The Inward Bow of a Real-Rooted Polynomial: Newton’s Inequalities, Machine-Checked in Lean 4

Inequalities for polynomials · real-rooted polynomials and log-concavity

Carles Marín

Statements

Theorem 1 — Newton’s inequalities, in Lean 4; Newton.newton_inequality 4

Lemma 1 — RealRooted.derivative , RealRooted.iterate_derivative 6

Lemma 2 — RealRooted.reverse 7

Lemma 3 — realRooted_discrim_coeff 8

Question 1 9

Question 2 9

Question 3 9

Question 4 10

Question 5 10

At a glance

What the paper contributes

A machine-checked, `sorry`-free proof in Lean 4 over Mathlib of **Newton's inequalities** (1707): for a real-rooted real polynomial of degree n, $e_k^2 \cdot C(n,k-1) \cdot C(n,k+1) \geq e_{k-1} \cdot e_{k+1} \cdot C(n,k)^2$ for $1 \leq k \leq n-1$; equivalently the normalized means $p_k = e_k / C(n,k)$ are log-concave.

- **The theorem** (Theorem 1): the symmetric-function form assumes p monic; the coefficient form `newton_inequality_coeff` assumes nothing.
- **The route** is the classical one (Sylvester; Hardy–Littlewood–Pólya): differentiate, reverse, differentiate again, down to a real-rooted quadratic, whose discriminant is the inequality once positive factorials are cleared.
- **Four reusable blocks**: real-rootedness under differentiation, iterated differentiation and reversal, and the coefficient-form discriminant for degree ≤ 2, which also absorbs the degenerate case $a_{i-1} = 0$.
- **Boundary**: the mathematics is classical; the paper claims the formalization (to the author's knowledge the first in any proof assistant, June 2026). The equality cases and Maclaurin's chain are named, not proved.

🚀 **Start here: Theorem 1**: Newton's inequalities for a monic real-rooted polynomial, and log-concavity of the normalized means.

Main results:

- **Theorem 1**: Newton's inequalities for a monic real-rooted polynomial, and log-concavity of the normalized means

9 results | ✅ **4** entirely in Lean | ◐ **0** partly in Lean

💻 **1** checked by a program | written proof only | (a result can carry several tags)

🗺️ Map of the results (click a box to go to it)

Dependencies in the written proof: an arrow A → B means that this proof of B uses A. Every result rests on its written proof; the tags record what was checked besides, and each statement's “Evidence & scope” block says how far. The colour of a box is its strongest tag, in the order above; the tags in the box give the rest. Boxes with a thick border are the main results; click a box to read the statement.

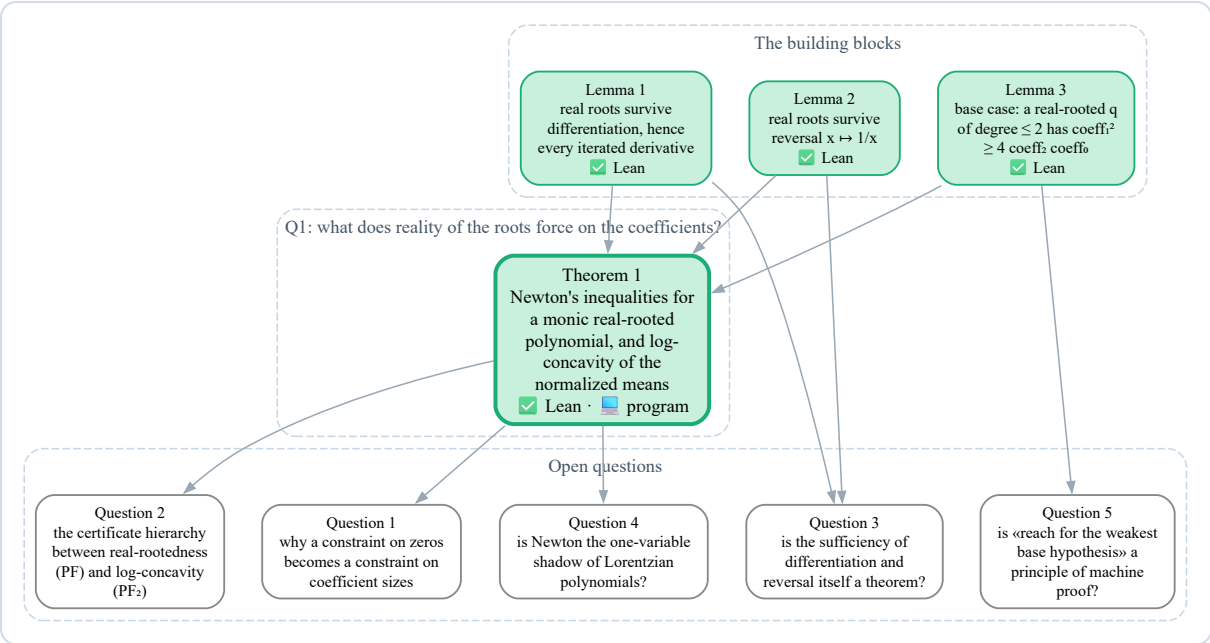

📐 Setting (copied from the paper, §1)

Write a real-rooted polynomial through its roots, $p(x) = \sum_{k=0}^{n} (-1)^k e_k \, x^{n-k}$, where e_k is the k-th *elementary symmetric function* of the roots, and form the *normalized means* $p_k = e_k / \binom{n}{k}$. *(§2)* I work over $\mathbb{R}$. A polynomial is *real-rooted* () when it splits into linear factors over $\mathbb{R}$— Mathlib's intrinsic ; equivalently it has $\deg p$ real roots counted with multiplicity. For the roots multiset $s = p.\texttt{roots}$, the elementary symmetric function is $e_k = \texttt{Multiset.esymm}\, s\, k$, and $\binom{n}{k} = \texttt{Nat.choose}$.

Introduction

No numbered statement in this section.

Its figures and tables are below; its text is in the paper (PDF).

🖼️ Figure 1

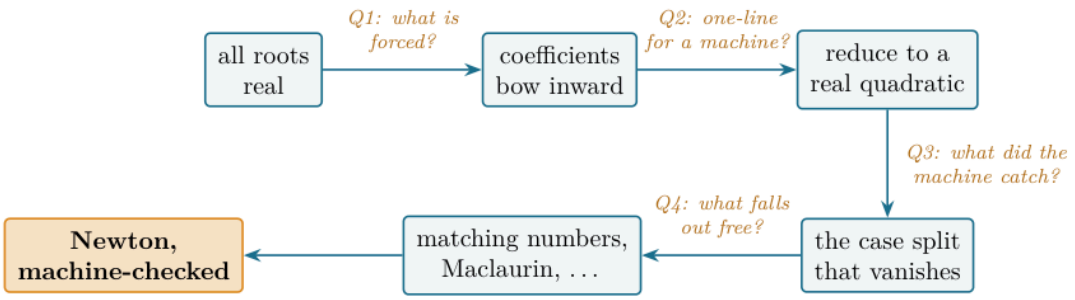

The reader's map. From a hypothesis on the roots to a certified inequality and its corollaries; each arrow is the question its section answers. Every node is a block of `sorry`-free Lean.

1 Q1: what does reality of the roots force on the coefficients?

Theorem 1. ✅ Lean

(Newton's inequalities, in Lean 4; `Newton.newton_inequality`)✓ 📕 p. 2

> *For a monic real-rooted real polynomial p of degree n and every $1 \le k \le n-1$,*
>
> $$e_k^2 \binom{n}{k-1}\binom{n}{k+1} \;\ge\; e_{k-1}\, e_{k+1} \binom{n}{k}^2 .$$
>
> *Equivalently (`newton_logConcave`) the normalized means $p_k = e_k / \binom{n}{k}$ are log-concave, $p_{k-1}\, p_{k+1} \le p_k^2$. Both are `sorry`-free; `#print axioms` on each reports only `propext`, `Classical.choice`, `Quot.sound`.*

Lean · ours show

- `Newton.newton_inequality`
- `Newton.newton_logConcave`
- `Newton.newton_inequality_coeff`

✅ Checked when this page was built: NewtonInequalities.lean compiles with Lean 4.35.0, no errors, no sorry. Axioms — newton_inequality, newton_logConcave, newton_inequality_coeff: propext, Classical.choice, Quot.sound.

Reproduce: Lean 4.35.0, Mathlib f0b4c8150aa0; in that Lean project, `lake env lean NewtonInequalities.lean` (SHA-256 of the compiled file: `69dd05da620a4eb0`…; a source page on this site may show it with comments edited, under its own hash).

Trust base: 🔒 kernel only. Lean: 3 results, standard axioms only

Proof: in this note · **Formalized:** Lean 4, by the author (public formalizations repository) · **Covers:** the whole statement · **Cite:** Theorem 1, p. 2, doi:10.5281/zenodo.20693064

🆕 **Classical and new:** classical: stated by Newton (1707), proved by Sylvester, textbook account in Hardy–Littlewood–Pólya; the paper claims the formalization, to the author's knowledge the first in any proof assistant (searches of June 2026, §6).

↳ Used in: 1 (p. 9), 2 (p. 9), 4 (p. 10)

Proof

Formalized in Lean (see the linked declarations).

💡 **In plain words (commentary, not part of the paper)**

1. If all roots of a polynomial are real, its coefficients cannot be arbitrary: the normalized means $p_k = e_k / \binom{n}{k}$ satisfy $p_{k-1}p_{k+1} \le p_k^2$, so on a log scale they bow inward. For $(x-1)(x-2)(x-4) = x^3 - 7x^2 + 14x - 8$ the two inequalities read $147 \ge 126$ and

$588 \geq 504$; $(x-1)^4$ gives equality everywhere, and x^2+1, whose roots are not real, fails ($0 \not\geq 4$). 2. Why it is true (the classical proof, here checked by Lean): differentiation (Rolle) and reversal $x \mapsto 1/x$ both keep all roots real. Differentiating $i-1$ times, reversing and differentiating again leaves a real-rooted polynomial of degree at most 2 whose three coefficients are positive factorial multiples of a_{i-1}, a_i, a_{i+1}; its discriminant is ≥ 0, and that is the inequality once the factorials are cleared. The base case is stated with deg ≤ 2, which also covers $a_{i-1} = 0$ without a separate case. 3. New and not new, as the paper says: the inequalities are Newton's (1707), proved by Sylvester, in Hardy–Littlewood–Pólya; the paper claims the formalization, to the author's knowledge the first in any proof assistant (§6). Lean checks the symmetric-function form for monic p (no loss: the e_k depend only on the roots) and the coefficient form with no normalization; the equality cases and Maclaurin's chain are not proved.

Evidence & scope

- **Program:** the worked instances: e_k of (x−1)(x−2)(x−4), the Newton gaps 21 and 84, the zero gaps of $(x-1)^4$ and the gap −4 of x^2+1 (numpy, with tolerance); script in the paper repository, not in the Zenodo record. `figures/newton_fig1_bow.py`

Figure 2

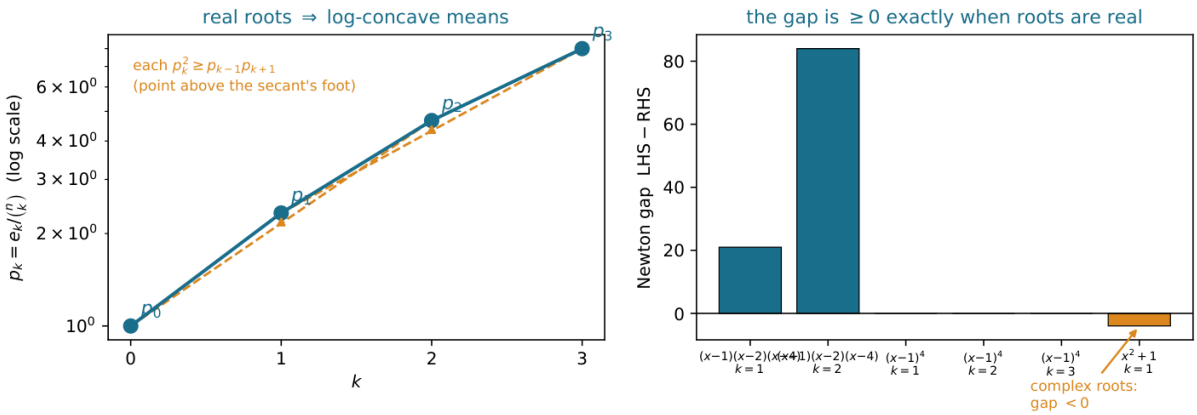

Left: the normalized means of $(x-1)(x-2)(x-4)$, plotted on a log axis; each interior mean sits above the geometric mean of its neighbours (the amber arrows), so the sequence is concave—it "bows inward." *Right:* the Newton gap $e_k^2\binom{n}{k-1}\binom{n}{k+1} - e_{k-1}e_{k+1}\binom{n}{k}^2$ at every interior k: positive for the real-rooted cubic, exactly zero for $(x-1)^4$ (the sharp case), and *negative* for x^2+1. The generating script asserts every plotted value against exact integer arithmetic.

2 Q2: how do you make it a one-line check for a machine?

No numbered statement in this section.

Its figures and tables are below; its text is in the paper, p. 3.

Figure 3

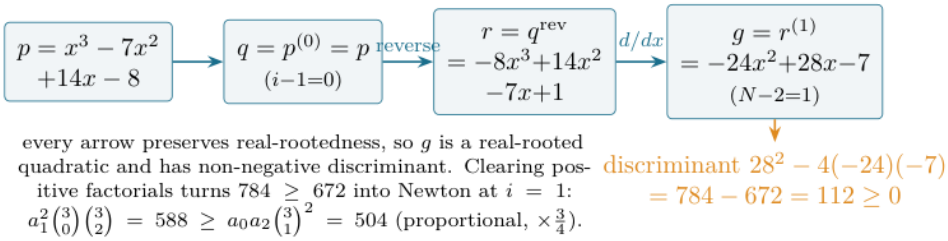

The differentiate–reverse reduction, worked on the running cubic at $i = 1$ (relating $a_0, a_1, a_2 = -8, 14, -7$). The final object is a quadratic; a real-rooted quadratic has non-negative discriminant; the factorials the proof cancels are the proportionality constant.

3 Q3: what did the machine catch that the textbook glides over?

No numbered statement in this section.

It is prose only (no statement, figure or table): read it in the paper, p. 5.

4 Q4: with Newton certified, what falls out for free?

No numbered statement in this section.

It is prose only (no statement, figure or table): read it in the paper, p. 5.

5 The building blocks

Lemma 1. Lean (`RealRooted.derivative`, `RealRooted.iterate_derivative`)✓

p. 6

If p is real-rooted then so is p', and hence every iterated derivative $p^{(k)}$.

Lean · ours **show**

- `Newton.RealRooted.derivative`

- `Newton.RealRooted.iterate_derivative`

✅ Checked when this page was built: NewtonInequalities.lean compiles with Lean 4.35.0, no errors, no sorry. Axioms — derivative, iterate_derivative: propext, Classical.choice, Quot.sound.
Reproduce: Lean 4.35.0, Mathlib f0b4c8150aa0; in that Lean project, `lake env lean NewtonInequalities.lean` (SHA-256 of the compiled file: `69dd05da620a4eb0…`; a source page on this site may show it with comments edited, under its own hash).

Trust base: 🔒 kernel only. Lean: 2 results, standard axioms only

Proof: in this note · **Formalized:** Lean 4, by the author (public formalizations repository) · **Covers:** the whole statement · **Cite:** Lemma 1, p. 6, doi:10.5281/zenodo.20693064

🆕 **Classical and new:** the paper: one of «four reusable building blocks» (§5); Rolle for polynomials is Mathlib's card_roots_le_derivative.

↳ Used in: Theorem 1 (p. 4), 3 (p. 9)

Proof

Formalized in Lean (see the linked declarations).

Lemma 2. ✅ Lean (`RealRooted.reverse`)✓ 📕 p. 6

If p is real-rooted then so is its reversal p^{rev}.

Lean · ours show

- `Newton.RealRooted.reverse`

✅ Checked when this page was built: NewtonInequalities.lean compiles with Lean 4.35.0, no errors, no sorry. Axioms — reverse: propext, Classical.choice, Quot.sound.
Reproduce: Lean 4.35.0, Mathlib f0b4c8150aa0; in that Lean project, `lake env lean NewtonInequalities.lean` (SHA-256 of the compiled file: `69dd05da620a4eb0…`; a source page on this site may show it with comments edited, under its own hash).

Trust base: 🔒 kernel only. Lean: 1 result, standard axioms only

Proof: in this note · **Formalized:** Lean 4, by the author (public formalizations repository) · **Covers:** the whole statement · **Cite:** Lemma 2, p. 6, doi:10.5281/zenodo.20693064

🆕 **Classical and new:** the paper: a reusable block; «the constant-term hypothesis the textbook attaches is unnecessary and is dropped» (§5).

↳ Used in: Theorem 1 (p. 4), 3 (p. 9)

Proof

Formalized in Lean (see the linked declarations).

Lemma 3. Lean (`realRooted_discrim_coeff`)✓ p. 6

> *A real-rooted q with* $\deg q \le 2$ *satisfies* $q.\texttt{coeff}\ 1^2 \ge 4\, q.\texttt{coeff}\ 2 \cdot q.\texttt{coeff}\ 0$.

Lean · ours show

- `Newton.realRooted_discrim_coeff`

✅ Checked when this page was built: NewtonInequalities.lean compiles with Lean 4.35.0, no errors, no sorry. Axioms — realRooted_discrim_coeff: propext, Classical.choice, Quot.sound.
Reproduce: Lean 4.35.0, Mathlib f0b4c8150aa0; in that Lean project, `lake env lean NewtonInequalities.lean` (SHA-256 of the compiled file: `69dd05da620a4eb0…`; a source page on this site may show it with comments edited, under its own hash).

Trust base: 🔒 kernel only. Lean: 1 result, standard axioms only

Proof: in this note · **Formalized:** Lean 4, by the author (public formalizations repository) · **Covers:** the whole statement · **Cite:** Lemma 3, p. 6, doi:10.5281/zenodo.20693064

🆕 **Classical and new:** the paper: the base case in coefficient form, «the source of the dissolved case split» (§3, §5).

↳ Used in: Theorem 1 (p. 4), 5 (p. 10)

Proof

> Formalized in Lean (see the linked declarations).

Lean name	statement	role
`newton_inequality`	$e_k^2 \binom{n}{k-1}\binom{n}{k+1} \ge e_{k-1}e_{k+1}\binom{n}{k}^2$	Theorem 1
`newton_logConcave`	$p_{k-1}p_{k+1} \le p_k^2$ (normalized means)	corollary
`newton_inequality_coeff`	$a_i^2 \binom{n}{i-1}\binom{n}{i+1} \ge a_{i-1}a_{i+1}\binom{n}{i}^2$	coefficient form
`realRooted_discrim_coeff`	real-rooted, $\deg \le 2 \Rightarrow a_1^2 \ge 4a_2a_0$	base case
`RealRooted.derivative`	p real-rooted $\Rightarrow p'$ real-rooted	block (Rolle)
`RealRooted.iterate_derivative`	p real-rooted $\Rightarrow p^{(k)}$ real-rooted	block
`RealRooted.reverse`	p real-rooted $\Rightarrow p^{\mathrm{rev}}$ real-rooted	block

Table 1: The load-bearing results, all `sorry`-free in `Newton.lean` of the repository (346 lines, on `RealStable.lean` and Mathlib). Axioms of every entry: `propext`, `Classical.choice`, `Quot.sound` only; no `native_decide`, no custom axiom.

6 The boundary of the contribution

No numbered statement in this section.

It is prose only (no statement, figure or table): read it in the paper, p. 7.

7 Open questions

Question 1. p. 7

Newton's inequalities are now a theorem in the library. The concrete next targets—Maclaurin's mean chain, the equality cases, log-concavity of the matching numbers—are named in Section 4; what follows are the questions I would rather carry away than tick off, one per turn of the argument, the last one the one to keep.

1. *Why zeros become magnitudes.* Reality of the roots is a statement about *where* a polynomial vanishes; log-concavity is a statement about the *size* of its coefficients. Newton is the first and cheapest bridge between the two. What is the deepest reason a constraint on a polynomial's zeros should turn, window by window, into a constraint on the magnitude of its coefficients?
2. *The certificate hierarchy.* For a sequence of non-negative reals, real-rootedness of the generating polynomial is equivalent to being a Pólya frequency sequence (Aissen–Schoenberg–Whitney), strictly stronger than log-concavity (which is PF_2). Newton derives the weak end from the strong.

Cite: Question 1, p. 7, doi:10.5281/zenodo.20693064

↳ Not used by another result of the paper.

Question 2. p. 7

What exactly lives in the gap between the rungs, and is there a *cheapest* certificate for each that a proof assistant could be taught to choose? *The closure, made explicit.* Differentiation and reversal each preserve real-rootedness, and that pair alone already exposes any three consecutive coefficients as a quadratic.

Cite: Question 2, p. 7, doi:10.5281/zenodo.20693064

↳ Not used by another result of the paper.

Question 3. p. 7

What is the full monoid of real-rootedness-preserving operations on coefficient windows, and is “reaches every quadratic window” a theorem one could state and check—i.e. is the sufficiency of *these two* stones itself a theorem? *Above the real-rooted floor.* Where no real-rooted polynomial exists—matroids—log-concavity needs Hodge theory [6] , with Brändén’s Lorentzian polynomials interpolating between the two worlds.

Cite: Question 3, p. 7, doi:10.5281/zenodo.20693064

↳ Not used by another result of the paper.

Question 4. p. 7

Is Newton the one-variable shadow of Lorentzian-ness, and what is the smallest formalized fragment that would make it one storey of a single building rather than an isolated classical fact? *One to keep.* The only thing formalizing added beyond transcription was that the *weakest honest* base-case hypothesis dissolved a case split (Section 3)—“without loss of generality” cost nothing, but only once a machine made me look.

Cite: Question 4, p. 7, doi:10.5281/zenodo.20693064

↳ Not used by another result of the paper.

Question 5. p. 7

In a theory three centuries old and re-proved many times, how often is the WLOG genuinely free, and is “always reach for the weakest base hypothesis” a transferable principle of machine proof or a happy accident of this one?

Cite: Question 5, p. 7, doi:10.5281/zenodo.20693064

↳ Not used by another result of the paper.

Question

Reproducing

No numbered statement in this section.

It is prose only (no statement, figure or table): read it in the paper (PDF).

📊 Tables

Every table of the paper, in its order: click one to unfold it, with where the paper cites it, how its numbers were checked, and its data as CSV. The typeset table is one click away, on its page of the PDF.

Table 1. The load-bearing results, all `sorry`-free in `Newton.lean` of the repository (346 lines, on `RealStable.lean` and Mathlib). Axioms of every entry: `propext`, `Classical.choice`, `Quot.sound` only; no `native_decide`, no custom axiom.

Lean name	statement	role
`newton_inequality`	$e_k^2 \binom{n}{k-1}\binom{n}{k+1} \geq e_{k-1}e_{k+1}\binom{n}{k}^2$	Theorem 1
`newton_logConcave`	$p_{k-1}p_{k+1} \leq p_k^2$ (normalized means)	corollary
`newton_inequality_coeff`	$a_i^2 \binom{n}{i-1}\binom{n}{i+1} \geq a_{i-1}a_{i+1}\binom{n}{i}^2$	coefficient form
`realRooted_discrim_coeff`	real-rooted, $\deg \leq 2 \Rightarrow a_1^2 \geq 4a_2a_0$	base case
`RealRooted.derivative`	p real-rooted $\Rightarrow p'$ real-rooted	block (Rolle)
`RealRooted.iterate_derivative`	p real-rooted $\Rightarrow p^{(k)}$ real-rooted	block
`RealRooted.reverse`	p real-rooted $\Rightarrow p^{\mathrm{rev}}$ real-rooted	block

p. 7 in the PDF · ⬇ CSV · cited in the opening, before §1 · cited in §4 Q4: with Newton certified, what falls out for free? · cited in §5 The building blocks · cited in §7 Open questions

References

1

I. Newton, *Arithmetica Universalis* (1707); the inequalities among the coefficients of a polynomial with real roots are stated there without proof. 🔎 find it

2

C. Maclaurin, A second letter … concerning the roots of equations, with the demonstration of other rules in algebra, *Phil. Trans. Roy. Soc. London* **36** (1729) 59–96. 🔎 find it

3

G. H. Hardy, J. E. Littlewood, G. Pólya, *Inequalities*, 2nd ed., Cambridge Univ. Press, 1952 (Newton's and Maclaurin's inequalities, §2.22). 🔎 find it

4 R. P. Stanley, Log-concave and unimodal sequences in algebra, combinatorics, and geometry, *Ann. New York Acad. Sci.* **576** (1989) 500–535. 🔎 find it

5 P. Brändén, Unimodality, log-concavity, real-rootedness and beyond, in *Handbook of Enumerative Combinatorics*, CRC Press, 2015, 437–483. 🔎 find it

6 K. Adiprasito, J. Huh, E. Katz, Hodge theory for combinatorial geometries, *Ann. of Math.* **188** (2018) 381–452. 🔎 find it

7 C. Marín, *Random Signs into Matchings: A Godsil–Gutman Identity, Formalized in Lean 4*, Zenodo, 2026. doi:10.5281/zenodo.20517350.

8 C. Marín, *Unfolding a Graph into a Tree: A Machine-Checked Proof of the Heilmann–Lieb Theorem in Lean 4*, Zenodo, 2026. https://github.com/karlesmarin/godsil-gutman-lean. 🔎 find it

9 The mathlib Community, The Lean mathematical library, in *CPP 2020*, ACM, 2020, 367–381. 🔎 find it