and removed warnings. Its recorded runtime and model cost therefore describe that final run, not the entire preceding development. Human involvement is recorded as theorem selection, orchestration, review requirements, and acceptance criteria. The accepted metadata and frozen proof blueprint preserve this distinction.

References

- Petter Brändén and June Huh, *Lorentzian polynomials*, *Annals of Mathematics* 192 (2020), 821–891. Theorem 4.14 gives the strongest Mason inequality.
- Nima Anari, Kuikui Liu, Shayan Oveis Gharan, and Cynthia Vinzant, *Log-Concave Polynomials III: Mason’s Ultra-Log-Concavity Conjecture for Independent Sets of Matroids* (2018), an independent proof through complete log-concavity.
- Swee Hong Chan and Igor Pak, *Introduction to the combinatorial atlas* (2022), especially the local-to-global framework and its application to Mason’s inequality.
- AxiomMath, [frozen proof blueprint and audit record](#), and [acceptance workflow](#).

Heron–Rota–Welsh

The theorem

Let M be a matroid on a finite ground set E , with rank function r_M . Its characteristic polynomial is

$$\chi_M(t) = \sum_{S \subseteq E} (-1)^{|S|} t^{r_M(E) - r_M(S)}.$$

Write $c_k(M) = |[t^k]\chi_M(t)|$. The **Heron–Rota–Welsh theorem** says that these unsigned coefficients are log-concave:

$$c_{k+1}(M)^2 \geq c_k(M)c_{k+2}(M) \quad (k \geq 0).$$

The Lean statement uses coefficients in **ascending degree**. This reverses the indexing often used in the literature, without changing the inequality. It quantifies over every finite matroid whose ground set is the whole ambient finite type; there is no representability assumption. Loops are allowed, and coefficients beyond the polynomial’s degree are zero. [Formal entry point](#).

Adiprasito, Huh, and Katz proved the theorem using the Hodge–Riemann relations for matroid Chow rings. Their [2018 paper](#) treats arbitrary matroids; the earlier [Huh–Katz paper](#) establishes the result for realizable matroids through intersection theory.

The accepted formal proof follows an explicit route through weighted bilinear forms on the lattice of flats. It proves the Hodge-index inequality needed for the coefficient sequence. It does not formalize the complete Chow-ring construction, hard Lefschetz theorem, or full Hodge–Riemann package of the original paper.

1. Replace subsets by flats

A flat is a subset closed under matroid closure. Every subset has a unique closure, so the defining sum for χ_M can be grouped by flats. For a flat F , define

$$\omega(F) = \sum_{\text{cl}(S)=F} (-1)^{|S|}.$$

For a loopless matroid, the proof identifies $\omega(F)$ with $\mu(\emptyset, F)$, the Möbius function of the flat lattice. Its cumulative sums are 1 at the bottom flat and 0 elsewhere, which characterizes this Möbius function. Consequently,

$$\chi_M(t) = \sum_F \mu(\emptyset, F) t^{r_M(E) - r_M(F)}.$$

The next ingredient is the sign rule $(-1)^{r_M(F)} \omega(F) \geq 0$. Thus the nonnegative weight $w(F) = |\omega(F)|$ can replace signed Möbius weights whenever the rank is fixed. The proof establishes this sign rule using cancellation after adjoining a distinguished element to a flat, followed by induction on rank.

Sources: [closure-fiber weights](#), [Möbius identification](#), [alternating signs](#).

2. Identify the reduced coefficients

For a nonempty loopless matroid of rank r , set $\bar{\chi}_M(t) = \chi_M(t)/(t-1)$ and

$$b_j(M) = |[t^{r-1-j}] \bar{\chi}_M(t)| \quad (0 \leq j < r).$$

Set $b_j(M) = 0$ for $j \geq r$. This extension makes the consecutive-coefficient formulas below meaningful at every $j \geq 0$; it also agrees with the avoiding-flat sum, since the top flat contains e and no flat has rank greater than r .

Fix an element $e \in E$. The proof obtains two descriptions of the same number:

$$b_j(M) = \sum_{\substack{r_M(F)=j \\ e \notin F}} w(F) = \sum_{\substack{r_M(F)=j+1 \\ e \in F}} w(F).$$

The first formula comes from the relation between the characteristic and reduced characteristic polynomials. The second uses closure with e : a flat avoiding e is sent to $\text{cl}(F \cup \{e\})$, increasing its rank by one. The relevant Möbius cancellation identity gives the equality of the weighted sums. This is a weighted identity, not an assertion that the two collections of flats have equal cardinality.

These formulas are the bridge from polynomial coefficients to the explicit bilinear form below. [Reduced polynomial coefficient identity](#), [avoiding-flat formula](#), [containing-flat formula](#).

3. Build a form on adjacent ranks

Let T be a flat, and let $s \geq 1$. On real-valued functions on the flats, define

$$\begin{aligned} Q_{s,T}(x, y) = & \sum_{\substack{F \leq T \\ r_M(\bar{F})=s}} w(F) x_F y_F + \sum_{\substack{G \leq T \\ r_M(G)=s+1}} w(G) x_G y_G \\ & - \sum_{\substack{F \leq T \\ r_M(\bar{F})=s}} w(F) \sum_{F \triangleleft G \leq T} (x_F - x_G)(y_F - y_G). \end{aligned}$$

Here $F \triangleleft G$ means that G covers F : no flat lies strictly between them. The last term is a weighted graph energy along cover relations. All ingredients are finite sums, ranks, containment, and the weights already defined.

Take $T = E$, and let u_F indicate that $e \in F$, while v_F indicates that $e \notin F$. The coefficient identities give

$$Q_{j+1,E}(u, u) = b_j, \quad Q_{j+1,E}(u, v) = b_{j+1}, \quad Q_{j+1,E}(v, v) = b_{j+2}.$$

For example, an avoiding flat has exactly one cover containing e , namely its closure after adjoining e . Its contribution to the graph energy therefore cancels its contribution on the lower diagonal. This gives the last pairing identity. The mixed identity also uses the Möbius-weighted cancellation from the preceding step.

Sources: [definition of the form](#), [containing-vector pairing](#), [avoiding-vector self-pairing](#).

4. Establish the Hodge-index inequality

The central theorem asserts that, whenever $Q_{s,T}(a, a) > 0$,

$$Q_{s,T}(a, a)Q_{s,T}(x, x) \leq Q_{s,T}(a, x)^2.$$

This reverse Cauchy–Schwarz inequality expresses that the form has at most one positive direction. It is proved for every rank and every ceiling flat containing a chosen element. [All-ranks theorem](#).

At rank one, call the rank-one flats points and the rank-two flats lines. Each pair of distinct points lies on a unique line, and a line’s weight is its number of points minus one. Expanding the form produces a square of the total point sum minus a sum of squares indexed by lines. On the hyperplane where the point sum is zero, the form is therefore nonpositive. Elementary bilinear algebra gives reverse Cauchy–Schwarz. [Atomic nonpositivity](#).

The inductive step studies the deformation

$$B_t = tQ_{s,T} + Q_{s+1,T}, \quad t > 0,$$

with reference vector $h_t(F) = t$ when $e \in F$ and $h_t(F) = 1$ otherwise. Only flats of ranks $s, s+1, s+2$ contribute. For each flat in this three-rank window, the proof constructs a local form. At the lower rank the required inequality reduces to an atomic interval calculation; at the middle rank it is an explicit two-functional calculation; at the upper rank it follows from the induction hypothesis applied below that flat. [Local forms](#).

5. Assemble the local inequalities

The passage from those local calculations to B_t is the main assembly argument. The local forms obey exchange identities and decompose the global form. They also have positive norms at h_t .

From these norms, the proof constructs a positive diagonal form D such that $K = D - B_t$ is a weighted graph Laplacian after rescaling coordinates by h_t . Its quadratic form is a sum of nonnegative edge energies. The cover graph on the three-rank window is connected, so its Laplacian equation is solvable for every forcing term of total mass zero. [Connectivity](#), [Laplacian solvability](#).

If $B_t(h_t, x) = 0$, the resulting forcing term has zero total mass. A solution supplies a vector z with $Kz = Dx$. Summing the local reverse Cauchy–Schwarz inequalities gives

$$B_t(z, z) \leq D(z - x, z - x).$$

Together with $Kz = Dx$, this implies $D(z, x) \leq D(x, x)$. Applying $K \geq 0$ to $x - z$ then yields

$$B_t(x, x) \leq D(z, x) - D(x, x) \leq 0.$$

Thus B_t is nonpositive on its primitive hyperplane. The proof derives reverse Cauchy–Schwarz for B_t and then sends t to zero. In Lean, the last passage is an explicit small-parameter argument bounding the error terms of a quadratic polynomial, rather than an appeal to a spectral continuity theorem. This completes the rank induction for $Q_{s+1,T}$.

6. Recover log-concavity

Apply the all-ranks inequality to the containing and avoiding indicators from Step 3. When $b_j > 0$, it gives immediately

$$b_j b_{j+2} \leq b_{j+1}^2.$$

If $b_j = 0$, the same inequality follows from nonnegativity. The formal implementation packages the form, the two vectors, the pairing identities, and primitive nonpositivity as a finite Hodge witness, then applies a general coefficient reduction. [Incidence witness](#), [reduced log-concavity](#).

To return to χ_M , adjoin one coloop: let $N = M \oplus U_{1,1}$. Direct-sum multiplicativity gives

$$\chi_N(t) = (t - 1)\chi_M(t), \quad \overline{\chi}_N(t) = \chi_M(t).$$

Reduced log-concavity for N is therefore exactly the desired coefficient inequality for M , after reversing indices. The proof also handles the boundary cases: beyond rank the upper-degree coefficient vanishes; if M has a loop, subsets cancel in pairs and the entire characteristic polynomial is zero. This coloop argument includes the empty ground set. [Coloop identity](#), [final reduction](#).

Formalization and provenance

The accepted overlay contains 14 Lean files, totaling 26,918 lines. That is the size of the preserved submission, including supporting developments and alternative constructions; it is not a measurement of a minimized dependency closure.

The acceptance record dates to September 5, 2026. Both the official and wide Comparator lanes passed [submission #23](#). The recorded environment uses Lean 4.33.0 and mathlib commit [6f1ef4e5dd604a435bddba4747b13970cd65d2a1](#).

The submitted provenance describes a supervised Codexiom run lineage involving several models. It credits GPT-6 Astra at maximum reasoning effort with the decisive incidence construction and final proof, building on earlier coefficient reductions and exploration. It records human involvement in theorem selection, orchestration, recovery, model switching, and acceptance criteria, with no manually constructed Lean proof text. These are provenance claims in the submission record; the mathematical acceptance evidence is the checked theorem and the successful Comparator runs.

For historical background and the original Chow-ring proof, see Adiprasito–Huh–Katz, [Hodge theory for combinatorial geometries](#), Annals of Mathematics 188 (2018), 381–452. The explanation above follows the accepted Lean declarations and makes no claim that this particular presentation is a new mathematical proof.