hexagon:2611.01006v1

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 this entire document, including this notice and the framing sections, without human intervention in the text. It combined mechanical work (a Lean build, dependency extraction and scripted checks) with AI work (OpenAI models via Codex and Anthropic Claude models). No human wrote, edited or checked the mathematical text of this document, including its framing sections.

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. Errors are possible. Corrections are welcome at https://github.com/lpkralv/alpoge-lean-expositions.

Abstract. Levent Alpöge's Lean archive contains a tableau bound for the quadratic Casimir on column-alternating tensors. This automatically generated exposition explains his existing theorem `casBound_holds` and its coordinate proof. Its contribution is expository. The theorem bounds the real quadratic form of the slot Casimir on a union of whole columns by the squared tensor norm times an explicit coefficient in the tableau label multiplicities. Alpöge's proof bounds the swaps from one outside box into a column together, using one squared norm, and then counts the contributions by label. An AI-run survey found no explanation of the theorem in nine write-ups of the proof. On a separate theorem, `typePrimesN`, two independent AI graders found that the pipeline captured seven of eight insights in a hidden key before review and all eight after revision, with no incorrect keyed insights in either version.

1. Introduction

The theorem `PositivityTheory.Rg.Cz.casBound_holds` proves the tableau Casimir bound, stated as `CasBound` in Alpöge's archive. Its only caller, `ClozelCohomologicalOfThree.lean`, uses it in the representation-theoretic input to the Weil-bound part of the proof. Alpöge published the Lean archive for his proof of the Riemann hypothesis and of the generalized Riemann hypothesis for Dirichlet L-functions on 20 November 2026. The Lean community certified the archive's statement and proof on 21 November 2026. The theorem's local import closure is 975 lines excluding Mathlib and contains no `sorry`, `admit` or `axiom`. [1, 2]

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. The survey found explanations for 25 theorems and outlines for 11. It found no coverage of the Casimir bound or of one bookkeeping result. A search of Rodgers's full text also found no treatment of the module. This paper explains Alpöge's existing result and proof. [3, 4, 5, 6, 7, 8, 9, 10, 11]

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.

Section 2 starts with the swap estimate and the source of the hypotheses, then states the bound. Propositions 1–9 develop the coordinate calculation and the tableau count. The concordance links each result to its Lean declaration. Section 3 describes how the text was produced, and Section 4 reports a comparison with a human write-up of another theorem in the archive.

2. The exposition

The tableau Casimir bound. This is an exposition of the tableau Casimir bound in Levent Alpöge's 2026 Lean archive, published on 2026-11-20 [1].

2.1. **The idea in brief.** Expand the quadratic Casimir in the original tensor coordinates, with the column signs kept in view. Bounding the cross-column transpositions one at a time misses the constraint they share: a whole column of target boxes alternates. The central move is to square the defect of the sum of swaps from one outside box into that column, which bounds the entire sum using the nonnegativity of one squared norm. Symmetry then lets us count from the smaller column in either order, and the terminal intervals of tableau labels turn that count and the exact within-column debit into the multiplicity formula.

The highest-weight proof (not in the Lean; checked by hand) bounds the Casimir scalar on each irreducible summand using dominance of highest weights. In Alpöge's coordinate proof, for one box outside an alternating column of size k, one forms $f - Yf$, where Y sums the swaps of that box with the column. The identity $\|f\|^2 - \mathrm{Re}\langle f, Yf\rangle = \|f - Yf\|^2/(k+1)$ bounds the entire sum using the nonnegativity of one squared norm. The Casimir expansion supplies these sums, and counting from the shorter column combines their bounds with the exact within-column signs to give the coefficient in tableau labels [12, lectures 18 and 25–26] and [13, Example 5.4], [14, §2].

The bound assumes a natural-number rank n, a finite slot set Q and a column set C, each with decidable equality, labels $b : Q \to \{0, \dots, n-1\}$ that are distinct and upward closed in each column of $\mathrm{col} : Q \to C$, a union S of whole columns, and a complex tensor f that changes sign under every swap of distinct slots in one column (`StepFormula.lean:88-90`, `KTypeData.lean:51-52`). It bounds the real quadratic form of the zero-background slot Casimir on f by $\sum_{i=0}^{n-1} m_i(m_i + 2i + 1 - n)\|f\|^2$, where m_i counts slots in S with label i (`StepFormula.lean:91-94`).

In the automorphic application, the K-type construction supplies the slots and tableau with an embedding label constant on each column and a formula for each embedding's label counts in terms of the infinitesimal weights, while the cochain construction carries the tensor's swap signs to the cochain values (`KTypeHolds.lean:27`, `KTypeData.lean:104-111`, `IsStarRep.lean:88-93`). The Lie-action formula uses the fiber S of each embedding, which is a union of columns, and the shift identities show how the diagonal twist s contributes $-2s|S| + ns^2$ and how the corrected multiplicity sum becomes the quadratic Casimir scalar of the infinitesimal weights (`KTypeData.lean:60`, `IsStarRep.lean:75-80,95-98`, `SumFinVal.lean:33-37,111-114`).

The archive passes the universal bound to coefficient gluing, whose output is an algebraic coefficient representation with the required trace and Casimir comparisons and a nonzero cochain (`ClozelCohomologicalOfThree.lean:11`, `ClozelCohomologicalOfRest.lean:19-21`, `IsStarRep.lean:102`, `ClozelCohomologicalOf.lean:37-49`). At rank a, coefficient gluing takes slots $\mathrm{Fin}(q)$, column indices $\mathbb{N}$, the supplied column and label maps, $S = \mathrm{fib}(\mathrm{lab}, \tau)$, and $f = Jv$ for each coefficient vector v. Intertwining expresses the coefficient Casimir as $J^\dagger C_b J$, where C_b is the twisted ambient Casimir at τ, and the isometry gives $\|Jv\|^2 = \|v\|^2$. The coefficient quadratic form on v is thus the ambient form on Jv. Adding the twist correction $(-2s|S| + as^2)\|v\|^2$ to the zero-background bound makes the right-hand side $c_\tau\|v\|^2$, where c_τ is the automorphic Casimir scalar (`CoefGlueHolds.lean:120-150`, `SumFinVal.`

`lean:111-114`). The `IsSmall` comparison uses this same c_τ as its ceiling, and a positive excess in the tableau coefficient would survive the twist and leave the estimate too weak for nonzero v (`ClozelCohomologicalOf.lean:39-43`). The separate Kuga argument uses unitarity up to twist and the opposite trace scalars to express its Casimir difference as twice the sum of the coefficient-minus-automorphic differences. The fiberwise upper bound makes this form nonpositive on the tensor product. Positivity then forces the cochain differential to vanish in every degree, and the nonzero cochain yields cohomology (`KgCasDiff.lean:37-52`, `KugaGlueHolds.lean:15-21,80-100`, `KugaAbsHolds.lean:137-181`, `IsStarRep.lean:67-70`). The assembled theorem gives cohomologicality for regular algebraic cuspidal representations of positive rank over totally complex number fields (`ClozelCohomologicalOf.lean:55-63`, `ClozelCohomological.lean:8-9`, `ClozelCohomologicalModel.lean:10-12`).

The exterior-transfer derivation below combines Lévy's slot expansion with Howe's exterior-algebra action. The calculation in Proposition 3 expresses the deficit for each pair of alternating columns as the squared norm of a transfer out of the shorter column. In tensor coordinates, for a source column c and a target column c' of height $k = h_{c'}$, the pair deficit is $h_c\|f\|^2 - \mathrm{Re}\langle f, X_{cc'}f\rangle = \frac{1}{k+1}\sum_{p\in S_c}\|f - Y_p f\|^2$, where $Y_p = \sum_{q\in S_{c'}} P_{pq}$. The Lean proof uses this normalization for the one-box inequality and the inherited labels of whole selected columns to express the bound in label multiplicities [15, Proposition 2.2] and [16, §4.2], [17, Theorem 8 and example (d)].

2.2. **Where the hypothesis comes from.** For a regular algebraic cuspidal representation, the K-type construction sorts the distinct integer weights at each embedding τ in decreasing order $l_\tau(w_\tau(i))$, and sets $s = \sum_{\tau,k}|l_\tau(k)|$, so the integer gaps make $\nu_\tau(i) = s - l_\tau(w_\tau(i)) - i$ nonnegative and nondecreasing in i. The finite slots are triples (τ, i, c) with $0 \le c < \nu_\tau(i)$, numbered as $\mathrm{Fin}(q)$, with natural-number column code for (τ, c) and label i, so a label determines at most one slot per column and raising the label keeps the same column available. The archive adds two copies of the pairs $c < i$ indexed by one embedding from each conjugate pair. The two copies together add i occurrences of label i at each embedding τ, giving the enlarged multiplicity $\nu_\tau(i) + i = s - l_\tau(w_\tau(i))$. Each transferred quantum determinant factor $u + l_\tau(w_\tau(i)) + \nu_\tau(i) + i$ is therefore $u + s$. Spherical descent uses the resulting eigenvalue $(u+s)^n$ to obtain a compact-fixed tensor. The archive takes its signed column sum, whose terms a same-column transposition reindexes with opposite sign. The slot data for the Casimir bound use the original multiplicities $\nu_\tau(i)$. Contracting the added pairs against Hermitian matrix entries and alternating the arguments carries every original-slot sign to the cochain values, since a swap of original slots affects only the input tensor, and Hermitian matrix tests detect the nonzero paired alternation. Evaluating a cochain value at a nonzero coordinate and choosing a complex linear functional that detects it gives a nonzero scalar tensor with those signs, and the ambient tensor action preserves the column-alternating space by commuting with swaps of slots labeled by the same embedding. The archive chooses a stable subspace of least dimension inside that alternating space with nonzero pairing against the scalar tensor, and uses an orthonormal basis to embed every coefficient vector v as a column-alternating tensor Jv. Coefficient gluing applies the bound at the positive automorphic rank with these slots, columns, labels and $f = Jv$, taking S to be each embedding fiber, which is a union of whole columns because the column code fixes the embedding. The proof of simultaneous Whittaker transfer, the individual steps of the spherical descent and the Young contraction proof used to ensure nonzero paired alternation are outside the scope of this exposition.

2.3. **Theorem.** In the statement below, $L_n = \{0, \dots, n-1\}$, with $L_0 = \varnothing$, and $\mathcal{V} = \mathbb{C}^{L_n^Q}$ has the coordinate Hermitian inner product. Here L_n^Q is the set of assignments $x : Q \to L_n$. For

$p, q \in Q$, P_{pq} interchanges the two slots, so $(P_{pq}f)(x) = f(x \circ (p\,q))$, including $P_{pp} = \mathrm{Id}$. The operator $E^{(p)}_{ii'}$ is the matrix unit $e_{i'} \mapsto e_i$ in slot p, acting as the identity in all other slots. The Casimir is $K = \sum_{i,i' \in L_n} G_{ii'} G_{i'i}$ with $G_{ii'} = \sum_{p \in S} E^{(p)}_{ii'}$, and $m_i = |\{p \in S : b(p) = i\}|$.

Theorem 10 (Tableau Casimir bound)**.**

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

The hypotheses are distinctness and upward closure of the tableau labels, selection by whole columns, and alternation of the tensor within every column. They are structural hypotheses of the inequality. No conjecture or separate theorem is assumed. When $S = \varnothing$ the generators and the multiplicities vanish, and when $f = 0$ both quadratic forms vanish. For $n = 0$ the label map forces $Q = \varnothing$, so the same empty-selection case applies.

2.4. **Literature context. Whole-space inequality (Theorem 10).** The highest-weight proof (not in the Lean; checked by hand) gives the full inequality by comparing irreducible summands. Alpöge's coordinate proof bounds the quadratic form on the given tensor itself. The active alternating space identifies with the tensor product of the column exterior powers, with the factors outside S carried along by the identity. In reversed basis order its highest constituent has row lengths $\lambda_j = m_{n-j}$ for $1 \leq j \leq n$, and the trace-form eigenvalue $\Gamma_n(\lambda) = \sum_{j=1}^{n} \lambda_j(\lambda_j + n + 1 - 2j)$ is the displayed coefficient. Constituent dominance and orthogonal decomposition give the upper bound for all alternating tensors. The normalization is for $\mathfrak{gl}_n$, including the scalar direction. For $n > 0$ the trace-form $\mathfrak{sl}_n$ operator would subtract d^2/n times the identity, where $d = |S|$ [18, Proposition 2.1], [12, lectures 18 and 25–26], [13, Example 5.4], [14, §2], [19, §1], [20, Proposition 3.2].

The transposition-content formula gives another route once the signed column condition restricts the allowed shapes. The necessity of dominance in the permutation-module occurrence theorem suffices after twisting by the sign character and conjugating partitions. The converse, established in the literature on Young permutation modules, gives a full occurrence criterion. Proposition 2 converts the content sums to trace-form eigenvalues, which the dominance comparison in Theorem 10 bounds by $\Gamma_n(\lambda)$. A second route works pairwise on exterior powers: the transfer-square calculation in Proposition 3 gives the two height bounds on a column interaction [21, introduction and references], [22, §6], [23, (2.1) and Proposition 6.2], [16, §4.2, Lemma 4.2.3 at $q = 1$], [24, §2], [25, (4.1) and Proposition 4.1, and (4.2) and Theorem 5.3, (5.9)], [17, Theorem 8 and example (d)].

Slot expansion (Propositions 1–2). Lévy contracts the matrix units into diagonal terms and swaps. Szőke uses the negative internal signs on Young-symmetrized tensors. Column alternation alone gives these signs in Alpöge's proof [15, Proposition 2.2] and [26, §III.B, (27), and §IV].

One-box defect (Proposition 3). On a tensor already alternating on the target, dividing the defect by $k + 1$ projects onto tensors alternating on the enlarged set. Alpöge's coordinate

proof derives the norm identity by multiplying swaps. The one-column factor in Jucys's factorization or Molev's branching formula gives this projector [27, (30) and (35)] and [28, (7)–(9)], [29, (3)], [30, (15) and the paragraph after (19)].

Summing into a column (Propositions 3 and 8). The enlarged-set projector bounds the sum from one box. On two exterior powers, the transfer-square calculation in Proposition 3 gives the same smaller-height estimate for the entire pair. Alpöge's proof sums the one-box inequality in the original slots [27, (30) and (35)] and [16, §4.2, Lemma 4.2.3 at $q = 1$], [30, (15) and the paragraph after (19)], [17, Theorem 8 and example (d)].

Choosing the source column (Propositions 6–8). The calculation in Proposition 3 gives a height bound in each transfer direction. The shorter height counts common rows. Alpöge's proof uses inclusion of the inherited label intervals to choose the source for both orders of the swap sum [16, §4.2, Lemma 4.2.3 at $q = 1$] and [31, ch. I, §1], [17, Theorem 8 and example (d)].

Recovering the coefficient (Propositions 4, 5 and 9). The row-and-column Casimir formula gives the coefficient after adjusting the scalar direction. In Alpöge's proof, upward closure supplies a box in the same column for each higher label, and whole-column selection puts each such box in S. Gross's earlier derivation also counts overlaps by pairing indicators for nested rows of the diagram [32, (2.1)–(2.2)] and [31, ch. I, §1], [33, Appendix A.1].

2.5. **Why the proof is built this way.**

- Theorem 10 bounds every column-alternating tensor. Evaluating the Casimir on the tensor of highest wedges tests a single vector.
- Propositions 1 and 2 expand the operator, and Proposition 9 retains the column signs, since an active column of height at least two excludes the fully symmetric active tensors that attain the ambient maximum.
- Proposition 3 groups swaps into one target column, since estimating them individually charges a product of column heights.
- Proposition 3 squares the defect, since a large negative eigenvalue can rule out a norm contraction while helping the upper bound.
- Proposition 7 uses symmetry and label inclusion to obtain both ordered bounds from the smaller source, since the fixed orientations give a sum of heights.
- Proposition 6 groups the cross-column terms before applying the equal-label comparison of Propositions 7 and 8, since individual unequal-label swaps can have positive expectation.
- Proposition 4 uses whole columns, and Proposition 9 combines its count with the count in Proposition 5, since arbitrary selections can leave holes above selected labels and invalidate the multiplicity formula.

2.6. **Notation.** All sums over slots and columns are finite. The column-index set C itself may be arbitrary. A displayed family indexed by c or i uses the same convention at any other index of the corresponding type.

Symbol	Meaning
n	Nonnegative dimension of the one-slot space.
L_n	Label and coordinate-value set $\{0, \dots, n-1\}$, empty at $n = 0$.
Q	Finite set of tensor slots or boxes.
p, q, q', p_1, p_2, p_3, p_*	Boxes in Q, with p_* the witness box in Proposition 7.
C	Set of column indices.
γ	Map assigning each box its column.

Symbol	Meaning
c, c'	Column indices.
S	Set of selected boxes.
d	Number of selected boxes, $\lvert S\rvert$.
S_c	Selected boxes in column c, namely $\{p \in S : \gamma(p) = c\}$.
C_S	Set of active column indices, $\gamma(S)$.
h_c	Height of selected column c, namely $\lvert S_c\rvert$.
ℓ_c	Least label in S_c, defined as n if S_c is empty.
b	Fixed tableau label map $Q \to L_n$.
i, i', i_*	Values in L_n. The value i_* is the witness label $b(p_*)$.
m_i	Number of selected boxes with tableau label i.
$\mathcal{T}(\gamma, b)$	Tableau condition. Two boxes in the same column with the same label coincide, and 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 $b(q) = i$.
$\mathcal{U}(\gamma, S)$	Whole-column selection. For all $p, q \in Q$ with $\gamma(p) = \gamma(q)$, membership of p and q in S is equivalent.
x	Variable coordinate assignment $Q \to L_n$, varied independently of b.
$x[p \leftarrow i']$	Assignment obtained from x by replacing its value at p by i'.
e_i	Standard basis vector of $\mathbb{C}^n$ with index i.
$\mathcal{V}$	Space $\mathbb{C}^{L_n^Q}$ of complex-valued tensor coordinates, with the coordinate Hermitian inner product.
f	Tensor in $\mathcal{V}$.
$\langle \cdot, \cdot \rangle$	Hermitian inner product, conjugate-linear in its first argument and linear in its second.
$\lVert \cdot \rVert$	Norm induced by that inner product.
Re	Real part of a complex scalar.
$\lvert 010\rangle$, $\lvert 100\rangle$, $\lvert 001\rangle$	Coordinate basis tensors for the indicated triples of coordinate values in the designated slot order.
Id	Identity on the operator's underlying space.
$(p\,q)$	Transposition of slots p, q, with the identity convention for $p = q$.
P_{pq}	Tensor operator of $(p\,q)$, sending $f(x)$ to $f(x \circ (p\,q))$ at coordinate x.
$E_{ii'}^{(p)}$	Matrix unit in slot p, taking $e_{i'}$ to e_i in that slot.
$G_{ii'}$	Sum of the matrix units $E_{ii'}^{(p)}$ over selected slots.
K	Trace-form quadratic Casimir $\sum_{i,i' \in L_n} G_{ii'} G_{i'i}$ with zero scalar background.
$z_{ii'}$, z	Scalar background entries for the comparison in Proposition 2, with $z = z_{00}$ when $n = 1$.
B	Set of target slots for the one-box estimate.
s	Number of slots in B.
Y	Sum of the swaps P_{pq} from the outside slot p into B.
g	Defect $f - Yf$.
N	Real scale in the counting argument, specialized to $\lVert f\rVert^2$ in the Casimir application.

Symbol	Meaning
R	Real symmetric kernel in the counting argument, specialized to $\mathrm{Re}\langle f, P_{pq} f\rangle$ at (p, q).
F	Arbitrary real kernel for the reindexing identity.
$\mathbf{1}$	Indicator of the subscripted condition.
$X_{cc'}$	Cross-column interaction, the sum of swaps between the two factors in their exterior-power realization.
λ	Highest weight of the tensor of highest wedges, with row lengths $\lambda_j = m_{n-j}$.
μ	Highest weight of a constituent of the tensor product of the active exterior powers.
j	Row or coordinate index of a weight, starting at 1.
k	Length of a weight prefix, from 1 through n.
Γ_n	Trace-form eigenvalue function on highest weights, specified in the literature context.
D_k	Difference between the length-k prefix sums of λ and μ.
v_j	Coefficient $\lambda_j + \mu_j + n + 1 - 2j$ in the comparison of the trace-form eigenvalues.
$\mathfrak{gl}_n$, $\mathfrak{sl}_n$	Lie algebras of all and of traceless complex $n \times n$ matrices, respectively.
$\mathbb{Z}_{\geq 0}$, $\mathbb{R}$, $\mathbb{C}$	Nonnegative integers, real numbers, and complex numbers, respectively.

Proposition 1 (Contraction of two slot matrix units)**.**

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

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

Proof. By the matrix-unit action, the outer operator forces $i = x(p)$ and reads the intermediate assignment $x[p \leftarrow i']$, obtained by replacing the value at p by 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. Its replacements put the original value of each slot in the other slot and preserve the entire assignment elsewhere. Even if $x(p) = x(q)$, there is still only one contributing term, because the two slots are distinct. □

Remarks. For $n = 0$ and a nonempty slot set, there is no assignment at which to evaluate the identity. At $n = 1$ the two branches have the same scalar effect.

Proposition 2 (Slot expansion of the zero-background Casimir)**.**

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

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

Proof. The generators here have zero scalar background. If a scalar background $z_{ii'}\mathrm{Id}$ were added, the mixed terms and the product of backgrounds would also enter the expansion. For

example, with $n = 1$, S a singleton and $z_{00} = z$, the sole generator would be $(1 + z)\mathrm{Id}$ and its square $(1 + z)^2\mathrm{Id}$, giving Id at $z = 0$ and 4Id at $z = 1$.

Distribute the two sums over S in each product and group the matrix-index contraction by the ordered slot pair (p, q). Proposition 1 gives $n\,\mathrm{Id}$ at $p = q$ and P_{pq} at $p \neq q$. Isolating $n\,\mathrm{Id}$ per box leaves exactly the distinct ordered pairs needed for the count. Including P_{pp} in addition to the scalar term adds $d\,\mathrm{Id}$, whereas replacing the scalar term by $\sum_{p\in S} P_{pp}$ loses $(n-1)d\,\mathrm{Id}$. A singleton selection has operator $n\,\mathrm{Id}$, and at $n = 1$ the diagonal swap alone happens to give the right scalar.

The ordered transposition sum on the entire active tensor power has largest eigenvalue $d(d-1)$ for $n > 0$, attained on fully symmetric tensors. Adding the diagonal term gives the ambient Casimir maximum $nd + d(d-1)$. With a single active column of height $d \leq n$, every distinct swap acts by -1 and the exact coefficient is $nd - d(d-1)$. Thus the ambient estimate has excess $2d(d-1)$, giving 6 versus 2 at $n = d = 2$. The content formula applies once the allowed shapes are restricted. Twisting the Young permutation module associated with the sorted column heights by the sign character converts occurrence to the column-sign condition, and dominance necessity, followed by conjugation of partitions, leaves only shapes dominated by λ. On each such shape μ, doubling the content sum and adding nd gives the trace-form eigenvalue $\Gamma_n(\mu)$. The dominance comparison in Theorem 10 gives $\Gamma_n(\mu) \leq \Gamma_n(\lambda)$, so a full occurrence criterion is unnecessary. The slot argument carries the signs directly into Proposition 9 [34, Theorem 27.3], [22, §6], [31, I.1, (1.11), and I.7, Example 2], [23, (2.1) and Proposition 6.2], [24, §2], [25, (4.1) and Proposition 4.1]. □

Remarks. There are $d(d-1)$ ordered off-diagonal terms. Each unordered pair contributes its swap twice. Empty S makes both terms of the expansion zero.

Proposition 3 (Swaps from one outside box into an alternating set)**.**

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

Proof. Estimating each matrix element by Cauchy–Schwarz, using that each swap is unitary, gives only $s\|f\|^2$ for $s = |B|$. This agrees with the desired bound at $s = 1$ and loses $(s-1)\|f\|^2$ for $s > 1$. For columns of heights $h_c, h_{c'}$, summing those separate estimates costs $2h_c h_{c'}\|f\|^2$ in the two orders. Grouping the swaps into a row costs one squared norm per source box and ultimately gives $2\min(h_c, h_{c'})\|f\|^2$. At equal heights the separate and grouped coefficients are $2h_c^2$ and $2h_c$, respectively. The exterior approach also obtains this pair bound. The calculation below relates the interaction $X_{cc'}$ on two wedges to the positive transfer squares $h_c\mathrm{Id} - X_{cc'}$ and $h_{c'}\mathrm{Id} - X_{cc'}$. The one-box estimate uses only the target's alternation and applies to any set B, a convenience since the application takes $B = S_c$ for a column other than the source column [16, §4.2, Lemma 4.2.3 at $q = 1$], [17, Theorem 8 and example (d)].

To compute the first transfer square, write A_i for contraction by e_i and $A_i^\dagger$ for exterior multiplication by e_i. The transfer from c to c' is $T = \sum_i A_i \otimes A_i^\dagger$. Using $A_i A_j^\dagger = \delta_{ij}\mathrm{Id} - A_j^\dagger A_i$ and $\sum_i A_i^\dagger A_i = h_c\mathrm{Id}$ on the source wedge gives $T^\dagger T = h_c\mathrm{Id} - \sum_{i,j} A_i^\dagger A_j \otimes A_j^\dagger A_i$. The operator $A_i^\dagger A_j$ is the matrix unit acting on the wedge, so the sum is $X_{cc'}$ by the matrix-unit contraction of Proposition 1. Reversing the transfer gives the square with $h_{c'}$.

Set $Y = \sum_{q\in B} P_{pq}$. A bound $\|Yf\| \leq \|f\|$ could give the desired upper bound by Cauchy–Schwarz, but when $n \geq s + 1$ there are nonzero tensors fully alternating on $B \cup \{p\}$, on which $Yf = -sf$. Their norm ratio is s, which exceeds 1 for $s > 1$, though their expectation is already below the upper bound. The sign needed is that of $\mathrm{Re}\langle f, f - Yf\rangle$. The exterior transfers certify

a one-sided bound on a pair of wedges. The one-box calculation below uses the square with just one source box and one target set.

Heuristic: As the target size s grows, the separate upper estimate grows linearly in s, and the fully alternating example has expectation $-s\|f\|^2$. In the square of the swap sum there are s diagonal terms and $s(s-1)$ off-diagonal terms, with a minus sign and multiplicity $s-1$ per target in the latter. Coefficient by coefficient, the normalized quadratic has leading form $\mathrm{Id}-Y$ on f, and the squared-defect formula will divide by a factor of leading size s.

The exact product identity for distinct $q, q' \in B$ is $P_{pq}P_{pq'} = P_{pq'}P_{qq'}$, the transposition relation on three distinct slots. It places the internal swap next to f, giving $-P_{pq'}f$, and the s diagonal products give sf. Each q' appears $s-1$ times in the off-diagonal part, so

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

This identity yields a constant independent of s. The calculation uses $p \notin B$. If $p \in B$ with $s \geq 1$, the self-swap is the identity and the other $s-1$ swaps negate f, so $Yf = (2-s)f$. The inequality still holds there, but at $s = 2$ with $f \neq 0$ the quadratic identity would predict $Y^2 f = 2f$ instead of zero.

Set $g = f - Yf$. The product identity gives $Yg = -sg$. 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 works for every $s \geq 0$. At $s = 0$ the sum is zero, $g = f$, and the divisor is 1. Equality holds exactly when $Yf = f$, and on a tensor fully alternating on $B \cup \{p\}$ one has $g = (s+1)f$. The squared norm $\|Yf\|^2$ alone gives no sign for the difference. □

Remarks. The same squared-defect identity follows from the orthogonal antisymmetrizer on $B \cup \{p\}$. Its signed group average splits into $s+1$ cosets of the subgroup permuting B, represented by the identity and the swaps P_{pq}. On a vector alternating on B the average equals $(f-Yf)/(s+1)$, so the equality comes from the norm identity for an orthogonal projection.

Proposition 4 (The number of higher labels in a selected column)**.**

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

Proof. The strict cutoff $b(p)$ identifies the increasing ordering of each distinct same-column pair. Using a weak inequality counts p itself and raises the answer by one. A common cutoff would give the same count for every box in a column containing all labels, whereas the count at label 0 is $n-1$ and the count at label $n-1$ is 0.

Map each selected box in the given column whose label exceeds $b(p)$ to its label. Injectivity alone gives an upper bound on the number of such boxes. For the upper Casimir estimate this goes in the wrong direction, because the count is multiplied by $-2\|f\|^2$. Upward closure supplies a box for every label in L_n above $b(p)$, and whole-column selection puts each such box in S. The map is thus a bijection onto the $n-1-b(p)$ labels $b(p)+1, \ldots, n-1$.

Consider using the same multiplicity formula after arbitrary selection inside a column. At $n = 2$, take an ambient column with labels $0, 1$ and select only its label-0 box. The selected box has zero selected boxes above it, although $n-1-b(p) = 1$. A nonzero alternating tensor on the ambient column exists, and by Proposition 2 the singleton Casimir is $2\,\mathrm{Id}$. The original selected label gives $m_0 = 1$, $m_1 = 0$, so the proposed bound would read $2\|f\|^2 \leq 0$. The height expression from the two-column exterior calculation assigns a singleton of height 1 the

coefficient $1(2+1-1)=2$. Its height remains valid for that selection, while the retained label 0 fails to encode the terminal interval for height 1 [15, Proposition 2.2] and [32, (2.1)–(2.2)], [33, Appendix A.1]. □

Remarks. In a variant (not in the Lean; checked by hand), selecting terminal subcolumns would preserve both the upward count and the nested-label comparison below. The whole-column hypothesis ensures these properties directly. At $n=0$ there is no $p\in S$ in the proposition.

Proposition 5 (The number of equal labels in other columns)**.**

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

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 label injectivity makes it a singleton. Subtracting one therefore removes exactly p. When $m_{b(p)}=1$, there is no other-column match. Using $m_{b(p)}$ as the other-column count would add one charge N per box, hence dN overall. If $p\notin S$, the same-column part is empty and the other-column count is $m_{b(p)}$. □

Remarks. The hypothesis includes upward closure as part of the tableau condition, although this count uses only injectivity in each column. The subset S need not be a union of whole columns here.

Proposition 6 (Partition of cross-column pairs)**.**

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

Proof. Each ordered cross-column pair belongs to the unique product $S_c\times S_{c'}$ with $c=\gamma(p)$ and $c'=\gamma(q)$. For distinct column indices, every pair in that product passes the cross-column test. For equal indices, the whole product is omitted by that test, including pairs of distinct boxes. This partitions exactly the domain of the sum. The use of an arbitrary kernel is a convenience: the proof needs the same partition for $F(p,q)=R(p,q)$ and for $F(p,q)=N\mathbf{1}_{b(p)=b(q)}$, so the two sums can be compared on identical ordered column pairs in a single grouping step. □

Remarks. The reindexing depends only on the column map and the chosen subset. If $S=\varnothing$, then $C_S=\varnothing$ and both sums are zero.

Proposition 7 (Comparison for two distinct columns)**.**

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

Proof. Applying the row bound in the displayed order gives h_cN, and in the reverse order it gives $h_{c'}N$. For $N>0$ and $h_c>h_{c'}$, the first estimate overcharges by $(h_c-h_{c'})N$. The two ordered sums are equal by symmetry of R, so using the smaller source twice yields $2h_{c'}N$ in

place of $(h_c + h_{c'})N$. At heights 3 and 1 the two estimates are $2N$ and $4N$. The transfer-square calculation in Proposition 3 gives the two height bounds for the same interaction. In the coordinate proof, symmetry permits both orientations, and label inclusion picks the orientation used for the count [16, §4.2, Lemma 4.2.3 at $q = 1$], [17, Theorem 8 and example (d)].

If one height grows while the other stays fixed, the sum of the fixed-orientation bounds $(h_c + h_{c'})N$ increases for $N > 0$, stays zero for $N = 0$, and decreases for $N < 0$. Once the shorter height is fixed, the number of common labels is fixed as well, and orienting both sums from the shorter column gives $2\min(h_c, h_{c'})N$ in total.

The label sets are terminal intervals, possibly empty, by upward closure and selection by whole columns. To see the orientation explicitly, first suppose every label in S_c has a match in $S_{c'}$. For each $p \in S_c$, there is then exactly one such match by label injectivity, so the row on the right sums to exactly N, which bounds the row on the left. Equal label sets use this branch.

Otherwise take $p_* \in S_c$ with label $i_* = b(p_*)$ missing from $S_{c'}$. A box $q \in S_{c'}$ of label i_* contradicts the choice of i_*. If $b(q) < i_*$, upward closure from q supplies i_* in its column, and selection by whole columns puts that box in $S_{c'}$, again a contradiction. Therefore $b(q) > i_*$, and upward closure from p_* and whole-column selection give q a match in S_c. Now interchange the sums and use $R(p, q) = R(q, p)$ to apply the row bound to each $q \in S_{c'}$, with its unique match giving exactly N on the right.

In either branch, the right side is $\min(h_c, h_{c'})N$, because the contained label set has one match per box. This is an exact count, so the proof works for $N = 0$ and $N < 0$ as well. In particular, for $N < 0$ the orientation is justified by the matching identity, independently of which product of a height and N is numerically smaller. □

Remarks. Write $\ell_c, \ell_{c'}$ for the least selected label in each column, with value n for an empty column. In the missing-label branch any i_* with $\ell_c \le i_* < \ell_{c'}$ works, since only its presence in S_c and its absence from $S_{c'}$ are needed. The statement allows empty columns for convenience, while Proposition 8 uses only indices in C_S. If either column is empty, both double sums vanish. If $p \in S$ and some column $c \neq \gamma(p)$ has $S_c = \varnothing$, the row bound for p and c reads $0 \le N$.

Proposition 8 (Cross-column comparison by equal labels)**.**

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

Proof. Replacing each cross-column entry by its own equal-label charge would use the pointwise comparison $R(p, q) \le N\mathbf{1}_{b(p)=b(q)}$. Take $n = 2$, with slots p_1, p_2 forming a column with labels $0, 1$, respectively, and a singleton p_3 with label 1. In this slot order take $f = (|010\rangle - |100\rangle)/\sqrt{2}$, which alternates in p_1, p_2 and has squared norm 1. The swap of p_1, p_3 fixes $|010\rangle$ and sends $|100\rangle$ to $|001\rangle$, which is orthogonal to f. Its expectation is $1/2 > 0$ despite unequal tableau labels. The swap of p_2, p_3 fixes $|100\rangle$ and sends $|010\rangle$ to $|001\rangle$, giving the other expectation $1/2$. The column total $1/2 + 1/2 = 1$ is exactly the single equal-label charge. The label b stays fixed in the count while the coordinate assignment x varies in the inner product. The transfer-square calculation in Proposition 3 bounds the full sum $X_{cc'}$. Its coefficient $\min(h_c, h_{c'})$ counts common labels for terminal intervals without an entrywise comparison [16, §4.2, Lemma 4.2.3 at $q = 1$], [17, Theorem 8 and example (d)].

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

The same bound can be counted by label using whole rows. Orient each unordered pair of active columns from the shorter column, with one choice of orientation in a tie. Each source label i occurs once there and also occurs in the target column. Its row contributes at most N, and symmetry lets that row bound account for both orders. Exactly $\binom{m_i}{2}$ unordered column pairs contain i in both columns, so the total upper bound from rows charged to label i is $2\binom{m_i}{2}N$. □

Remarks. With at most one active column the cross-column domain is empty, so both sides vanish for every real N.

Proposition 9 (The tableau count for a symmetric kernel)**.**

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

$$\textit{Then} \qquad ndN + \sum_{p \in S} \sum_{\substack{q \in S \\ q \neq p}} R(p,q) \leq \left(\sum_{i \in L_n} m_i(m_i + 2i + 1 - n) \right) N.$$

Proof. The formulation for an arbitrary symmetric real kernel and a real scale of either sign is a convenience. The application uses $R(p,q) = \mathrm{Re}\langle f, P_{pq} f \rangle$ and $N = \|f\|^2 \geq 0$. Here the within-column contribution is an exact identity multiplied by N, and the cross-column bound already holds in its signed form.

The signed internal contribution from a column of height h_c is $-h_c(h_c - 1)N$, of magnitude $h_c(h_c - 1)|N|$. In the tensor application $N = \|f\|^2 \geq 0$, so for a tall column the internal debit has leading magnitude $h_c^2 N$, compared with a diagonal charge $nh_c N$. Across two equally tall columns the separate estimates grow as $2h_c^2 N$, while the whole-column estimate grows as $2h_c N$. The latter counts only common labels. These are scales for the terms separately. When $h_c = n$ the quadratic part of the diagonal and internal contributions cancels and their combined charge is only nN.

Split the distinct ordered pairs into same-column and cross-column pairs. The case $q = p$ has no term in this sum. Within a column the hypothesis supplies $-N$ for distinct pairs only, and across columns the terms retain the kernel $R(p,q)$. Assigning $-N$ to the diagonal as well would add an erroneous debit of dN, already changing the answer for a singleton when $N > 0$.

Each unordered pair of distinct boxes in a column has exactly one increasing-label ordering, by label injectivity. Charge both orders to its lower label. Proposition 4 counts $n - 1 - b(p)$ such partners for p, giving the exact signed contribution

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

For $N \geq 0$, dropping the nonpositive terms loses $2N \sum_{p \in S}(n - 1 - b(p))$. A single active column of height h_c has no cross-column contribution and gives exactly $nh_c - h_c(h_c - 1)$ as the coefficient of N. Dropping the signs gives the coefficient nh_c.

For the remaining ordered pairs, Proposition 8 bounds the sum by one N for each equal-label match across columns. Proposition 5 gives $m_{b(p)} - 1$ such matches from p. Adding the

diagonal n, the internal debit $-2(n-1-b(p))$, and the cross-column charge $m_{b(p)}-1$ yields $m_{b(p)}+2b(p)+1-n$ per box. Summing by label puts m_i identical terms together and gives the right side of the proposition. □

Remarks. With $S=\varnothing$ both sides vanish. With $N=0$ the within-column and diagonal parts vanish and the cross-column part is nonpositive. The latter part can be strictly negative for an abstract kernel even at $N=0$.

2.7. **Theorem 10. Application to the slot Casimir.**

Proof. Evaluating K on the tensor of highest wedges identifies a distinguished eigenvalue. That direct test alone gives no bound on every alternating tensor. At $n=2$ with two singleton columns, Proposition 2 gives $K=4\,\mathrm{Id}+2P_{p_1p_2}$, where p_1,p_2 are the two slots. Both the symmetric and alternating tensors are admissible, with eigenvalues 6 and 2, respectively.

Highest-weight proof (not in the Lean; checked by hand). For any constituent highest weight μ, a basis wedge in column c uses distinct indices, so at most $\min(k,h_c)$ of its indices can lie among the first k basis vectors in 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 the degrees agree. Put $D_k=\sum_{j=1}^{k}(\lambda_j-\mu_j)$ and $v_j=\lambda_j+\mu_j+n+1-2j$. The difference of the trace-form eigenvalues is the sum of $(\lambda_j-\mu_j)v_j$, which telescopes to $\sum_{k=1}^{n-1}D_k(v_k-v_{k+1})$. Each D_k is nonnegative and $v_k-v_{k+1}\geq 2$ since both weights are dominant. Orthogonal decomposition now bounds the quadratic form on arbitrary sums of constituents, with the factors outside S acting as spectators. This comparison needs no constituent multiplicities from exterior Pieri. Alpöge's coordinate proof follows [18, Proposition 2.1], [12, lectures 18 and 25–26], [13, Example 5.4], [14, §2], [20, Proposition 3.2].

Set $N=\|f\|^2$ and $R(p,q)=\mathrm{Re}\langle f,P_{pq}f\rangle$. By Proposition 2 and linearity of the real part of the inner product in its second argument over the reals,

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

Keeping exact within-column terms but bounding cross-column swaps individually would cost $h_ch_{c'}N$ for the ordered pair of columns c,c'. For two full columns at $n=2$, the diagonal gives $8N$ and the internal terms give $-4N$. The individual estimates add $8N$ and give $12N$, whereas the grouped bounds add $4N$ and give $8N$. The reduction from $h_ch_{c'}$ charges to $\min(h_c,h_{c'})$ comes from Propositions 3 and 7.

The kernel is symmetric because $P_{pq}=P_{qp}$. The global alternation hypothesis first gives the exact value $R(p,q)=-N$ for distinct same-column boxes. It is used a second time with only the alternation within a target column S_c. For $p\in S$ and $c\neq\gamma(p)$, that column lies entirely away from p, so Proposition 3 bounds $\sum_{q\in S_c}R(p,q)$ by N. Empty target columns also satisfy the bound, which then reads $0\leq N$. These are exactly the hypotheses of Proposition 9. Its conclusion is the asserted inequality. □

Remarks. The coefficient can also be written by heights as

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

where the second sum is over two-element subsets. The first sum is the diagonal plus the exact internal contribution. The second counts each cross-column interaction twice. Sharpness (not in the Lean; checked by hand). For $n>0$, the tensor of nested highest wedges attains this coefficient by the trace-form eigenvalue formula. If $S=\varnothing$ or $f=0$, both sides are zero, and

$n = 0$ forces $Q = S = \varnothing$ [13, Example 5.4], [18, Proposition 2.1], [33, Appendix A.1], [14, §2], [20, Proposition 3.2].

2.8. **Routine and classical facts.**

- Finite sums can be reordered, expanded by distributivity, or collected according to the values of a map on their index set.
- In the matrix-unit action, $E_{ii'}^{(p)} f$ at x is $f(x[p \leftarrow i'])$ when $x(p) = i$, and is zero otherwise. A zero scalar background gives only the sum of the slot matrix units.
- A slot permutation acts by permuting the coordinate basis and products of the operators follow composition of the permutations. A transposition is a unitary involution, hence self-adjoint.
- The Hermitian norm is nonnegative, and the real part of a Hermitian matrix element is bounded above by the product of the two norms by Cauchy–Schwarz.
- The interval of labels strictly above $i \in L_n$ has $n - 1 - i$ elements. In a column with distinct labels, a prescribed label contributes exactly once when present and zero times when absent.
- Within one column, each unordered pair of distinct boxes corresponds to two ordered pairs and to exactly one pair in increasing label order.
- For any real-valued function of the tableau label, the sum of its values over selected boxes uses the value at label i exactly m_i times.
- The signed average over a finite unitary permutation group is the orthogonal projector onto its sign subspace.
- Normalized alternation identifies $\bigwedge^{h_c} \mathbb{C}^n$ isometrically with the alternating tensors of degree h_c. Each unit basis wedge maps to the signed sum of its $h_c!$ tensor orderings divided by $\sqrt{h_c!}$ [18, §2].
- Taking invariant orthogonal complements gives an orthogonal decomposition of a finite-dimensional unitary representation into irreducibles. The trace-form quadratic Casimir acts on the highest-weight constituent of weight μ by $\Gamma_n(\mu)$ [12, lectures 18 and 25–26], [13, Example 5.4], [14, §2].
- Schur–Weyl duality identifies the tensor-power constituents indexed by partitions. On the symmetric-group constituent the sum of all transpositions has scalar equal to the sum of the contents of its boxes [34, Theorem 27.3], [23, (2.1) and Proposition 6.2], [24, §2], [25, (4.1) and Proposition 4.1].
- Occurrence in a Young permutation module forces dominance of its indexing partition by the constituent's partition. Tensoring with sign conjugates the partition, reversing dominance [22, Theorem 6.5 and §6.9], [31, I.1, (1.11), and I.7, Example 2].
- On a tensor product of exterior powers the exterior transfer moves a basis vector from one factor into the other by contraction and wedge multiplication. The calculation in Proposition 3 identifies the squared norm form in each direction with the corresponding height times the identity minus the interaction [16, §4.2, Lemma 4.2.3 at $q = 1$], [17, Theorem 8 and example (d)].

2.9. **Concordance.** Module-local declarations are listed by their source name within the indicated module.

Displayed result	**Declaration**	**Source**
Theorem 10. Tableau Casimir bound	`PositivityTheory.Rg.Cz.casBound_holds`	`CasBoundHolds.lean:9-32`. Statement in `StepFormula.lean:88-94`.

Displayed result	Declaration	Source
Proposition 1. Contraction of two slot matrix units	`PositivityTheory.Rg.Cz.pair_apply`	`Expand.lean:32-55`
Proposition 2. Slot expansion	`PositivityTheory.Rg.Cz.cas_expand'`	`Expand.lean:66-92`
Proposition 3. One outside box	`PositivityTheory.Rg.Cz.reSwapSum_le'`	`ReSwapSumLe.lean:29-83`
Proposition 4. Higher labels	`PositivityTheory.Rg.Cz.count_up`	`CountSame.lean:22-44`
Proposition 5. Equal labels	`PositivityTheory.Rg.Cz.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`
Proposition 8. Cross-column comparison	`PositivityTheory.Rg.Cz.count_cross`	`CountCross.lean:85-102`. Uses the lemmas for Propositions 6 and 7.
Proposition 9. Tableau count	`PositivityTheory.Rg.Cz.cas_count`	`CountCross.lean:104-153`

3. How this document was produced

The automated run took place on 23 and 24 November 2026. A Python driver ran the steps. The mechanical work built the theorem's Lean import closure, extracted its dependency graph, produced a plan from that graph, and ran scripted 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, plan and premise classification used `gpt-6.1-terra`. The strategy pass and one of the independent reviews used `gpt-6.2-astra`. The Codex 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`, as did the second independent review.

After the build and premise classification, the process extracted reasons for individual proof steps and examined why obvious alternatives fail. It searched for related work and precedents for individual ideas, excluding write-ups of Alpöge's proof from the literature search. It drafted the exposition, produced a digest, and added the thesis and labels. It traced downstream uses condition by condition under a stopping rule and followed the Lean declarations that supply the hypotheses. It then added context on the structure shared with classical proofs, followed citations one hop, and audited quantifiers, necessity and signs.

The first revision applied 22 findings from the two independent reviews and rejected one after checking the evidence. A third review by Claude Opus 5.5 led to three further edits, and its other three findings had already been addressed. A fresh audit caught one local overstatement that the revision fixed. 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. A fetch-check of about 50 citations found 31 that needed correction or narrowing. The process corrected 28 and removed three. It fetched every remaining citation, and a final check judged the exposition ready.

4. Validation on `typePrimesN`

The validation used `UpperTheory.T3.typePrimesN`, a step in the same archive (1,194 lines). Jeremy Rodgers explains it in Section 3.18 of his reconstruction (version 2, 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 reader extracted eight key insights from Rodgers and other write-ups for a hidden answer key. [3, Section 3.18]

Two AI graders, `gpt-6.2-astra` and `gpt-6.2-sol`, graded the exposition independently. The two graders found seven insights captured, one missed and zero wrong before review, and eight captured, zero missed and zero wrong after revision. The missing insight was why the two primes need disjoint exceptional sets. Disjointness gives a usable local test in every residue characteristic. Both reviewers found this omission without seeing the key. The comparison covers one theorem, and AI models extracted the key and assessed the exposition.

5. Limitations

No human expert has read this document. The review and grading used AI models that may share blind spots. The literature step required extensive correction before the citation check was complete. The audit covered the theorem and its callers. Claims elsewhere in the archive remain outside that audit.

The highest-weight proof and the two remarks labelled **not in the Lean; checked by hand** were checked by an AI model. The label describes a mathematical check outside Lean. The supplied Lean archive has no licence file. The paper 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, 20 November 2026, `lean.tar.gz.part01` through `lean.tar.gz.part14`, Lean 4.33.1. Archive links and `SHA256SUMS` at `https://alpo.ge`. First part at `https://alpo.ge/lean.tar.gz.part01`. SHA-256 of part 01 `86974b1c242ca7b6748f3299df4fc64ee0f6e06c776ac964fb64888101a3de77`.

[2] L. Alpöge, Announcement and remark accompanying the Lean archive, `https://alpo.ge/remark.pdf`.

[3] J. Rodgers, *The Riemann Hypothesis Argument of Levent Alpöge and Anthropic's Claude: A Full Human-Readable Reconstruction from Lean*, version 2, `https://www.everythingequation.com/publications/riemann-hypothesis/paper.pdf`.

[4] R. Zhang, *A Preliminary Note on Alpöge's Proposed Proof of the Generalized Riemann Hypothesis*, `https://www.ihes.fr/~rzhang/files/GRH%20note.pdf`.

[5] S. Slijepčević, *Outline of Alpöge's proof of the generalised Riemann hypothesis*, Hexagon 2611.00812v1, `https://hexagonmath.org/pdf/2611.00812v1`.

[6] M. Jaswal, *The Riemann hypothesis for Dirichlet L-functions [after Levent Alpöge]*, Hexagon 2611.00799v1, `https://hexagonmath.org/pdf/2611.00799v1`.

[7] M. O. Rubinstein, *The Alpöge–Claude proof of the Riemann hypothesis explained for the Riemann zeta function*, working draft, version 2, `https://www.math.uwaterloo.ca/~mrubinst/alpoge_claude_grh/zeta-paper_v2.pdf`.

[8] *Notes on the Riemann Hypothesis*, distributed by T. Dupuy, `https://tdupu.github.io/riemann-public/riemann-book.pdf`. Repository and attribution at `https://github.com/tdupu/riemann-public`.

[9] G. 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, https://hexagonmath.org/pdf/2611.00881v1.
[10] J. Adriazola, Reconstruction of Alpöge's Riemann hypothesis proof, https://github.com/jimjamzola/RH-Alpoge-Reconstruction.
[11] Ulam, companion to Alpöge's Riemann hypothesis proof, https://ulam.ai.
[12] P. Etingof, *Lie Groups and Lie Algebras*, MIT OpenCourseWare 18.745, Fall 2020, https://ocw.mit.edu/courses/18-745-lie-groups-and-lie-algebras-i-fall-2020/mit18_745_f20_lec_full.pdf.
[13] Z. Yang, New approaches to $\mathfrak{gl}_N$ weight system, Izv. Math. 87 (2023), no. 6, 1255–1270, arXiv:2202.12225v2, https://arxiv.org/pdf/2202.12225.
[14] 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/links/31651a58ff26049b02381fcda398189c/im2520_eng.pdf.
[15] T. Lévy, Schur–Weyl duality and the heat kernel measure on the unitary group, Adv. Math. 218 (2008), no. 2, 537–575, https://arxiv.org/pdf/math/0703690.
[16] 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/pdf/1210.6437.
[17] R. Howe, Remarks on classical invariant theory, Trans. Amer. Math. Soc. 313 (1989), no. 2, 539–570, https://www.math.utah.edu/~ptrapa/math-library/howe/howe-remarks-invariant-theory-TAMS-1989.pdf.
[18] 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.
[19] K. Baur, Cartan components and decomposable tensors, Transformation Groups 8(4) (2003), 309–319.
[20] 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.
[21] M. Hazewinkel and A. C. F. Vorst, On the Snapper/Liebler-Vitale/Lam theorem on permutation representations of the symmetric group, J. Pure Appl. Algebra 23 (1982), 29–32, https://ir.cwi.nl/pub/10134/10134D.pdf.
[22] 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.
[23] 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.
[24] G. Frobenius, Über die Charaktere der symmetrischen Gruppe, Sitzungsber. Königl. Preuss. Akad. Wiss. Berlin (1900), 516–534, https://faculty.fordham.edu/dswinarski/snchar/Frobenius1900.pdf.
[25] 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/pdf/0902.0727.
[26] É. Szőke, A quantummechanical derivation of the eigenvalues of the quadratic Casimir operator of the algebra $SU(n)$ in Young tableau representation, author-uploaded preprint, 20 September 2018, https://www.researchgate.net/publication/327777204_Eigenvalues_of_the_quadratic_Casimir_operator_of_SUn_in_Young_tableaux_representation.
[27] P. Navrátil, G. P. Kamuntavičius and B. R. Barrett, Few-nucleon systems in translationally invariant harmonic oscillator basis, Phys. Rev. C 61 (2000), 044001, https://arxiv.org/pdf/nucl-th/9907054.
[28] A. I. Molev, On the fusion procedure for the symmetric group, Rep. Math. Phys. 61 (2008), 181–188, arXiv:math/0612207v5, https://arxiv.org/pdf/math/0612207.
[29] A. A. Jucys, Factorization of Young projection operators for the symmetric group, Lietuvos Fizikos Rinkinys 11 (1971), no. 1, 5–10 (Russian), https://ncatlab.org/nlab/files/Jucys-1971.pdf.
[30] P. Navrátil and B. R. Barrett, Four-nucleon shell-model calculations in a Faddeev-like approach, Phys. Rev. C 59 (1999), 1906–1918, https://arxiv.org/pdf/nucl-th/9812062.
[31] I. G. Macdonald, *Symmetric Functions and Hall Polynomials*, 2nd ed., Clarendon Press, Oxford, 1995, https://math.berkeley.edu/~corteel/MATH249/macdonald.pdf.
[32] D. J. Gross and W. Taylor IV, Two dimensional QCD is a string theory, Nucl. Phys. B 400 (1993), 181–210, https://arxiv.org/pdf/hep-th/9301068.
[33] D. J. Gross, Two dimensional QCD as a string theory, Nucl. Phys. B 400 (1993), 161–180, https://arxiv.org/pdf/hep-th/9212149.

[34] P. Etingof, *Representations of* GL_n*, I*, MIT OpenCourseWare 18.755, Spring 2024, https://ocw.mit.edu/courses/18-755-lie-groups-and-lie-algebras-ii-spring-2024/mit18_755_s24_lec01.pdf.