hexagon:2611.01006v2

AN AUTOMATICALLY GENERATED EXPOSITION OF THE CASIMIR BOUND `casBound_holds` IN ALPÖGE'S LEAN PROOF OF THE RIEMANN HYPOTHESIS

LOUIS P. SLOTHOUBER

About this document

An automated process generated all of this document, including this notice. It combined mechanical work (a Lean build, dependency extraction and scripted checks) with AI work (OpenAI models through Codex and Anthropic Claude models). No human wrote, edited or checked the text of version 1. An AI model drafted version 2 after the AI review of version 1 (Section 8). No human checked version 2 either.

Louis P. Slothouber, a computer scientist, set the goal of building a tool to explain AI-generated Lean proofs, approved the plan and method choices proposed by the AI agents and decided to publish. Louis P. Slothouber decided to revise version 1 after the review. Report errors at https://github.com/lpkralv/alpoge-lean-expositions.

Abstract. Take a finite set of boxes arranged in columns and a tensor f in a tensor power of $\mathbb{C}^n$, one factor per box, that changes sign when two boxes of one column are exchanged. Let K be the quadratic Casimir operator of $\mathfrak{gl}_n$ acting on the factors in a union of whole columns, of heights $h_c \le n$. Then

$$\mathrm{Re}\langle f, Kf\rangle \le \Big(\sum_c h_c(n+1-h_c) + 2\sum_{\{c,c'\}} \min(h_c, h_{c'})\Big)\|f\|^2.$$

The coefficient is the Casimir eigenvalue on the Cartan component, the summand generated by the product of highest-weight vectors. The Lean theorem `casBound_holds` gives the equivalent bound by tableau labels and proves it with matrix units, swaps and counting. If f alternates on s boxes and Y sums the swaps of an outside box with those boxes, then $\|f\|^2 - \mathrm{Re}\langle f, Yf\rangle = \|f - Yf\|^2/(s+1)$. Summing this identity over column pairs gives the deficit in the bound. In the archive, a scalar shift brings the bound to the required automorphic Casimir scalar.

1. Introduction

1.1. **The coordinate proof in outline.** Fix an integer $n \ge 0$ and a finite set of boxes arranged in columns. Attach a copy of $\mathbb{C}^n$ to every box, and let f be a tensor in the product of these copies that changes sign whenever two boxes of the same column are exchanged. Select a union S of whole columns, write $h_c \le n$ for the height of the selected column c, and let K be the quadratic Casimir operator of $\mathfrak{gl}_n$ acting on the factors in S. The bound is

$$\mathrm{Re}\langle f, Kf\rangle \le \Big(\sum_c h_c(n+1-h_c) + 2\sum_{\{c,c'\}} \min(h_c, h_{c'})\Big)\|f\|^2, \tag{1}$$

Date: Version 2. Version 1 appeared as Zenodo 10.5281/zenodo.24514672 and Hexagon 2611.01006v1.

2020 *Mathematics Subject Classification.* Primary 17B10. Secondary 22E47, 05E10, 11M26, 68V15.

Key words and phrases. Casimir operator, column alternation, tensor representations, Young tableaux, Riemann hypothesis, Lean 4, automated mathematical exposition.

ORCID 0009-0005-4535-3306. Email lpslot@me.com.

where the first sum runs over the selected columns and the second over unordered pairs of distinct selected columns.

The box labels used in Lean can be assigned independently of f and K. A column of height h carries labels $n-h,\dots,n-1$, so two columns have one common label for every box in the shorter column. Section 3.6 derives (1) from the labelled statement.

The coefficient is the largest Casimir eigenvalue in the product of the column exterior powers. It occurs on the Cartan component, the irreducible summand generated by the tensor of highest-weight vectors [13, §1, p. 309]. Section 5.1 computes the eigenvalue.

Alpöge's Lean archive states the bound by labels on the boxes (Theorem A), under the name `CasBound`, and proves it as `casBound_holds` using matrix units, swaps and counting. In its application, the archive subtracts a scalar from each diagonal generator of the tensor action. This shift brings the bound on the coefficient representation's Casimir to the automorphic Casimir scalar. Any excess in the bound survives the shift (Section 6).

Let f alternate on a set B of s boxes, let p be a box outside B, and let Y be the sum of the swaps of p with the boxes of B. Then

$$\|f\|^2 - \mathrm{Re}\langle f, Yf\rangle = \frac{\|f - Yf\|^2}{s+1} \geq 0. \tag{2}$$

So the swaps from one box into an alternating column contribute at most $\|f\|^2$ in total, however tall the column, although a single one of them can contribute a positive amount. The Casimir is n times the number of selected boxes times the identity, plus the swaps over ordered pairs of distinct selected boxes. A swap inside a column contributes $-\|f\|^2$. For two columns, apply (2) to each box of the shorter one. This bounds the swaps between the two columns by $\min(h_c, h_{c'})\|f\|^2$ in each of the two orders. Adding up gives (1). In Theorem B (supplementary), the difference between the right and left sides of (1) is a sum of squared norms $\|f - Y_p f\|^2$ with positive weights.

1.2. **The archive and its formal status.** On 20 November 2026 Levent Alpöge released a Lean archive covering the Riemann hypothesis and its generalization to Dirichlet L-functions [3, 1]. The remark that accompanies it describes the work as joint with an AI model: "I believe I have, working with Claude, discovered a proof of the Riemann Hypothesis", and "Claude and I have also formalized the arguments without assumptions." [2]. Below, "the archive's proof" means the proof in that archive, with the authorship its remark describes.

In *Lean verification of the Riemann Hypothesis* (Lean community blog, 21 November 2026), the Lean FRO and the Mathlib Initiative certified the correctness of the formal statement and proof of Mathlib's `RiemannHypothesis` [3]. That certification concerns ordinary RH. The proof term for `riemannHypothesis` in `Main.lean` depends on `casBound_holds` through the archive's proof of `EveryCharacter`. Table 3 gives the chain of declarations.

The theorem has one caller in the archive, `ClozelCohomologicalOfThree.lean`, which uses it in the representation-theoretic input to the Weil-bound part of the proof (Section 6). Its import closure outside Mathlib consists of ten files with 975 lines. Of these, 466 lines are the five files that prove the bound, and the rest are definitions and other declarations brought in by imports. The ten files contain no `sorry`, `admit` or `axiom`. In the review of version 1 (Section 8), two AI reviewers each recompiled the ten files with Lean 4.33.1, one of them against Mathlib revision `0df444a3`, and ran `#print axioms` on `casBound_holds`. Both reported the three axioms `propext`, `Classical.choice` and `Quot.sound`.

An AI-run coverage survey on 24 November 2026 compared 38 proved theorems on the archive's route to `Main.lean` with nine write-ups of the proof [4, 5, 6, 7, 8, 9, 10, 11, 12]. It found explanations for 25 theorems and outlines for 11. A search of the text of Rodgers's

version 2 for `casBound` found no match. None of the nine copies examined explains the tableau Casimir bound.

1.3. **Reading the Lean references.** Theorem A and Propositions 1–9 state results from the Lean archive. Their proofs follow the Lean proofs. Section 7 gives the declarations and line numbers.

Supplementary marks arguments and illustrations supplied by the exposition. Section 10 lists them and records the AI checks.

2. Setting and statement

2.1. **Boxes, columns and labels.** Throughout, $n \geq 0$ is the dimension of the defining space, and $L_n = \{0, \dots, n-1\}$, which is empty for $n = 0$. The combinatorial data are a finite set Q, whose elements are called boxes or slots, a set C of column indices, a map $\gamma : Q \to C$ giving the column of each box, a map $b : Q \to L_n$ giving its label, and a set $S \subseteq Q$ of selected boxes. The set C may be infinite; all sums below are finite. The box labels b are fixed. A tensor coordinate assignment x specifies a basis index at each box. We write

$$d = |S|, \qquad S_c = \{p \in S : \gamma(p) = c\}, \quad h_c = |S_c|,$$
$$C_S = \gamma(S), \quad m_i = |\{p \in S : b(p) = i\}|.$$

So h_c is the height of the selected column c, the set C_S consists of the columns that meet S, called the active columns, and m_i is the multiplicity of the label i among the selected boxes. Two conditions on these data are used.

- The *tableau condition* $\mathcal{T}(\gamma, b)$. Distinctness: two boxes in the same column with the same label coincide. Upward closure: for every $p \in Q$ and $i \in L_n$ with $b(p) < i$ there is $q \in Q$ with $\gamma(q) = \gamma(p)$ and $b(q) = i$.
- *Whole-column selection* $\mathcal{U}(\gamma, S)$. For all $p, q \in Q$ with $\gamma(p) = \gamma(q)$, the box p is in S if and only if q is.

Observation (supplementary). Assume $\mathcal{T}(\gamma, b)$ and $\mathcal{U}(\gamma, S)$. Then for every active column c the labels of the boxes in S_c are exactly $n - h_c, \dots, n-1$, each carried by one box. In particular $h_c \leq n$, two active columns $c \neq c'$ have exactly $\min(h_c, h_{c'})$ labels in common, and m_i is the number of active columns of height at least $n - i$.

Proof. Let ℓ_c be the least label in S_c. Distinctness makes b injective on S_c. Upward closure gives, for each i with $\ell_c < i \leq n-1$, a box of column c with label i, and whole-column selection puts that box in S. So the labels of S_c are exactly $\ell_c, \dots, n-1$ and $h_c = n - \ell_c$. The other claims follow, since the label i occurs in S_c exactly when $n - h_c \leq i$. □

Propositions 4 and 7 use distinctness, upward closure and whole-column selection directly. By the Observation, the data on S can also be read by heights. Read the m_i from $i = n-1$ downward to obtain the partition conjugate to the sorted heights (supplementary) [15, ch. I, §1, (1.3)].

2.2. **Tensors, swaps and the Casimir.** Let L_n^Q be the set of assignments $x : Q \to L_n$, and let $\mathcal{V} = \mathbb{C}^{L_n^Q}$ be the space of complex functions on it. This is the tensor power of $\mathbb{C}^n$ with one factor for each box, written in coordinates. The assignment x varies and the label map b stays fixed. The inner product is $\langle f, g \rangle = \sum_x \overline{f(x)}\, g(x)$, conjugate-linear in the first argument, with norm $\|f\|$.

For $p, q \in Q$ the *swap* P_{pq} is defined by $(P_{pq} f)(x) = f(x \circ (p\,q))$, where $(p\,q)$ is the transposition of the two slots, and $P_{pp} = \mathrm{Id}$. Each swap permutes the coordinate basis, so it is a unitary

label	c_1	c_2	c_3	c_4	multiplicity
0					$m_0 = 0$
1	1				$m_1 = 1$
2	2	2	2		$m_2 = 3$
3	3	3	3	3	$m_3 = 4$
height	3	2	2	1	

FIGURE 1. Four selected columns at rank $n = 4$, with heights $3, 2, 2, 1$. Each box shows its label. A column of height h carries the labels $n - h, \dots, n - 1$, so the columns are aligned at the label $n - 1$, and m_i is the number of boxes in the row of label i. Reading upward, the rows have lengths $4, 3, 1$, with row index $j = n - i$. Columns c_1, c_4 share the label 3, giving one match in each order. Here the coefficient of Theorem A is $0 + 0 + 12 + 28 = 40$ by labels, and $(6 + 6 + 6 + 4) + 2 \cdot 9 = 40$ by heights as in Theorem B. This illustration is supplementary.

involution and hence self-adjoint, and $P_{pq} = P_{qp}$. For two columns $c \neq c'$ we write $X_{cc'} = \sum_{p \in S_c} \sum_{q \in S_{c'}} P_{pq}$ for the sum of the swaps between them.

For $i, i' \in L_n$ the *matrix unit in slot* p is

$$(E_{ii'}^{(p)} f)(x) = \begin{cases} f(x[p \leftarrow i']), & x(p) = i, \\ 0, & x(p) \neq i, \end{cases}$$

where $x[p \leftarrow i']$ is the assignment obtained from x by replacing its value at p by i'. It sends the basis vector $e_{i'}$ to e_i in slot p and is the identity in the other slots. The matrix units of $\mathfrak{gl}_n$ act on the factors in S by $G_{ii'} = \sum_{p \in S} E_{ii'}^{(p)}$, and the Casimir is

$$K = \sum_{i, i' \in L_n} G_{ii'} G_{i'i}.$$

This is the quadratic Casimir of $\mathfrak{gl}_n$ for the trace form, with the scalar direction included (supplementary) [25, Example 5.4]. If $\mathcal{V}'$ is the tensor space of the unselected slots, then $K = K_S \otimes I_{\mathcal{V}'}$, where K_S acts on the selected slots. The Lean statement uses the coefficient module $\mathbb{C}$ and sets its term in each generator to zero.

2.3. **The theorem.**

Theorem A (Tableau Casimir bound; in the Lean as `casBound_holds`)**.**

For every $n \in \mathbb{Z}_{\geq 0}$, *finite set* Q, *set* C, $\gamma : Q \to C$, $b : Q \to L_n$, $S \subseteq Q$, $f \in \mathcal{V}$, *assume*

$$\begin{aligned} &\forall p, q \in Q, \quad \gamma(p) = \gamma(q) \ \wedge \ b(p) = b(q) \implies p = q, \\ &\forall p \in Q, \ \forall i \in L_n, \quad b(p) < i \implies \exists q \in Q : \ \gamma(q) = \gamma(p), \ b(q) = i, \\ &\forall p, q \in Q, \quad \gamma(p) = \gamma(q) \implies (p \in S \iff q \in S), \\ &\forall p, q \in Q, \quad \gamma(p) = \gamma(q) \ \wedge \ p \neq q \implies P_{pq} f = -f. \end{aligned}$$

Then

$$\operatorname{Re}\langle f, Kf \rangle \leq \left(\sum_{i \in L_n} m_i (m_i + 2i + 1 - n) \right) \|f\|^2.$$

The first two hypotheses are the tableau condition $\mathcal{T}(\gamma, b)$, the third is whole-column selection $\mathcal{U}(\gamma, S)$, and the fourth says that f alternates within every column. This is the statement of the Lean theorem, written in the notation above. The following comments are supplementary.

- The operator K is self-adjoint, because the adjoint of $G_{ii'}$ is $G_{i'i}$. So $\langle f, Kf\rangle$ is real, and the real part is there to match the Lean statement.
- The hypotheses are stated on all of Q. The proof uses the tableau condition and the alternation only for boxes of S.
- For $n = 0$ the label map forces $Q = \varnothing$. If $S = \varnothing$ or $f = 0$, both sides are zero.

2.4. **The form by heights, and the bound as a sum of squares.**

Theorem B (Heights form and sum of squares; supplementary)**.** *Assume the hypotheses of Theorem A.*

(a) *The coefficient can be written by heights:*

$$\sum_{i \in L_n} m_i(m_i + 2i + 1 - n) = \sum_{c \in C_S} h_c(n + 1 - h_c) + 2 \sum_{\{c,c'\} \subseteq C_S} \min(h_c, h_{c'}),$$

where the second sum is over two-element subsets.

(b) *For each two-element subset $\{c, c'\}$ of C_S name the two columns so that $h_c \le h_{c'}$, with either naming in a tie, and for $p \in S_c$ put $Y_p = \sum_{q \in S_{c'}} P_{pq}$ for that pair. Then*

$$\Big(\sum_{i \in L_n} m_i(m_i + 2i + 1 - n)\Big)\|f\|^2 - \mathrm{Re}\langle f, Kf\rangle = 2 \sum_{\{c,c'\} \subseteq C_S} \frac{1}{h_{c'} + 1} \sum_{p \in S_c} \|f - Y_p f\|^2.$$

(c) *Equality holds in Theorem A exactly when $Y_p f = f$ for every such pair and every $p \in S_c$.*

In part (a), a column of height h contributes nh from the diagonal part of K and $-h(h-1)$ from the swaps inside it, and an unordered pair of columns contributes at most twice the smaller height. Part (b) writes the deficit as a sum of squares. Each vector $(f - Y_p f)/(h_{c'} + 1)$ is the antisymmetrization of f over the taller column together with one box p of the shorter one (Section 3.2). The Lean proves the inequality of Theorem A and counts by labels throughout; heights do not appear in it. Theorem B is proved in Section 3.6 from the steps of the Lean proof.

Running example (supplementary). Take $n = 2$ and three boxes, all selected: p_1, p_2 form a column with labels $0, 1$, and p_3 is a column by itself with label 1. The heights are 2 and 1, and $m_0 = 1$, $m_1 = 2$. The coefficient is $1 \cdot 0 + 2 \cdot 3 = 6$ by labels and $2 + 2 + 2 \cdot 1 = 6$ by heights. Write $|x_1 x_2 x_3\rangle$ for the coordinate basis tensor at the assignment $p_j \mapsto x_j$, and take

$$f = \frac{|010\rangle - |100\rangle}{\sqrt{2}},$$

which alternates in p_1, p_2 and has norm 1. At dimension two, every tensor alternating in p_1, p_2 attains the bound, since $\bigwedge^3 \mathbb{C}^2 = 0$ makes the enlarged antisymmetrization vanish.

Keep the two columns and take $n = 3$. The labels are now $1, 2$ in the first column and 2 in the second. Take the unit alternating tensor

$$u = \frac{|012\rangle + |120\rangle + |201\rangle - |021\rangle - |102\rangle - |210\rangle}{\sqrt{6}}.$$

Every swap sends u to $-u$. By Proposition 2, the Casimir expectation is $3 \cdot 3 + 2(-1-1-1) = 3$. The height coefficient is $2(4-2)+1(4-1)+2\cdot 1 = 9$, also $1\cdot 1 + 2 \cdot 4 = 9$ by labels. For p_3, the only box in the shorter column, $Y_{p_3} u = -2u$. The one-order deficit is $\|u - Y_{p_3} u\|^2/(2+1) = 9/3 = 3$. Theorem B(b) counts both orders, giving the deficit $2 \cdot 3 = 6 = 9 - 3$.

3. The coordinate proof

First expand the Casimir into swaps. Within-column swaps give exactly $-\|f\|^2$ per ordered pair. To bound cross-column swaps, sum the one-box estimate from the shorter column, counting one match per common label. The final step collects the diagonal, within-column and cross-column terms by label.

3.1. **Expanding the Casimir.**

Proposition 1 (Contraction of two slot matrix units; Lean: `pair_apply`)**.**

$$\text{For every } n \in \mathbb{Z}_{\geq 0},\ \text{finite set } Q,\ p, q \in Q,\ f \in \mathcal{V},\ x \in L_n^Q,$$

$$\sum_{i,i' \in L_n} (E_{ii'}^{(p)} E_{i'i}^{(q)} f)(x) = \begin{cases} nf(x), & p = q, \\ f(x \circ (p\,q)), & p \neq q. \end{cases}$$

Proof. By the definition of the matrix units, the outer operator forces $i = x(p)$ and reads the intermediate assignment $x[p \leftarrow i']$. The inner operator contributes only when that intermediate assignment has value i' at q, and then replaces the value at q by i.

When $p = q$ the first replacement makes the inner condition automatic for all n values of i', and the second restores $i = x(p)$, giving $nf(x)$. When $p \neq q$ the first replacement leaves $x(q)$ intact, so the inner condition forces $i' = x(q)$ and exactly one term survives, also when $x(p) = x(q)$. Its replacements put the original value of each of the two slots in the other and change nothing elsewhere. □

Proposition 2 (Slot expansion of the Casimir; Lean: `cas_expand'`)**.**

$$\text{For every } n \in \mathbb{Z}_{\geq 0},\ \text{finite set } Q,\ S \subseteq Q,$$

$$K = nd\,\mathrm{Id} + \sum_{p \in S} \sum_{\substack{q \in S \\ q \neq p}} P_{pq}.$$

Proof. Distribute the two sums over S in each product $G_{ii'}G_{i'i}$ and group the contraction over i, i' by the ordered pair of slots (p, q). Proposition 1 gives $n\,\mathrm{Id}$ for each of the d pairs with $p = q$ and P_{pq} for each of the $d(d-1)$ pairs with $p \neq q$. □

Each unordered pair of selected boxes contributes its swap twice. In the two-dimensional example, $K = 6\,\mathrm{Id} + 2(P_{p_1p_2} + P_{p_1p_3} + P_{p_2p_3})$.

Remarks (supplementary). The expansion sets the coefficient-module term to zero. In the application a scalar shift contributes mixed and quadratic terms, which are computed separately in Section 6.

In the last display of the proof of Proposition 2.2, Lévy sums over ordered pairs of distinct slots. His Laplacian has the opposite sign to the Casimir used here [18, proof of Proposition 2.2, p. 6 of the preprint].

3.2. **One box against an alternating set.**

Let $B \subseteq Q$ have s elements, let $p \in Q \setminus B$, and put $Y = \sum_{q \in B} P_{pq}$. Suppose f alternates on B, that is, $P_{qq'}f = -f$ for distinct $q, q' \in B$. Orthogonal projection onto tensors alternating on the enlarged set gives the squared norm in (2) (supplementary).

Let A be the orthogonal projector onto the tensors that alternate on the enlarged set $B \cup \{p\}$. It is the signed average of the slot permutations of $B \cup \{p\}$. Those $(s+1)!$ permutations fall into $s+1$ cosets of the subgroup permuting B, represented by the identity and the s transpositions

$(p\,q)$ with $q \in B$, which have sign -1. Because f already alternates on B, each coset contributes one term, and

$$Af = \frac{f - Yf}{s+1}.$$

For an orthogonal projector, $\langle f, Af\rangle = \|Af\|^2 \geq 0$. Written out, this is (2): $\|f\|^2 - \mathrm{Re}\langle f, Yf\rangle = \|f - Yf\|^2/(s+1)$. So the defect $f - Yf$ is, up to the factor $s+1$, the part of f that alternates on the column together with the new box, and the divisor $s+1$ is the number of cosets. In the theory of the symmetric group this is the one-column case of the factorization of Young projectors [19, (3)] and of the branching formula for them [22, §2, (7)–(9)].

Proposition 3 states the inequality. Its Lean proof obtains the identity (2) as an intermediate calculation by multiplying swaps (`ReSwapSumLe.lean`, lines 65–68 and 80–81).

Proposition 3 (Swaps from one outside box into an alternating set; Lean: `re_swapSum_le'`)**.**

$$\textit{For every } n \in \mathbb{Z}_{\geq 0},\ \textit{finite set } Q,\ B \subseteq Q,\ p \in Q \setminus B,\ f \in \mathcal{V},$$
$$(\forall q, q' \in B,\quad q \neq q' \implies P_{qq'} f = -f) \implies \sum_{q \in B} \mathrm{Re}\langle f, P_{pq} f\rangle \leq \|f\|^2.$$

Proof. Set $s = |B|$, $Y = \sum_{q\in B} P_{pq}$ and $g = f - Yf$. For distinct $q, q' \in B$ the three slots p, q, q' are distinct, because $p \notin B$, and the transposition relation on three distinct slots gives $P_{pq}P_{pq'} = P_{pq'}P_{qq'}$. This places the swap inside B next to f, so $P_{pq}P_{pq'}f = -P_{pq'}f$. Also $P_{pq}P_{pq}f = f$. Hence, for each $q \in B$,

$$P_{pq}(Yf) = f - \sum_{q' \in B,\ q' \neq q} P_{pq'} f = f - Yf + P_{pq} f.$$

Subtracting this equality from $P_{pq}f$ gives $P_{pq}g = -g$, so g alternates on the enlarged set, as in the projector calculation. Summing over $q \in B$, the s diagonal products give sf, and each q' appears $s-1$ times in the off-diagonal part, so

$$Y^2 f = sf - (s-1)Yf.$$

It follows that $Yg = Yf - Y^2 f = -s(f - Yf) = -sg$. Each term of Y is a self-adjoint involution, so Y is self-adjoint and

$$\|g\|^2 = \mathrm{Re}\langle (\mathrm{Id} - Y)f, g\rangle = \mathrm{Re}\langle f, (\mathrm{Id} - Y)g\rangle = (s+1)\,\mathrm{Re}\langle f, g\rangle,$$
$$\|f\|^2 - \sum_{q \in B} \mathrm{Re}\langle f, P_{pq} f\rangle = \frac{\|f - Yf\|^2}{s+1} \geq 0.$$

The division is valid for every $s \geq 0$. At $s = 0$ the sum is empty, $g = f$ and the divisor is 1. □

The proposition uses alternation only on the target set B, and B can be any set not containing p. In the proof of Theorem A it is applied with $B = S_c$ for a column c other than the column of p. In the two-dimensional example take $p = p_3$ and $B = \{p_1, p_2\}$. The swap $P_{p_1p_3}$ fixes $|010\rangle$ and sends $|100\rangle$ to $|001\rangle$, and $P_{p_2p_3}$ fixes $|100\rangle$ and sends $|010\rangle$ to $|001\rangle$. So each of the two swaps has expectation $\mathrm{Re}\langle f, Pf\rangle = 1/2$, their sum is $1 = \|f\|^2$, and $Yf = f$: the example is an equality case.

Remarks (supplementary). By the displayed identity, equality holds in Proposition 3 exactly when $Yf = f$. At the other extreme, on a tensor that alternates on all of $B \cup \{p\}$ one has $Yf = -sf$ and $g = (s+1)f$.

The operator Y commutes with every swap inside B, so it preserves the space of tensors alternating on B. On that space the identity $Y^2 = s - (s-1)Y$ factors as $(Y - \mathrm{Id})(Y + s\,\mathrm{Id}) = 0$, so Y is self-adjoint with spectrum in $\{1, -s\}$. For $s = 0$, $Y = 0$, so the polynomial relation

and the spectral inclusion still hold. This gives $Y \le \mathrm{Id}$ for all s. Equivalently, $(\mathrm{Id} - Y)^2 f = (s+1)(\mathrm{Id} - Y)f$.

On tensors alternating on all of $B \cup \{p\}$, $Yf = -sf$. For $s > 1$ and $f \neq 0$, the norm $s\|f\|$ exceeds $\|f\|$, while the quadratic form is $-s\|f\|^2$.

3.3. **Counting labels in one column and across columns.** Write $N = \|f\|^2$. The diagonal contributes nN per selected box. Counting both orders at the lower label gives the within-column term $-2(n-1-i)N$ at a box of label i. The cross-column estimate counts one match for each of the $m_i - 1$ other boxes of that label, after the swaps have been grouped by column pair.

Proposition 4 (The number of higher labels in a selected column; Lean: `count_up`)**.**

$$\text{For every } n \in \mathbb{Z}_{\ge 0},\ \text{finite set } Q,\ \text{set } C,\ \gamma : Q \to C,\ b : Q \to L_n,\ S \subseteq Q,\ p \in S,$$
$$\mathcal{T}(\gamma, b)\ \wedge\ \mathcal{U}(\gamma, S)\ \Longrightarrow\ |\{q \in S : \gamma(q) = \gamma(p),\ b(q) > b(p)\}| = n - 1 - b(p).$$

Proof. Map each selected box in the column of p whose label exceeds $b(p)$ to its label. The map is injective by distinctness. Upward closure supplies a box of that column for every label in L_n above $b(p)$, and whole-column selection puts each such box in S. The map is therefore a bijection onto the $n - 1 - b(p)$ labels $b(p) + 1, \dots, n - 1$. □

Remarks (supplementary). Injectivity alone gives only an upper bound for this count. For the Casimir estimate that is the wrong direction, because the count is multiplied by $-2\|f\|^2$. Upward closure and whole-column selection are what make the count exact. Section 4 gives an example where a partly selected column breaks the theorem.

Proposition 5 (The number of equal labels in other columns; Lean: `count_eqb`)**.**

$$\text{For every } n \in \mathbb{Z}_{\ge 0},\ \text{finite set } Q,\ \text{set } C,\ \gamma : Q \to C,\ b : Q \to L_n,\ S \subseteq Q,\ p \in S,$$
$$\mathcal{T}(\gamma, b)\ \Longrightarrow\ m_{b(p)} = 1 + |\{q \in S : \gamma(q) \neq \gamma(p),\ b(q) = b(p)\}|.$$

Proof. Partition the selected boxes of label $b(p)$ into those in the column of p and those in other columns. The first part contains p, because $p \in S$, and distinctness makes it the singleton $\{p\}$. □

The Lean statement assumes the full tableau condition, as here, although the proof uses only distinctness, and S need not be a union of whole columns in this proposition.

3.4. **Two columns.** The shorter column's labels are contained in the taller column's labels. Each box in the shorter column has one match and supplies one application of the one-box estimate. Proposition 7 finds that column directly from the tableau hypotheses. Proposition 8 compares sums over columns, where swaps of differently labelled boxes can contribute.

Proposition 6 partitions a sum for an arbitrary function F. Propositions 7 and 8 use a real symmetric kernel R on pairs of boxes and allow N of either sign, as in Lean. In the proof of Theorem A they are applied with $R(p,q) = \mathrm{Re}\langle f, P_{pq} f\rangle$ and $N = \|f\|^2$. The hypothesis

$$\forall p \in S,\ \forall c \in C \setminus \{\gamma(p)\}, \qquad \sum_{q \in S_c} R(p,q) \le N$$

is called the *one-box bound*: it is what Proposition 3 gives for a box p and a column c other than its own.

Proposition 6 (Partition of cross-column pairs; Lean: `sum_sum_ite_col_ne_eq_sum_by_col`, private)**.**

$$\text{For every finite set } Q,\ \text{set } C,\ \gamma : Q \to C,\ S \subseteq Q,\ F : Q \times Q \to \mathbb{R},$$
$$\sum_{p \in S} \sum_{\substack{q \in S \\ \gamma(q) \neq \gamma(p)}} F(p,q) = \sum_{c \in C_S} \sum_{\substack{c' \in C_S \\ c' \neq c}} \sum_{p \in S_c} \sum_{q \in S_{c'}} F(p,q).$$

Proof. The products $S_c \times S_{c'}$ with $c \neq c'$ in C_S partition the ordered pairs of selected boxes in different columns: the pair (p,q) lies in the product with $c = \gamma(p)$ and $c' = \gamma(q)$ and in no other. □

The proposition is applied twice, to $F = R$ and to $F(p,q) = N\mathbf{1}_{b(p)=b(q)}$, where $\mathbf{1}$ is the indicator of the stated condition. The two sums are then compared on the same ordered pairs of columns.

Proposition 7 (Comparison for two distinct columns; Lean: `sum_filter_col_le_sum_ite_of_bound`, private)**.**

$$\text{For every } n \in \mathbb{Z}_{\geq 0},\ \text{finite set } Q,\ \text{set } C,\ \gamma : Q \to C,\ b : Q \to L_n,\ S \subseteq Q,$$
$$R : Q \times Q \to \mathbb{R},\ N \in \mathbb{R},$$
$$\text{assume } \mathcal{T}(\gamma, b), \quad \mathcal{U}(\gamma, S), \quad \forall p, q \in Q, \quad R(p,q) = R(q,p),$$
$$\forall p \in S,\ \forall c \in C \setminus \{\gamma(p)\}, \qquad \sum_{q \in S_c} R(p,q) \leq N.$$
$$\text{Then for every } c, c' \in C \text{ with } c \neq c',$$
$$\sum_{p \in S_c} \sum_{q \in S_{c'}} R(p,q) \leq \sum_{p \in S_c} \sum_{q \in S_{c'}} N\mathbf{1}_{b(p)=b(q)}.$$

Proof. Say that a box $p \in S_c$ has a *match* in $S_{c'}$ if some $q \in S_{c'}$ has $b(q) = b(p)$. By distinctness a match is unique.

First suppose every box of S_c has a match in $S_{c'}$. For each $p \in S_c$ the inner sum on the right is then exactly N, and the one-box bound for p and the column c' says that the inner sum on the left is at most N. Summing over $p \in S_c$ gives the claim.

Otherwise take $p_* \in S_c$ whose label $i_* = b(p_*)$ does not occur in $S_{c'}$. Let $q \in S_{c'}$. Its label is not i_*. If $b(q) < i_*$, upward closure from q supplies a box of label i_* in column c', and whole-column selection puts that box in $S_{c'}$, a contradiction. Therefore $b(q) > i_*$, and upward closure from p_* together with whole-column selection gives q a match in S_c. Now interchange the two sums and use $R(p,q) = R(q,p)$. For each $q \in S_{c'}$, the one-box bound for q and the column c bounds $\sum_{p \in S_c} R(p,q)$ by N, and the unique match of q gives exactly N on the right. Summing over $q \in S_{c'}$ gives the claim. □

The statement and the proof are for every real N, because the proof compares exact counts of matches. In the application $N = \|f\|^2 \geq 0$.

Remarks (supplementary). In the second case of the proof any label i_* with $\ell_c \leq i_* < \ell_{c'}$ works. Here ℓ_a is the least label in S_a when S_a is nonempty and is n when S_a is empty. The argument uses only the presence of i_* in S_c and its absence from $S_{c'}$.

The statement allows empty columns. If S_c or $S_{c'}$ is empty, both double sums vanish. If $p \in S$ and some column $c \neq \gamma(p)$ has $S_c = \varnothing$, the one-box bound for p and c reads $0 \leq N$. Proposition 8 uses only columns in C_S.

Proposition 8 (Cross-column comparison by equal labels; Lean: `count_cross`)**.**

$$\begin{gathered}\textit{For every } n \in \mathbb{Z}_{\geq 0},\ \textit{finite set } Q,\ \textit{set } C,\ \gamma : Q \to C,\ b : Q \to L_n,\ S \subseteq Q,\\ R : Q \times Q \to \mathbb{R},\ N \in \mathbb{R},\\ \textit{assume } \mathcal{T}(\gamma, b), \quad \mathcal{U}(\gamma, S), \quad \forall p, q \in Q, \quad R(p,q) = R(q,p),\\ \forall p \in S,\ \forall c \in C \setminus \{\gamma(p)\}, \qquad \sum_{q \in S_c} R(p,q) \leq N.\\ \textit{Then} \qquad \sum_{p \in S} \sum_{\substack{q \in S \\ \gamma(q) \neq \gamma(p)}} R(p,q) \leq \sum_{p \in S} \sum_{\substack{q \in S \\ \gamma(q) \neq \gamma(p),\ b(q) = b(p)}} N.\end{gathered}$$

Proof. Apply Proposition 6 to both sides, with $F = R$ on the left and $F(p,q) = N\mathbf{1}_{b(p)=b(q)}$ on the right. Each side then runs over the same ordered pairs of distinct active columns. Proposition 7 gives the inequality for every such pair, and addition proves the result. □

Remarks (supplementary). In the two-dimensional example the boxes p_1 and p_3 have the unequal labels 0 and 1, yet the swap $P_{p_1p_3}$ has expectation $1/2 > 0$. The swap $P_{p_2p_3}$ of the two boxes of label 1 also has expectation $1/2$. The total $1/2 + 1/2 = 1$ over the column is exactly N times the one equal-label match. The label b is fixed in the count, while the coordinate assignment x varies in the inner product.

To count the bound by label, orient each unordered pair of active columns from the shorter column, with either choice in a tie. Each label i of the source column occurs once there and once in the target column. The one-box bound for its box contributes at most N, and symmetry lets the same bound serve for both orders. Exactly $\binom{m_i}{2}$ unordered pairs of columns contain the label i in both columns, so the total attached to the label i is $2\binom{m_i}{2}N = m_i(m_i - 1)N$.

3.5. **Assembly.**

Proposition 9 (The tableau count for a symmetric kernel; Lean: `cas_count`)**.**

$$\begin{gathered}\textit{For every } n \in \mathbb{Z}_{\geq 0},\ \textit{finite set } Q,\ \textit{set } C,\ \gamma : Q \to C,\ b : Q \to L_n,\ S \subseteq Q,\\ R : Q \times Q \to \mathbb{R},\ N \in \mathbb{R},\\ \textit{assume } \mathcal{T}(\gamma, b), \quad \mathcal{U}(\gamma, S), \quad \forall p, q \in Q, \quad R(p,q) = R(q,p),\\ \forall p, q \in Q, \quad \gamma(p) = \gamma(q) \ \wedge\ p \neq q \implies R(p,q) = -N,\\ \forall p \in S,\ \forall c \in C \setminus \{\gamma(p)\}, \qquad \sum_{q \in S_c} R(p,q) \leq N.\\ \textit{Then} \qquad ndN + \sum_{p \in S} \sum_{\substack{q \in S \\ q \neq p}} R(p,q) \leq \left(\sum_{i \in L_n} m_i(m_i + 2i + 1 - n) \right) N.\end{gathered}$$

Proof. Split the ordered pairs of distinct selected boxes into same-column pairs and cross-column pairs.

Within a column the hypothesis gives $R(p,q) = -N$. Each unordered pair of distinct boxes in a column has exactly one ordering with increasing label, by distinctness. Count both of its orders at the box with the lower label. Proposition 4 says that p has $n - 1 - b(p)$ partners of higher label, so the same-column pairs contribute exactly

$$-2N \sum_{p \in S} (n - 1 - b(p)).$$

For the cross-column pairs, Proposition 8 bounds the sum by one N for each equal-label match across columns, and Proposition 5 says that p has $m_{b(p)} - 1$ such matches. The three coefficients

at label i, in units of N, are

$$n - 2(n-1-i) + (m_i - 1) = m_i + 2i + 1 - n.$$

These terms account for the total after the column-pair comparison. Grouping the m_i boxes of each label i gives the right side. □

Remarks (supplementary). By heights, a column of height h_c contributes nh_cN from the diagonal and exactly $-h_c(h_c-1)N$ from its internal pairs, together $h_c(n+1-h_c)N$. A single active column has no cross-column pairs, so its coefficient is exactly $nh_c - h_c(h_c-1)$, and at $h_c = n$ this is n.

Proof of Theorem A. Set $N = \|f\|^2$ and $R(p,q) = \mathrm{Re}\langle f, P_{pq} f\rangle$. By Proposition 2, and because the real part of the inner product is real-linear in its second argument,

$$\mathrm{Re}\langle f, Kf\rangle = ndN + \sum_{p\in S} \sum_{\substack{q\in S \\ q\neq p}} R(p,q).$$

The kernel R is symmetric because $P_{pq} = P_{qp}$. The alternation hypothesis gives $R(p,q) = -N$ for distinct boxes in one column. For $p \in S$ and a column $c \neq \gamma(p)$, the tensor f alternates on S_c and $p \notin S_c$, so Proposition 3 with $B = S_c$ gives $\sum_{q\in S_c} R(p,q) \leq N$. For an empty S_c this reads $0 \leq N$. These are the hypotheses of Proposition 9, and its conclusion is the theorem. □

In the two-dimensional example $N = 1$. The diagonal term is $nd = 6$, the two ordered pairs inside the first column give -2, and the four ordered cross-column pairs give $2 \cdot (1/2 + 1/2) = 2$. So $\mathrm{Re}\langle f, Kf\rangle = 6 - 2 + 2 = 6$, the coefficient, and f attains the bound.

3.6. **Proof of Theorem B.** This subsection is supplementary.

Part (a). Split the coefficient into $\sum_i m_i[n - 2(n-1-i)]$ and $\sum_i m_i(m_i-1)$. For the first sum a column of height h contributes $nh - 2(0 + \cdots + (h-1)) = h(n+1-h)$. The second sum counts ordered equal-label matches. Distinctness puts the matched boxes in different columns. By the Observation,

$$\sum_i m_i(m_i-1) = 2 \sum_{\{c,c'\}\subseteq C_S} \min(h_c, h_{c'}).$$

Adding the two sums proves (a).

Part (b). Let $c \neq c'$ be active columns. The tensor f alternates on $S_{c'}$, and each $p \in S_c$ lies outside $S_{c'}$, so the identity at the end of the proof of Proposition 3 applies with $B = S_{c'}$ and $s = h_{c'}$. Summing it over $p \in S_c$ gives the *pair-deficit identity*

$$h_c\|f\|^2 - \mathrm{Re}\langle f, X_{cc'} f\rangle = \frac{1}{h_{c'}+1} \sum_{p\in S_c} \|f - Y_p f\|^2, \qquad Y_p = \sum_{q\in S_{c'}} P_{pq}. \tag{3}$$

Use either column as the source c. Partition the ordered pairs in Proposition 2 by column. Each within-column swap contributes $-\|f\|^2$, so

$$\mathrm{Re}\langle f, Kf\rangle = \Big(nd - \sum_{c\in C_S} h_c(h_c-1)\Big)\|f\|^2 + \sum_{c\in C_S} \sum_{\substack{c'\in C_S \\ c'\neq c}} \mathrm{Re}\langle f, X_{cc'} f\rangle,$$

Since $X_{cc'} = X_{c'c}$, each unordered pair occurs twice. Substitute (3) with $h_c \leq h_{c'}$ for both occurrences. Since $nd - \sum_c h_c(h_c-1) = \sum_c h_c(n+1-h_c)$, part (a) turns the result into part (b).

Part (c). The right side of (b) is a sum of nonnegative terms.

Derivation of (1). If any column has height greater than n, alternation forces $f = 0$. Otherwise give the boxes in each column of height h the labels $n-h, \ldots, n-1$, one per box. This

satisfies the tableau condition and leaves f and K unchanged. Apply Theorem A and use part (a) for its coefficient.

4. Why simpler estimates fail

These comparisons are supplementary. Throughout, $N = \|f\|^2$.

4.1. **The column signs must be kept.** For $n > 0$ the sum of the $d(d-1)$ ordered swaps of selected boxes has largest eigenvalue $d(d-1)$ on the whole tensor space, attained on fully symmetric tensors. With the diagonal term, the largest eigenvalue of K on the whole space is $nd + d(d-1)$. With a single active column of height $d \le n$, every swap of distinct selected boxes acts on f by -1 and the exact coefficient is $nd - d(d-1)$. An estimate that ignores the alternation therefore has excess $2d(d-1)$. At $n = d = 2$ the two coefficients are 6 and 2, an excess of 4.

Alternation within columns does not exclude all symmetric tensors. At $n = 2$ with two columns of one box each, Proposition 2 gives $K = 4\,\mathrm{Id} + 2P_{p_1p_2}$, where p_1, p_2 are the two boxes. The alternation hypothesis is empty for such columns, so both the symmetric and the alternating tensors are admissible, with eigenvalues 6 and 2. The coefficient of Theorem A is 6 here. So the theorem has to bound every admissible tensor, and computing K on one distinguished vector does not do that.

4.2. **The swaps must be grouped by column.** Each swap is unitary, so Cauchy–Schwarz bounds each $\mathrm{Re}\langle f, P_{pq} f\rangle$ by N. For one box and a target set of s boxes this gives sN, which agrees with Proposition 3 at $s = 1$ and loses $(s-1)N$ for $s > 1$. For two columns of heights h_c and $h_{c'}$, the separate estimates total $2h_c h_{c'} N$ over the two orders. Applying the one-box bound in each order as it stands gives $(h_c + h_{c'})N$. Using the symmetry $P_{pq} = P_{qp}$ to apply it from the shorter column in both orders gives $2\min(h_c, h_{c'})N$. At equal heights the first and last are $2h_c^2 N$ and $2h_c N$. At heights 3 and 1 the last two are $4N$ and $2N$.

For two full columns at $n = 2$, the diagonal gives $8N$ and the pairs inside the columns give $-4N$. The separate estimates add $8N$, for a total of $12N$. The grouped bound adds $4N$, for a total of $8N$, which is the coefficient of Theorem A.

4.3. **Whole columns are needed.** At $n = 2$ take a column with labels $0, 1$ and select only its box of label 0. The selected box has no selected box above it, although $n - 1 - b(p) = 1$, so the count of Proposition 4 fails. The theorem fails with it. The tensor $e_0 \otimes e_1 - e_1 \otimes e_0$ alternates on the column. By Proposition 2 the Casimir of the one selected box is $2\,\mathrm{Id}$. The multiplicities $m_0 = 1$, $m_1 = 0$ give the coefficient 0, and the bound would read $2\|f\|^2 \le 0$.

The height-one coefficient is $1 \cdot (2 + 1 - 1) = 2$. The retained label is 0, whereas a tableau column of height one at dimension two carries label 1. If each selected subcolumn is closed upward in the labels, Theorem A still applies after reassigning the unselected boxes. Put each unselected box in a new column of its own and give it label $n - 1$. The new data meet the tableau, selection and alternation hypotheses. The selected slots, K and the m_i are unchanged.

5. The classical proofs

The following proofs and comparisons are supplementary. The literature search did not find the archive's coordinate proof in the sources cited here.

5.1. **The highest-weight proof.** For each active column c, normalized alternation identifies $\bigwedge^{h_c}\mathbb{C}^n$ with the tensors on the slots of S_c that alternate. A basis wedge of height h maps to the signed sum of its $h!$ orderings divided by $\sqrt{h!}$. The orderings are distinct orthonormal basis tensors. So every tensor of Theorem A lies in

$$\Big(\bigotimes_{c\in C_S}\bigwedge^{h_c}\mathbb{C}^n\Big)\otimes\mathcal{V}',$$

where $\mathcal{V}'$ is the tensor space of the unselected slots, and K acts on the first factor and as the identity on $\mathcal{V}'$.

For the Schur-module construction from column exterior powers and exchange relations, see Fulton, §8.1, especially (4) [16]. For the Young-symmetrizer construction of Schur functors, see Fulton–Harris, §6.1 [17].

The tensor product is a finite-dimensional unitary representation of $U(n)$, so it decomposes as an orthogonal direct sum of irreducibles (Etingof, Proposition 11.13 [23]). The operator K is the Casimir of the complexified Lie algebra action. On a constituent of highest weight $\mu=(\mu_1\ge\cdots\ge\mu_n)$ it acts by the scalar

$$\Gamma_n(\mu)=\sum_{j=1}^{n}\mu_j(\mu_j+n+1-2j).$$

Etingof gives highest-weight theory and the Casimir scalar $(\mu,\mu+2\rho)$ on a highest-weight representation of a semisimple Lie algebra [23, lectures 25–26, Lemma 26.6]. Yang computes the Harish-Chandra projection of $\sum_{i,j}E_{ij}E_{ji}$ as $\sum_i E_{ii}^2+\sum_{i<j}(E_{ii}-E_{jj})$, equivalently $\sum_i(E_{ii}^2+(n+1-2i)E_{ii})$, which is $\Gamma_n(\mu)$ at the highest weight [25, Example 5.4]. Perelomov and Popov give the Casimir eigenvalues for the unitary group [26, §2]. The normalization here is that of $\mathfrak{gl}_n$, with the scalar direction included. For $n>0$ the trace-form operator of $\mathfrak{sl}_n$ would subtract d^2/n times the identity; Gross and Taylor's formula is written with that subtraction [27, (2.1)].

Weights are taken with respect to the reversed basis order $e_{n-1},e_{n-2},\dots,e_0$, so the j-th basis vector is e_{n-j}. The order is reversed because a column of height h carries the labels $n-h,\dots,n-1$: its highest wedge $e_{n-1}\wedge\cdots\wedge e_{n-h}$ uses the first h basis vectors in this order. Put

$$\lambda_j=m_{n-j}\qquad(1\le j\le n).$$

By the Observation, λ_j is the number of active columns of height at least j. So λ is a partition, it is the weight of the tensor product of the highest wedges of the active columns, and, putting $j=n-i$,

$$\Gamma_n(\lambda)=\sum_{j=1}^{n}\lambda_j(\lambda_j+n+1-2j)=\sum_{i\in L_n}m_i(m_i+2i+1-n).$$

Proposition 10 (Dominance comparison; supplementary)**.** *Let $n\ge1$.*

(i) *If μ is the highest weight of an irreducible constituent of $\bigotimes_{c\in C_S}\bigwedge^{h_c}\mathbb{C}^n$, then $\mu_1+\cdots+\mu_k\le\lambda_1+\cdots+\lambda_k$ for $1\le k\le n$, with equality at $k=n$.*

(ii) *If μ is a dominant weight satisfying the inequalities and the equality of (i), then $\Gamma_n(\mu)\le\Gamma_n(\lambda)$, with equality only for $\mu=\lambda$.*

Proof. (i) The weight μ is a weight of the product of exterior powers, so it is a sum of weights of basis wedges, one for each active column. A basis wedge in column c uses distinct basis vectors, so at most $\min(k,h_c)$ of them lie among the first k in the reversed order. Thus the first k coordinates of μ sum to at most $\sum_{c\in C_S}\min(k,h_c)=\sum_{j=1}^{k}\lambda_j$, with equality at $k=n$ because both weights have degree d.

(ii) Put $D_k = \sum_{j=1}^{k}(\lambda_j - \mu_j)$ and $v_j = \lambda_j + \mu_j + n + 1 - 2j$. Then $D_k \geq 0$ and $D_n = 0$. The difference $\Gamma_n(\lambda) - \Gamma_n(\mu)$ is the sum of $(\lambda_j - \mu_j)v_j$, which by summation by parts equals

$$\sum_{k=1}^{n-1} D_k(v_k - v_{k+1}),$$

the boundary term vanishing because $D_n = 0$. Since both weights are dominant, $v_k - v_{k+1} \geq 2$. So the difference is nonnegative, and it vanishes only if every D_k is zero, that is, only if $\mu = \lambda$. □

The comparison of the shifted squared norms, with equality only at the highest weight, is the type A case of Humphreys, §13.4, Lemmas B–C [24]. The central contribution is constant here because the total weight is fixed. Etingof uses the same Casimir comparison for the summand generated by highest-weight vectors, then treats exterior powers and determinant twists [33, Proposition 27.1 and §§27.2–27.3].

Now decompose a tensor f of Theorem A orthogonally along the constituents. The quadratic form $\mathrm{Re}\langle f, Kf\rangle$ is the sum of $\Gamma_n(\mu)$ times the squared norms of the components, and Proposition 10 bounds each $\Gamma_n(\mu)$ by $\Gamma_n(\lambda)$. This proves Theorem A for $n \geq 1$. The comparison uses only which weights occur. The exterior Pieri rule gives the multiplicities (Gleizer–Postnikov, Proposition 2.1 [28]).

The bound is sharp. For $n > 0$, the tensor product of the highest wedges of the active columns, tensored with any nonzero column-alternating tensor on the unselected slots, is a highest-weight vector of weight λ. So K multiplies it by $\Gamma_n(\lambda)$, and the coefficient of Theorem A is attained. In particular, alternating the basis tensor $x = b$ within each column gives a nonzero product of these wedges. Let $A_{\mathrm{out}} \subseteq \mathcal{V}'$ be the column-alternating subspace on the unselected slots. By Proposition 10, the admissible equality space is $L_\lambda \otimes A_{\mathrm{out}}$, where L_λ is the Cartan component, the unique summand of largest highest weight, generated by the tensor of highest-weight vectors [13, §1, p. 309]. For two irreducible factors the Cartan component has highest weight equal to the sum of their highest weights and occurs once (Baur, 2003, §1, p. 309 [13]). The same statement for more factors follows by iteration. See also [14, Chapter II, Proposition 3.2, p. 21].

The passage between the two forms of the coefficient in Theorem B(a), by labels, which are the rows of the diagram, and by heights, which are its columns, also has precedents. Gross converts the Casimir formula between rows and columns by pairing indicators of nested rows and counting overlaps [29, Appendix A.1], and Gross and Taylor state the row and column formula [27, (2.1)–(2.2)].

5.2. **Exterior transfer.** Exterior transfer gives another formula for the pair deficit (3). Fix two active columns $c \neq c'$ and work in $\bigwedge^{h_c}\mathbb{C}^n \otimes \bigwedge^{h_{c'}}\mathbb{C}^n$, realized by normalized alternation as in Section 5.1. Let A_i be contraction by e_i and $A_i^\dagger$ exterior multiplication by e_i. For the inner products that make normalized alternation isometric, $A_i^\dagger$ is the adjoint of A_i. The *transfer* from c to c' is $T = \sum_i A_i \otimes A_i^\dagger$, with target degrees $(h_c - 1, h_{c'} + 1)$. Exterior degree $n+1$ is zero. The transfer removes a basis vector from the first wedge and inserts it in the second. The transfer in the opposite direction is $T' = \sum_i A_i^\dagger \otimes A_i$.

Proposition 11 (Transfer squares; supplementary)**.** *On* $\bigwedge^{h_c}\mathbb{C}^n \otimes \bigwedge^{h_{c'}}\mathbb{C}^n$,

$$T^\dagger T = h_c\,\mathrm{Id} - X_{cc'}, \qquad T'^\dagger T' = h_{c'}\,\mathrm{Id} - X_{cc'}.$$

Consequently $\mathrm{Re}\langle f, X_{cc'} f\rangle \leq \min(h_c, h_{c'})\|f\|^2$ *for every tensor* f *that alternates on* S_c *and on* $S_{c'}$.

Proof. Contraction and exterior multiplication satisfy $A_i A_j^\dagger = \delta_{ij}\mathrm{Id} - A_j^\dagger A_i$, and $\sum_i A_i^\dagger A_i = h_c\,\mathrm{Id}$ on $\bigwedge^{h_c}\mathbb{C}^n$. Hence

$$T^\dagger T = \sum_{i,j} A_i^\dagger A_j \otimes A_i A_j^\dagger = h_c\,\mathrm{Id} - \sum_{i,j} A_i^\dagger A_j \otimes A_j^\dagger A_i.$$

The operator $A_i^\dagger A_j$ is the action on the wedge induced by the matrix unit E_{ij} of $\mathfrak{gl}_n$. In exterior degree one it acts as the matrix unit itself. Under normalized alternation the induced action is the sum $\sum_{p\in S_c} E_{ij}^{(p)}$ over the slots of the column. The double sum is therefore $\sum_{p\in S_c}\sum_{q\in S_{c'}}\sum_{i,j} E_{ij}^{(p)}E_{ji}^{(q)}$, which is $X_{cc'}$ by Proposition 1, since $p \neq q$. Exchanging the roles of the two columns gives the identity for T'. Both squares are positive semidefinite, which gives the bound. □

Comparing Proposition 11 with (3) gives, for f as in the proposition,

$$\|Tf\|^2 = h_c\|f\|^2 - \mathrm{Re}\langle f, X_{cc'}f\rangle = \frac{1}{h_{c'}+1}\sum_{p\in S_c}\|f - Y_p f\|^2.$$

Thus $Y_p f = f$ for every box of the shorter column exactly when the transfer out of the shorter column vanishes. This is the one-box exchange condition. The corresponding exterior-power relations in the Schur-module presentation are given by Fulton [16, §8.1] and Akin–Buchsbaum–Weyman [38, §II]. On a product of highest wedges, exterior multiplication repeats a basis vector in the taller column, so the transfer vanishes.

For the commuting actions and exterior transfers, see Cautis–Kamnitzer–Morrison, §4.2, taking $q = 1$ in Lemma 4.2.3 [30]. Howe gives the commuting actions in [31, Theorem 8 and example (d)] and treats skew Howe duality in [32]. The archive uses slot matrix units for the contraction and counts labels to assemble the bound. After contraction its pair estimate uses slot permutations.

5.3. **The content formula.** By Proposition 2, $K = nd\,\mathrm{Id} + 2Z$, where Z is the sum of the transpositions of the d selected slots, each unordered pair taken once. Schur–Weyl duality decomposes the tensor power on the selected slots into constituents indexed by partitions μ of d with at most n rows [33, Theorem 27.3]. The content of a box is its column index minus its row index. On the constituent of shape μ, the operator Z acts by the sum of the contents of the boxes of μ [34, (4.1) and Proposition 4.1]. Okounkov and Vershik obtain the same scalar from the operators of their (2.1), using their Proposition 6.2 and Theorem 6.4 [35]. Twice the content sum is $\sum_j \mu_j(\mu_j + 1 - 2j)$, so adding nd gives $\Gamma_n(\mu)$.

In the projector calculation, the sum Y is a Jucys–Murphy element. In Young's seminormal representation its eigenvalue is the content of the box added at the branching step [22, §2, (7)–(9)]. For a column of positive height s, one can add a box at the right of the first row (content 1) or below the column (content $-s$). The empty partition has one addable box, of content 0, and $Y = 0$. Jucys studied symmetric polynomials in these elements and the center [20]. Murphy gives a construction of the seminormal representation [21].

Column alternation restricts the shapes. A shape μ that supports the column signs occurs in the representation induced from the sign character of the product of the symmetric groups of the columns. Twisting by the sign character turns this into occurrence of the conjugate shape μ' in the Young permutation module of the sorted column heights α. Occurrence there forces μ' to dominate α: this is the necessity of dominance in Young's rule [36, Theorem 6.5 and §6.9]. Conjugation reverses dominance, and conjugation of the shape corresponds to the sign twist [15, I.1, (1.11), and I.7, Example 2]. So μ is dominated by the conjugate of α, which is λ, and

Proposition 10(ii) gives $\Gamma_n(\mu) \le \Gamma_n(\lambda)$. For the converse in Young's rule, see [36, §6.10] and the introduction of [37].

6. The role of the theorem in the archive

This section paraphrases the Lean declarations in Table 3.

6.1. **The use of the bound.** The one caller of the theorem passes it to a step that the archive calls coefficient gluing. That step is part of the archive's proof that regular algebraic cuspidal representations of positive rank over totally complex number fields are cohomological. The result belongs to the classical comparison of regular algebraicity and archimedean cohomology for cuspidal representations of GL_n (Clozel, Lemme 3.14 [39], as used at the start of the proof of Grobner–Raghuram's Theorem 6.3 [40]).

Coefficient gluing applies Theorem A at the rank of the representation, written n here, with the following data. The slots are numbered by an initial segment of the natural numbers, and the columns are indexed by natural numbers. Each slot is attached to an embedding τ of the number field, and S is the fiber of one embedding, the set of all slots attached to τ. The fiber is a union of whole columns, because the column index determines the embedding. The tensor is $f = Jv$, where v is an arbitrary vector of a finite-dimensional coefficient representation and J is an isometric embedding of that representation into the column-alternating tensors, constructed in Section 6.2.

The embedding J intertwines the generators. So the Casimir of the coefficient representation at τ is $J^\dagger K_\tau^{\mathrm{tw}} J$, where K_τ^{tw} is the Casimir at τ on the tensor space, formed from twisted generators. Since $\|Jv\| = \|v\|$, the quadratic form of the coefficient Casimir on v is the form of K_τ^{tw} on Jv. The twisted generators are the $G_{ii'}$ of Section 2.2 with a scalar t, the diagonal twist, subtracted on the diagonal. A shift identity in the archive shows that the twist changes the Casimir by the scalar $-2t|S| + nt^2$. On this fiber $m_i = \nu_\tau(i)$ (Section 6.2). The archive's identity `arith` gives

$$\sum_{i=0}^{n-1} \nu_\tau(i)(\nu_\tau(i) + 2i + 1 - n) - 2t \sum_{i=0}^{n-1} \nu_\tau(i) + nt^2 = \kappa_\tau,$$

where κ_τ is the automorphic Casimir scalar at τ. The condition `IsSmall` requires that ceiling for every coefficient vector. Any positive excess in the tableau coefficient would survive the shift.

With the automorphic action unitary up to twist and the trace scalars opposite, the archive obtains the Casimir difference

$$D = 2 \sum_\tau (I \otimes K_{\mathrm{coeff},\tau} - \kappa_\tau I), \qquad K_{\mathrm{coeff},\tau} = J^\dagger K_\tau^{\mathrm{tw}} J.$$

The fiber estimates give $\mathrm{Re}\langle Dv, v\rangle_B \le 0$ for the archive's invariant Hermitian form B on the product of the automorphic and coefficient modules. In the Kuga calculation on an invariant cochain ω of degree q, the real quadratic form $\mathrm{Re}\langle \omega, \mathrm{kDel}\, d\omega\rangle$ is the sum of these Casimir forms on the values of ω minus a nonnegative term. In the archive's cochain normalization, $\|d\omega\|^2 = (q+1)\,\mathrm{Re}\langle \omega, \mathrm{kDel}\, d\omega\rangle$. It follows that d vanishes in every degree. The nonzero cochain constructed below then represents a nonzero cohomology class. For the classical Laplacian identity, see Borel–Wallach, ch. II, §2.4 and Theorem 2.5(iii) [41].

6.2. **Where the hypotheses come from.** For a regular algebraic cuspidal representation the weights $l_\tau(\cdot)$ at each embedding τ are distinct integers. The K-type construction sorts them in decreasing order, $l_\tau(w_\tau(0)) > l_\tau(w_\tau(1)) > \cdots$, sets $t = \sum_{\tau,j} |l_\tau(j)|$, and puts

$$\nu_\tau(i) = t - l_\tau(w_\tau(i)) - i.$$

The first value $\nu_\tau(0)$ is nonnegative by the choice of t. The difference $\nu_\tau(i+1)-\nu_\tau(i)$ is the gap between two successive weights minus one, which is nonnegative because the weights are distinct integers. So ν_τ is nonnegative and nondecreasing. The slots are the triples (τ,i,r) with $0 \le r < \nu_\tau(i)$. The column of (τ,i,r) is coded by the pair (τ,r), and its label is i. A label then determines at most one slot in a column, which is distinctness. Since ν_τ is nondecreasing, $r < \nu_\tau(i)$ implies $r < \nu_\tau(i')$ for every $i' > i$, which is upward closure. For the fiber of τ the multiplicities are $m_i = \nu_\tau(i)$.

To obtain a compact-group invariant tensor, the archive enlarges the slot set by two copies of the pairs $r < i$, indexed by one embedding from each conjugate pair. The two copies together add i occurrences of the label i at each embedding τ, so the enlarged multiplicity is $\nu_\tau(i)+i = t - l_\tau(w_\tau(i))$. Each factor $u + l_\tau(w_\tau(i)) + \nu_\tau(i) + i$ of the transferred quantum determinant is therefore $u+t$. A step called spherical descent uses the resulting eigenvalue $(u+t)^n$ to obtain a tensor fixed by the compact group. The archive takes the signed column sum of that tensor. A transposition inside a column reindexes the terms of this sum with the opposite sign, and this is the source of the alternation hypothesis. The slot data for the Casimir bound use the original multiplicities $\nu_\tau(i)$.

The Young step supplies a nonzero paired alternation (`altPair`). The cochain contracts the added pairs against Hermitian matrix entries and alternates in its arguments. In `om_ne`, the archive shows that if all these tests vanished, every coordinate of the paired alternation would vanish. The cochain is therefore nonzero and inherits the swap signs on the original slots. Evaluating at a nonzero cochain value and applying a detecting linear functional gives a nonzero scalar tensor with the same signs.

The tensor action commutes with the within-column swaps. The archive chooses a stable subspace of least dimension among those pairing nontrivially with the scalar tensor in the alternating space. An orthonormal basis of that subspace gives the intertwining isometry J from the coefficient module and ensures that its image alternates on each column.

The row lengths in the reversed basis order are $\lambda_j = \nu_\tau(n-j)$. Subtracting the diagonal scalar t changes the highest weight to $\lambda - t(1,\dots,1)$, as in a twist by $\det^{-t}$. With $\rho_j = (n+1-2j)/2$,

$$\lambda_j + \rho_j - t = -\big(l_\tau(w_\tau(n-j)) + (n-1)/2\big).$$

Thus the shifted weights have infinitesimal character opposite to the automorphic one in the archive's normalization. The Casimir scalar is the squared norm of the shifted weight plus ρ, less $\|\rho\|^2$.

6.3. **From cohomology to the main theorem.** The archive supplies Clozel cohomologicality as a field of `published_results_eigen`. A second field uses `regularAlgebraicPurity_of_clozel` to pair cohomologicality with Clozel purity. This packages the purity assertion for regular algebraic representations among the inputs to the argument using theta integrals and local matching. The resulting `EveryCharacter` theorem gives nonvanishing to the right of the critical line for all Dirichlet characters, except at the pole of the principal character at $s = 1$. The ordinary Riemann hypothesis follows by applying `Formalism.VW.rh_of_every` to that result.

7. Concordance with the Lean

The declarations in Tables 1 and 2 are in the namespace `PositivityTheory.Rg.Cz`.

The five proof files `CasBoundHolds.lean`, `Expand.lean`, `ReSwapSumLe.lean`, `CountSame.lean` and `CountCross.lean` are in the folder `Development/WeilBound/RepresentationTheoryGL/CasimirAlgebra/KugaCoordinates/CasimirBounds/`.

The files `KTypeData.lean` and `StepFormula.lean` are in `Development/WeilBound/RepresentationTheoryGL/CasimirAlgebra/TensorRelationsGL/`. Line numbers refer to the archive of 20 November 2026. Private declarations are listed by their source name.

Table 1: Results and their Lean declarations.

Result	**Declaration**	**File and lines**
Theorem A. Tableau Casimir bound	`casBound_holds`	`CasBoundHolds.lean:9-32`. Statement `CasBound` in `StepFormula.lean:88-94`. Expansion at lines 11–19, symmetry 21–22, same-column value 23–24, one-box bound 25–32.
Proposition 1. Contraction of two slot matrix units	`pair_apply`	`Expand.lean:32-55`
Proposition 2. Slot expansion	`cas_expand'`	`Expand.lean:66-92`
Proposition 3. One outside box	`re_swapSum_le'`	`ReSwapSumLe.lean:29-83`. Product identity at lines 47–50, $P_{pq}(Yf)$ 37–55, $Y^2 f$ 56–60, $Yg = -sg$ 62–64, $\|g\|^2 = (s+1)\langle f, g\rangle$ 65–68, conclusion 74–83.
Proposition 4. Higher labels	`count_up`	`CountSame.lean:22-44`
Proposition 5. Equal labels	`count_eqb`	`CountSame.lean:46-69`
Proposition 6. Partition of cross-column pairs	`sum_sum_ite_col_ne_eq_sum_by_col` (private)	`CountCross.lean:59-83`
Proposition 7. Comparison for two distinct columns	`sum_filter_col_le_sum_ite_of_bound` (private)	`CountCross.lean:28-57`. First case at lines 35–39, second case 40–49, interchange and symmetry 50–57.
Proposition 8. Cross-column comparison	`count_cross`	`CountCross.lean:85-102`. Uses the lemmas for Propositions 6 and 7.
Proposition 9. Tableau count	`cas_count`	`CountCross.lean:104-153`. Split into same-column and cross-column pairs at lines 112–124, same-column count 126–131, equal labels 132–137, Proposition 8 at 138, sum per box 146–153.
Theorem B, Propositions 10 and 11, the Observation	Supplementary.	

Table 2: Definitions and auxiliary lemmas in the import closure.

Declaration	**File and lines**	**Role in this document**
`TV`, `eAt`, `tE`, `pm`	`KTypeData.lean:12-27`	The space $\mathcal{V}$, the matrix unit $E_{ii'}^{(p)}$, the generators with a coefficient-module term, and the action of a slot permutation (Section 2.2).
`IsTab`	`KTypeData.lean:51-52`	The tableau condition $\mathcal{T}(\gamma, b)$.
`CasBound`	`StepFormula.lean:88-94`	The statement of Theorem A: hypotheses at lines 88–90, conclusion 91–94.
`eAt_applyC`	`Expand.lean:12-15`	The coordinate formula for a matrix unit.
`tE_zero_apply`	`Expand.lean:17-20`	With the coefficient-module term set to zero a generator is the sum of the slot matrix units.
`sum4`	`Expand.lean:22-30`	Reordering of the fourfold finite sum in Proposition 2.
`pm_pm`	`ReSwapSumLe.lean:15-16`	Products of slot permutations compose.
`dot_pm_swap`	`ReSwapSumLe.lean:18-27`	A swap is self-adjoint (Proposition 3).
`count_rhs`	`CountSame.lean:11-20`	Summing by label, the last step of Proposition 9.
`count_same`	`CountSame.lean:71-88`	Each pair of distinct boxes in a column has one ordering with increasing label (Proposition 9).
`isWH_coeff_one_polynomial` (private)	`CountCross.lean:13-26`	A box with a match in a column has exactly one, so the inner sum on the right of Proposition 7 equals N.

Table 3: Lean locations for Section 6. The Lean writes a for the rank, s for the twist and $\mathrm{Fin}(q)$ for the slot set. The locations in the second part were reported by an AI reviewer of version 1.

Statement	**Lean location**
The sole caller, `clozelCohomological_of_three`	`ClozelCohomologicalOfThree.lean:11`
The bound is passed to coefficient gluing, whose output is an algebraic coefficient representation with the required trace and Casimir comparisons and a nonzero cochain	`ClozelCohomologicalOfRest.lean:19-21`, `IsStarRep.lean:102`, `ClozelCohomologicalOf.lean:37-49`
The application of Theorem A: slots, columns, S the fiber of τ, $f = Jv$; intertwining, isometry and the twist correction	`CoefGlueHolds.lean:120-150`, with the call at line 144
The fiber of an embedding	`KTypeData.lean:60`

Statement	Lean location
The K-type construction supplies the slots and the tableau, with the embedding constant on each column and a formula for the label counts; the cochain construction carries the swap signs to the cochain values	`KTypeHolds.lean:27`, `KTypeData.lean:104-111`, `IsStarRep.lean:88-93`
The Lie-action formula on the fiber (`BigCx`)	`IsStarRep.lean:75-80`
The trace and Casimir scalar of the infinitesimal weights (`CasVal`)	`IsStarRep.lean:95-98`
Shift identity `cas_shift`: the twist contributes $-2t\|S\| + nt^2$	`SumFinVal.lean:33-37`
Identity `arith`: the corrected multiplicity sum is the Casimir scalar of the infinitesimal weights	`SumFinVal.lean:111-114`
The comparison `IsSmall` with ceiling κ_τ	`ClozelCohomologicalOf.lean:37-43`
The Kuga argument: Casimir difference, positivity, vanishing differential, cohomology	`KgCasDiff.lean:37-52`, `KugaGlueHolds.lean:15-21,80-100`, `KugaAbsHolds.lean:137-181`, `IsStarRep.lean:67-70`
The assembled theorem on cohomological representations	`ClozelCohomologicalOf.lean:55-63`, `ClozelCohomological.lean:8-9`, `ClozelCohomologicalModel.lean:10-12`
Sorting of the weights, the twist t and the monotonicity of ν_τ (`sort_data`)	`PKTypeA.lean:207-257`
The slots (τ, i, r) and the column code	`PKTypeA.lean:13` and `PKTypeA.lean:38`
The label count $m_i = \nu_\tau(i)$ on the fiber of τ	`PKTypeA.lean:119` (`cntQ`), `CoefGlueHolds.lean:97-101`
The added pairs $r < i$ over one embedding from each conjugate pair	`PKTypeA.lean:15` and `PKTypeA.lean:259-264`
The i extra occurrences of the label i	`PKTypeA.lean:135-182`
The eigenvalue $(u + t)^n$	`PKType.lean:26-47`
The Hermitian tests for nonvanishing of the cochain	`CochainExpansions/CochainHolds.lean:49-105`
From a nonzero cochain to a nonzero scalar tensor	`CoefGlueHolds.lean:63-84`
A least-dimensional stable subspace satisfying the nonzero pairing condition, and the matrix of its orthonormal basis	`MatrixConjugateSums/Summand.lean:44-56,79-80`
Complex cohomologicality, a consumer of `clozelCohomological`	`ModularityArgument/ClozelCohomology/ClozelCohC.lean:11,24`
Open-image literature input, a consumer of `clozelCohomological`	`TemperedProductOfSquareIntegrable.lean:277-278`
Literature input `clozel_cohomological_of_literature`, a consumer of `clozelCohomological`	`EveryCharacterOfLiteratureOfVWin.lean:327-328`
Cohomologicality and regular algebraic purity enter `published_results_eigen`	Same file, lines 493–494, 666–667 and 689
`every_character_of_literature_of_vWin` takes that package as an argument to the theta-integral result	Same file, lines 810–811

Statement	Lean location
`everyCharacter_of_row` and `everyCharacter`	`CharacterBoxes/MainConclusion.lean:7-11`
`allChi` and `rhOfAll`	`Supplements/Kernels/MainTheorems.lean:7-9`
`riemannHypothesis` uses `rhOfAll`	`Main.lean:4`

8. How this document was produced

8.1. **Version 1.** The automated run took place on 23 and 24 November 2026. Scripts built the theorem's Lean import closure, extracted its dependency graph, produced a plan from that graph and ran checks. The archive uses Lean 4.33.1. The source module is `Development/WeilBound/RepresentationTheoryGL/CasimirAlgebra/KugaCoordinates/CasimirBounds/CasBoundHolds.lean`.

Claude Opus 5.5 directed the process, selected the target, wrote the instructions for the steps and reviewed the exposition. OpenAI models ran through Codex. The build, the plan and the premise classification used `gpt-6.1-terra`. The strategy pass and one of the two first reviews used `gpt-6.2-astra`. The steps for proof-step extraction, related work, drafting, the digest, literature search, scope, supplier tracing, the classical proof structure, citation chasing, audit and revision used `gpt-6.2-sol`, and so did the other first review. One of the two first reviewers was therefore the model that drafted the text.

After the build and the premise classification, the drafting steps recorded the reason for each proof step and compared it with the obvious alternatives. They searched for related work and for precedents of individual ideas, with write-ups of the archive's proof excluded from the search. They drafted the exposition, traced the downstream uses of the theorem and the Lean declarations that supply its hypotheses, added the classical context, followed references in the sources found, and audited quantifiers, necessity and signs.

The first revision applied 22 findings from the two first reviews and rejected one after checking the evidence. A third review, by Claude Opus 5.5, led to three further edits; its other three findings had already been addressed. A later audit led to one more correction, of a local overstatement. Claude Sonnet 5.5 checked findings and ran the release check. It found no mathematical error, matched all nine propositions to their Lean lemmas and verified 12 hand computations. The citation check fetched about 50 citations and flagged 31 of them. The process corrected or narrowed 28 and removed three. The revision logs, the fix log and the citation check that record these counts are in the folder `casbound/review-record` of the release repository, https://github.com/lpkralv/alpoge-lean-expositions.

8.2. **The review of version 1 and the changes in version 2.** On 27 November 2026, AI models completed seven reports on version 1. Fable and Astra wrote the two adversarial verifications. A Claude Fable 5.2 model and an OpenAI Astra model wrote the two referee reports. Opus and Sol each wrote a cold read, and Sol audited the prose. A Claude Opus model consolidated the reports into 55 findings and checked disputed facts against the sources. The revision log identifies Claude Fable 5.2 as the model that drafted version 2.

No reviewer found a mathematical error. Both verifiers recompiled the Lean closure, and one of them matched every statement of version 1 to its Lean declaration. The reports found eight defects, all in citations, attribution or the description of the process, and version 2 corrects them:

(1) Version 1 said that the Lean community had certified "the archive's statement and proof" and cited only the archive and its remark. The certification is the joint statement [3], and it covers the Lean statement and proof of Mathlib's `RiemannHypothesis`.
(2) Version 1 labelled supplementary arguments "checked by hand". AI models checked those arguments. The label is now "supplementary".
(3) Lists of citations had been attached to the last sentence of a paragraph, so that statements about the archive's proof appeared to cite unrelated literature. In one place version 1 said that the coordinate proof "follows" five classical sources. Each citation now accompanies the claim it supports.
(4) Mathematics that is not in the Lean stood unmarked inside proofs tied to Lean lemmas. Supplementary results and blocks now carry a mark and are listed in Section 10.
(5) Version 1 called its graders and first reviewers independent, although one of each was the model that drafted the text. The word is dropped.
(6) The title of Rodgers's reconstruction [4] was wrong for the version cited.
(7) Reference [1] said that the archive links and checksums were on the home page. The accompanying remark gives the download locations [2]. Version 2 reports the part 01 digest measured from a retained local copy.
(8) A pinpoint citation for the isometry of normalized alternation named a section that does not state it. The reason is now given in Section 5.1.

The reviewers also asked for a reorganization, and version 2 makes it. The statement and its classical meaning come first, the proof runs without interruption, alternatives and counterexamples are in remarks and in Section 4, the application is in Section 6, and Lean line references are in Section 7. Theorem 10 of version 1 is now Theorem A. Propositions 1–9 keep their numbers. Version 1 used *background* for the coefficient-module term. Theorem B, the Observation, the running example and Figure 1 are new in version 2. Propositions 10 and 11 state as separate results the dominance comparison and the transfer computation that version 1 had inside other proofs. The abstract no longer reports the test of Section 9. Two physics papers, an unrefereed preprint and a historical paper that were cited for standard facts have been removed from the references, together with links to third-party copies. The statements and proofs of Theorem A and Propositions 1–9 are mathematically unchanged.

The draft drew on version 1, the consolidated findings, the Lean sources, the version 1 citation log and the review reports.

8.3. **Review of version 2.** On 27 November 2026, Fable and Astra checked the revised mathematics and its Lean concordance in two adversarial delta reviews. Codex Sol checked specific sources, bibliographic data and provenance claims. The v1 citation check and this targeted v2 check are the citation audits on record. On the resulting 25-page draft, Fable and Astra separately recommended minor revision, Sol and Terra each wrote a cold read, and Sol audited the prose. The reports cover the draft before the final edits described in `ROUND2.md` and `EDIT-LOG.md`.

9. A test of the method on `typePrimesN`

The method was tested once, on another theorem of the same archive, `UpperTheory.T3.typePrimesN`, whose import closure has 1,194 lines. The test concerns the pipeline. It is not evidence about the present text.

Jeremy Rodgers explains that theorem in Section 3.18 of his reconstruction [4, Section 3.18, pp. 57–59]. The pipeline wrote its exposition without access to write-ups of the proof. Its exclusions covered Rodgers's work, `academia.edu`, `ulam.ai`, `hexagonmath.org` and `alpo.ge`. A separate AI model extracted eight points from Rodgers and other write-ups as an answer key,

which the pipeline did not see. So the test compares the exposition with an answer key that an AI model extracted from published write-ups.

Two AI graders, `gpt-6.2-astra` and `gpt-6.2-sol`, graded the exposition against the key, separately from each other. They were not independent of the pipeline, since `gpt-6.2-sol` is also the model that drafted the text. The graders ran on 24 November 2026, and some pipeline steps were run again between rounds. Four grading rounds ran before review. In the last round before review and in two earlier rounds with retained grade reports, both graders recorded seven of eight points captured and none wrong. After revision, both recorded eight of eight for the exposition. The point missed before the review was why the two primes need disjoint exceptional sets: disjointness gives a usable local test in every residue characteristic. Both reviews in that run raised the omission without seeing the key.

The release repository holds the exposition and the two grades made after the revision, in the folder `validation-typeprimes`. The package of version 1 did not include the grades made before the review.

The test covers one theorem. AI models chose the eight points, wrote the exposition and graded it, and one of the graders was the drafting model. The scores measure agreement between an exposition and eight points chosen by an AI reader. They do not show that a mathematician can learn the argument from the text.

10. Limitations

The recorded reviews and grading used AI models, which may share blind spots. The model `gpt-6.2-sol` drafted version 1, was one of its two first reviewers and one of the two graders of Section 9, and a Sol model wrote two of the seven reports on version 1. The literature step of version 1 needed extensive correction, and the review of version 1 still found the citation defects listed in Section 8. The targeted source checks of version 2 included the theorem and its callers.

The following material is supplementary to the Lean proof.

- The derivation of the label-free bound in the abstract and (1), the Observation of Section 2.1 and its partition interpretation, Figure 1, the trace-form identification of K, and the three comments after Theorem A.
- Theorem B: the form by heights, the sum of squares, the equality case, and the pair-deficit identity (3).
- The running example.
- In Section 3: the explanation by the antisymmetrizing projector before Proposition 3, and every item headed Remarks, including the equality case and the spectrum of Y in Proposition 3, the reading of Proposition 7 by heights, and the count by label after Proposition 8.
- Section 4: the comparisons with weaker estimates, the counterexample with a partly selected column and the variant with terminal subcolumns.
- Section 5: the realization by normalized alternation, the dominance comparison (Proposition 10), the highest-weight proof, sharpness and the equality statement, the transfer squares (Proposition 11) and their comparison with the pair deficit, and the content-formula route.

During production of version 1, an AI model verified 12 hand computations. In the review of version 1, AI reviewers rechecked the highest-weight proof, the transfer squares, the pair-deficit identity, the projector formula, sharpness and the form by heights, algebraically and numerically on small cases. The Fable referee proposed the global sum of squares in Theorem B(b). The pair-deficit identity already appeared in version 1. One reviewer checked the global identity numerically in 13 configurations, and the consolidating model checked it algebraically. One

reviewer checked the variant with terminal subcolumns in 74 two-column cases, and another argued it algebraically. Statements written for version 2 that follow from these by short arguments include the Observation, Theorem B(c), the equality statement of Section 5.1, the numbers in Figure 1 and the two-dimensional running-example calculations. The version 2 delta reviews checked these before the final edit round. The draft reviewed by the referees contained only the two-dimensional example.

The supplied Lean archive has no licence file. This document cites the archive and gives the theorem in mathematical notation, with identifiers for locating the formal results. Reports of errors can go to `https://github.com/lpkralv/alpoge-lean-expositions` or lpslot@me.com.

References

[1] L. Alpöge, Lean archive for the Riemann hypothesis and GRH for Dirichlet L-functions, released 20 November 2026, Lean 4.33.1. The archive was distributed as `lean.tar.gz.part01` through `lean.tar.gz.part14`. Alpöge's accompanying remark gives the download locations [2]. The SHA-256 of part 01 in the retained copy of the 20 November download is `86974b1c242ca7b6748f3299df4fc64ee0f6e06c776ac964fb64888101a3de77`.

[2] L. Alpöge, remark accompanying the Lean archive, undated, `https://alpo.ge/remark.pdf`, read November 2026.

[3] Lean FRO and Mathlib Initiative, *Lean verification of the Riemann Hypothesis*, Lean community blog, 21 November 2026, `https://leanprover-community.github.io/blog/posts/riemann-hypothesis/`.

[4] Jeremy Rodgers, *The ordinary Riemann Hypothesis argument in the supplied Lean source: A mathematical reconstruction for referee checking*, version 2 (repaired referee edition), 22 November 2026, DOI 10.5281/zenodo.24473551, `https://www.everythingequation.com/publications/riemann-hypothesis/paper.pdf`.

[5] Robin Zhang, *A Preliminary Note on Alpöge's Proposed Proof of the Generalized Riemann Hypothesis*, Hexagon 2611.00809v2, 22 November 2026, IHÉS copy consulted 24 November 2026. `https://www.ihes.fr/~rzhang/files/GRH%20note.pdf`.

[6] Siniša Slijepčević, *Outline of Alpöge's proof of the generalised Riemann hypothesis*, Hexagon 2611.00812v1, 21 November 2026, `https://hexagonmath.org/pdf/2611.00812v1`.

[7] Manvir Jaswal, *The Riemann hypothesis for Dirichlet L-functions [after Levent Alpöge]*, Hexagon 2611.00799v1, 21 November 2026, `https://hexagonmath.org/pdf/2611.00799v1`.

[8] Claude Fable 5.2, Codex Astra 6.2, and Gemini Carbon, at the behest of Michael O. Rubinstein, *The Alpöge–Claude proof of the Riemann hypothesis explained for the Riemann zeta function*, working draft, version 2, 22 November 2026, `https://www.math.uwaterloo.ca/~mrubinst/alpoge_claude_grh/zeta-paper_v2.pdf`.

[9] *Notes on Alpöge's Proof of The Riemann Hypothesis*, AI-generated notes printed by Taylor Dupuy, 21 November 2026. `https://tdupu.github.io/riemann-public/riemann-book.pdf`.

[10] Gilheum Park, *Full-Archive AST Stratification, Pinned Mathlib Boundary Audit, and Deep Hub Decompilation of the 17.37M-Line Alpöge–Claude Riemann Hypothesis Formalization*, Hexagon 2611.00881v1, 23 November 2026, `https://hexagonmath.org/pdf/2611.00881v1`.

[11] Jimmie Adriazola, *The Alpöge Proof of the Riemann Hypothesis*, 287-page working manuscript, snapshot of 24 November 2026, RH-Alpoge-Reconstruction, GitHub. `https://github.com/jimjamzola/RH-Alpoge-Reconstruction`.

[12] *Dirichlet L-functions and the critical line*, 23-page companion hosted by Ulam, undated, consulted 24 November 2026. `https://www.ulam.ai/research/rh-short.pdf`.

[13] K. Baur, Cartan components and decomposable tensors, Transformation Groups 8 (2003), no. 4, 309–319, DOI 10.1007/s00031-003-1203-2.

[14] K. Baur, *Two Contributions to the Representation Theory of Algebraic Groups*, doctoral dissertation, University of Basel, 2002, `https://edoc.unibas.ch/entities/publication/1af5150b-0f45-4b80-b0b4-843a4b08af55/full`.

[15] I. G. Macdonald, *Symmetric Functions and Hall Polynomials*, 2nd ed., Clarendon Press, Oxford, 1995.

[16] William Fulton, *Young Tableaux: With Applications to Representation Theory and Geometry*, London Mathematical Society Student Texts 35, Cambridge University Press, 1997.

[17] William Fulton and Joe Harris, *Representation Theory: A First Course*, Graduate Texts in Mathematics 129, Springer, 1991.

[18] T. Lévy, Schur–Weyl duality and the heat kernel measure on the unitary group, Adv. Math. 218 (2008), no. 2, 537–575, arXiv:math/0703690, https://arxiv.org/abs/math/0703690.
[19] A. A. Jucys, Factorization of Young projection operators for the symmetric group, Lietuvos Fizikos Rinkinys 11 (1971), no. 1, 5–10 (Russian).
[20] A.-A. A. Jucys, Symmetric polynomials and the center of the symmetric group ring, Reports on Mathematical Physics 5 (1974), no. 1, 107–112, DOI 10.1016/0034-4877(74)90019-6.
[21] G. E. Murphy, A new construction of Young's seminormal representation of the symmetric groups, Journal of Algebra 69 (1981), no. 2, 287–297, DOI 10.1016/0021-8693(81)90205-2.
[22] A. I. Molev, On the fusion procedure for the symmetric group, Rep. Math. Phys. 61 (2008), 181–188, arXiv:math/0612207v5, https://arxiv.org/abs/math/0612207.
[23] P. Etingof, *Lie Groups and Lie Algebras*, lecture notes for MIT 18.745, Fall 2020, MIT OpenCourseWare, https://ocw.mit.edu/courses/18-745-lie-groups-and-lie-algebras-i-fall-2020/mit18_745_f20_lec_full.pdf.
[24] J. E. Humphreys, *Introduction to Lie Algebras and Representation Theory*, Graduate Texts in Mathematics 9, Springer, 1972.
[25] Z. Yang, New approaches to $\mathfrak{gl}_N$ weight system, Izv. Math. 87 (2023), no. 6, 1255–1270, doi:10.4213/im9379e, arXiv:2202.12225v2.
[26] A. M. Perelomov and V. S. Popov, Casimir operators for semisimple Lie groups, Math. USSR-Izv. 2 (1968), no. 6, 1313–1335, https://www.mathnet.ru/eng/im2520.
[27] D. J. Gross and W. Taylor IV, Two dimensional QCD is a string theory, Nucl. Phys. B 400 (1993), 181–210, arXiv:hep-th/9301068, https://arxiv.org/abs/hep-th/9301068.
[28] O. Gleizer and A. Postnikov, Littlewood-Richardson coefficients via Yang-Baxter equation, Int. Math. Res. Not. 2000 (2000), no. 14, 741–774, https://math.mit.edu/~apost/papers/LR.pdf.
[29] D. J. Gross, Two dimensional QCD as a string theory, Nucl. Phys. B 400 (1993), 161–180, arXiv:hep-th/9212149, https://arxiv.org/abs/hep-th/9212149.
[30] S. Cautis, J. Kamnitzer and S. Morrison, Webs and quantum skew Howe duality, Math. Ann. 360 (2014), 351–390, arXiv:1210.6437v4, https://arxiv.org/abs/1210.6437.
[31] R. Howe, Remarks on classical invariant theory, Trans. Amer. Math. Soc. 313 (1989), no. 2, 539–570, https://www.jstor.org/stable/2001418.
[32] Roger Howe, Perspectives on invariant theory: Schur duality, multiplicity-free actions and beyond, in *The Schur Lectures (1992)*, eds. Ilya Piatetski-Shapiro and Stephen Gelbart, Israel Mathematical Conference Proceedings 8, Bar-Ilan University, 1995, 1–182.
[33] P. Etingof, *Representations of* GL_n*, I*, lecture notes for MIT 18.755, Spring 2024, MIT OpenCourseWare, https://ocw.mit.edu/courses/18-755-lie-groups-and-lie-algebras-ii-spring-2024/mit18_755_s24_lec01.pdf.
[34] F. Cesi, On the eigenvalues of Cayley graphs on the symmetric group generated by a complete multipartite set of transpositions, J. Algebraic Combin. 32 (2010), no. 2, 155–185, arXiv:0902.0727v2, https://arxiv.org/abs/0902.0727.
[35] A. Okounkov and A. Vershik, A new approach to representation theory of symmetric groups, Selecta Math. (N.S.) 2 (1996), 581–605, ESI preprint 333, https://www.esi.ac.at/preprints/esi333.pdf.
[36] M. Hazewinkel and C. F. Martin, Representations of the symmetric group, the specialization order, systems and Grassmann manifolds, L'Enseign. Math. 29 (1983), 53–87, https://ir.cwi.nl/pub/10125/10125D.pdf.
[37] Michiel Hazewinkel and Ton Vorst, On the Snapper/Liebler-Vitale/Lam theorem on permutation representations of the symmetric group, *Journal of Pure and Applied Algebra* 23 (1982), 29–32, https://ir.cwi.nl/pub/10133/10133D.pdf.
[38] K. Akin, D. A. Buchsbaum and J. Weyman, Schur functors and Schur complexes, Adv. Math. 44 (1982), no. 3, 207–278, doi:10.1016/0001-8708(82)90039-1.
[39] Laurent Clozel, Motifs et formes automorphes: applications du principe de fonctorialité, in *Automorphic Forms, Shimura Varieties, and L-functions*, vol. I, eds. L. Clozel and J. S. Milne, Perspectives in Mathematics 10, Academic Press, 1990, 77–159.
[40] H. Grobner and A. Raghuram, On some arithmetic properties of automorphic forms of GL_m over a division algebra, Int. J. Number Theory 10 (2014), no. 4, 963–1013, doi:10.1142/S1793042114500110, https://arxiv.org/abs/1102.1872.
[41] Armand Borel and Nolan R. Wallach, *Continuous Cohomology, Discrete Subgroups, and Representations of Reductive Groups*, 2nd ed., Mathematical Surveys and Monographs 67, American Mathematical Society, 2000.