hexagon:2611.00219v2

MONOMIAL IMMANANTS OF TOTALLY NONNEGATIVE MATRICES

QIYUE TANG

Abstract. We prove that every monomial immanant is nonnegative on the totally nonnegative real matrices, establishing a conjecture of Stembridge. We identify the corresponding normalized symmetric function of annular holonomy with a single cluster theta function.

1. Introduction

A real matrix is *totally nonnegative* if all its minors are nonnegative, and *strictly totally positive* if all its minors are positive. For a class function f on the symmetric group S_n, its immanant is the polynomial

$$\mathrm{Imm}_f(A) = \sum_{w\in S_n} f(w)\prod_{i=1}^{n} a_{i,w(i)}, \qquad A = (a_{ij})_{i,j=1}^{n}.$$

The determinant and permanent correspond respectively to the sign and trivial characters. The immanants of the irreducible characters χ^λ, indexed by partitions λ of n, interpolate between these two examples. Stembridge proved that these irreducible immanants are nonnegative on totally nonnegative matrices; in fact,

$$\mathrm{Imm}_{\chi^\lambda}(A) \geq \chi^\lambda(1)\det A$$

for every such A [20, Corollaries 3.3–3.4]. His proof combines a factorization of totally nonnegative matrices with positivity in the group algebra of S_n.

The Frobenius characteristic identifies class functions with homogeneous symmetric functions of degree n:

$$\mathrm{ch}(f) = \sum_{\rho\vdash n} \frac{f(\rho)}{z_\rho} p_\rho, \qquad z_\rho = \prod_{j\geq 1} j^{m_j(\rho)} m_j(\rho)!. \tag{1.1}$$

Here $p_j(x) = \sum_{a\geq 1} x_a^j$ is the jth power sum, $p_\rho = \prod_i p_{\rho_i}$, $m_j(\rho)$ is the multiplicity of j in ρ, and $f(\rho)$ is the common value of f on that cycle type. We write $\lambda\vdash n$ for a weakly decreasing sequence of positive integers of sum n, and append zeros when regarding it as an n-tuple. The characteristic of χ^λ is the Schur function s_λ.

For a partition $\lambda\vdash n$, let ϕ^λ be the virtual character determined by

$$\mathrm{ch}(\phi^\lambda) = m_\lambda,$$

where m_λ is the sum of the distinct monomials whose exponent sequence is a permutation of λ, padded with zeros. The polynomial $\mathrm{Imm}_{\phi^\lambda}$ is the *monomial immanant* indexed by λ. The change of basis

$$s_\lambda = \sum_{\mu\vdash n} K_{\lambda\mu} m_\mu, \qquad \chi^\lambda = \sum_{\mu\vdash n} K_{\lambda\mu}\phi^\mu,$$

has nonnegative integral Kostka coefficients $K_{\lambda\mu}$. Thus nonnegativity of the monomial immanants implies nonnegativity of every irreducible immanant. In *Some conjectures for immanants*,

Date: November 9, 2026.
2020 *Mathematics Subject Classification.* 15A15, 05E05, 13F60.
Key words and phrases. Immanant, total nonnegativity, symmetric function, cluster variety, theta function.

Stembridge formulated this stronger assertion [21, Conjecture 2.1]. We prove it for every partition and every matrix size.

Theorem 1.1. *Let $n \geq 1$ and $\lambda \vdash n$. For every totally nonnegative real $n \times n$ matrix A,*

$$\mathrm{Imm}_{\phi^\lambda}(A) \geq 0.$$

Stembridge's terminology in [20, 21] calls a matrix totally positive when all its minors are nonnegative; this is the class called totally nonnegative here. His formulation also singles out the monomial characters among all class functions. Indeed, let J_μ be the block diagonal matrix whose blocks are all-one square matrices of orders $\mu_1, \mu_2, \ldots$. Stembridge proved [21, Proposition 2.3] that

$$\mathrm{Imm}_{\phi^\lambda}(J_\mu) = \left(\prod_i \mu_i!\right) \delta_{\lambda\mu}. \tag{1.2}$$

Here $\delta_{\lambda\mu}$ is the Kronecker delta. To see the normalization, only permutations in the Young subgroup $S_{\mu_1} \times S_{\mu_2} \times \cdots$ contribute. Their sum is the order of this subgroup times the scalar product of ϕ^λ with its induced trivial character. Under the Frobenius characteristic the latter character is $h_\mu = \prod_i h_{\mu_i}$, where $h_j(x) = \sum_{a_1 \leq \cdots \leq a_j} x_{a_1} \cdots x_{a_j}$. The Hall scalar product is determined by $\langle p_\rho, p_\sigma \rangle = z_\rho \delta_{\rho\sigma}$; the bases m_λ and h_μ are dual for this product. Since J_μ is totally nonnegative, these tests force every class function whose immanant is nonnegative on all totally nonnegative matrices to have nonnegative coefficients in the basis $\{\phi^\lambda\}_{\lambda \vdash n}$. The theorem proves that this necessary condition is sufficient.

The conjecture belongs to a group of positivity questions involving Jacobi–Trudi matrices and Hecke algebra characters. For a partition β with at most n parts, its Jacobi–Trudi matrix is

$$H_\beta = (h_{\beta_i - i + j})_{i,j=1}^n, \qquad h_0 = 1, \qquad h_j = 0 \quad (j < 0).$$

Goulden and Jackson conjectured that irreducible immanants of these matrices have nonnegative coefficients in the underlying variables [10]. Greene proved this conjecture by studying Young's seminormal representations [11]. Haiman then proved the stronger Schur positivity of $\mathrm{Imm}_{\chi^\lambda}(H_\beta) - \chi^\lambda(1) \det H_\beta$ [13, Theorem 1.4]; Schur positivity means nonnegative coefficients when the symmetric function is expanded in the Schur basis.

Haiman also introduced Hecke algebra deformations of the monomial characters and conjectured positivity of their values on normalized Kazhdan–Lusztig basis elements [13, Conjecture 2.1]. At $q = 1$, this gives a nonnegative expansion of each monomial immanant in the Kazhdan–Lusztig immanant basis. The latter immanants are nonnegative on totally nonnegative matrices [17, Proposition 2.2], so this specialization implies Stembridge's conjecture. The graded conjecture additionally asks for positivity in the Hecke parameter q. These formulations connect the matrix inequality with canonical bases and give stronger coefficient assertions for particular shapes.

Several families were known before the general result. Stembridge proved the conjecture for rectangular partitions and for $(2, 1^{n-2})$ [21, Theorems 2.7–2.8]. Clearman, Shelton, and Skandera gave nonnegative expansions in Temperley–Lieb immanants and path-tableau interpretations for partitions with $\lambda_1 \leq 2$ [2, Theorems 5.6–5.7]. More recently, Dong, Wu, and Gao proved the inequality for every hook $(a, 1^k)$, $a + k = n$, under the hypothesis that all minors of order at most $k + 1$ are nonnegative [4, Theorem 1.2]. McNulty proved Kazhdan–Lusztig positivity at $q = 1$ for all hooks [14], and separately verified Haiman's graded conjecture for $n \leq 10$ by exact computation [15].

We prove the theorem by studying the normalized spectral function $m_\lambda(\mathrm{eig}\, g)/\det g$ on PGL_n. Pulling this function back by the holonomy of a marked annulus places it in a cluster Poisson algebra. A finite sequence of mutations reduces its matrix evaluation to a tridiagonal matrix. Its

characteristic coefficients are the independent-set polynomials of a path with alternating anchor and bridge vertices.

The identification of monomial symmetric functions evaluated on triangular interval graphs with single theta sections was proved in [16, §5.4, Theorem 5.8]. We use the version of its rotations with one bridge on each level [16, §§5.1–5.2] and establish the comparison with the full annular theta basis. A finite Newton-polytope lemma shows that every extremal exponent of a finite family belongs to a unique theta function. It follows that each summand of the annular spectral function descends to the path lattice. Rotating the bridges through $n+1$ periods forces every label to be a dominant sum of anchor vectors. A finite bound on bend increments then determines the coefficient of each such label and identifies the spectral function with a single theta function. Theta positivity in the original annular chart and a simultaneous constant term give the immanant inequality.

We use the theta-function construction and positivity theorems of Gross–Hacking–Keel–Kontsevich [12] and Davison–Mandel [3], together with the surface results of Goncharov–Shen [8, 9] and Shen [19]. The finite form of valuative independence needed here is proved in Section 5. The annular charts, all-rank word reduction, lattice comparisons, and coefficient extraction are given explicitly. We use $\mathbb{N} = \mathbb{Z}_{\geq 0}$. For a Laurent polynomial f, $[X^a]f$ denotes the coefficient of the monomial X^a; the same convention applies to any explicitly named variables.

2. The Frobenius coefficient identity

If g is a complex $n \times n$ matrix, a symmetric polynomial in $\operatorname{eig} g$ means its expression in the coefficients of the characteristic polynomial of g. This convention also applies when g is not diagonalizable. For an invertible matrix put

$$F_\lambda(g) = \frac{m_\lambda(\operatorname{eig} g)}{\det g}. \tag{2.1}$$

The numerator is polynomial in the matrix entries. The denominator is invertible on GL_n, and both numerator and denominator have scalar degree n. Thus F_λ is a regular conjugation-invariant function on PGL_n.

Lemma 2.1. *Let $Z = \operatorname{diag}(z_1, \dots, z_n)$. For every $n \times n$ matrix A,*

$$[z_1 \cdots z_n] m_\lambda(\operatorname{eig}(ZA)) = \operatorname{Imm}_{\phi^\lambda}(A). \tag{2.2}$$

If A is invertible, then

$$\operatorname{Imm}_{\phi^\lambda}(A) = \det(A)\, [z_1^0 \cdots z_n^0] F_\lambda(ZA). \tag{2.3}$$

Proof. The power sum $p_j(\operatorname{eig}(ZA))$ is $\operatorname{tr}((ZA)^j)$. A term of squarefree row degree $z_1 \cdots z_n$ in a product of traces $p_\rho(\operatorname{eig}(ZA))$ specifies a permutation of cycle type ρ. Each cycle of length j has j choices of starting point in its trace, and the $m_j(\rho)$ traces of that length can be assigned to the cycles in $m_j(\rho)!$ ways. Therefore

$$[z_1 \cdots z_n] p_\rho(\operatorname{eig}(ZA)) = z_\rho \sum_{\substack{w \in S_n \\ \operatorname{type}(w) = \rho}} \prod_{i=1}^{n} a_{i,w(i)}.$$

Apply (1.1) with $f = \phi^\lambda$; the factors z_ρ cancel. This proves (2.2). For invertible A, the identity

$$\det(ZA) = (z_1 \cdots z_n) \det(A)$$

then gives (2.3) by division by the indicated row monomial. □

It remains to establish positive Laurent coefficients for F_λ in a chart whose row weights can be read off directly. We assume $n \geq 2$ until the proof of Theorem 1.1 is completed.

3. An annular chart

Assume $n \geq 2$, put $r = n - 1$, and let $G = \mathrm{PGL}_n(\mathbb{C})$. Let S be an oriented annulus with one marked point on each boundary component. We use its moduli space $\mathcal{X}_{G,S}$ of framed local systems: a flag is specified at each marked point, and no boundary pinning is specified. Write B^+, B^- for the opposite standard Borel subgroups, $H = B^+ \cap B^-$, and w_0 for the longest Weyl group element. Choose a representative $\dot{w}_0 \in G$.

3.1. **The surface chart and its character lattice.** Cutting the annulus along an arc between its marked points and transporting the two flags to a common fiber gives the description

$$(g, F_0, F_1)/G,$$

where g is the holonomy around the oriented core. On the locus where F_0, F_1 are opposite, normalize them to the standard pair. Their stabilizer is H, which acts on g by conjugation. The two boundary genericity conditions say that a standard flag and its image under g are opposite, for each of the two standard flags. Thus the resulting open part of the moduli space is

$$G^{w_0,w_0}/\operatorname{Ad} H, \qquad G^{w_0,w_0} = B^+\dot{w}_0B^+ \cap B^-\dot{w}_0B^-. \tag{3.1}$$

The double-word coordinates below parametrize a dense seed torus in this quotient.

We recall how the surface cluster coordinates restrict to this chart. For a root $i \in \{1, \dots, r\}$ put

$$E_i = I + e_{i,i+1}, \qquad F_i = I + e_{i+1,i}, \qquad H_i(t) = \operatorname{diag}(\underbrace{t, \dots, t}_{i}, 1, \dots, 1).$$

The matrix units are denoted by e_{ij}; the displayed matrices represent elements of G. A signed word uses E_i for the letter i and F_i for the letter $-i$. Its coordinates are attached to intervals between successive occurrences of each root, with an additional interval at each end. Evaluation inserts H_i with the corresponding coordinate in each root-i interval. Since H_i commutes with root factors of every other root, its position within that interval is immaterial. These are the double-word coordinates of [5, 9].

The compatibility with surface gluing includes an orientation convention which we shall need. If p, q are the boundary-oriented pinnings of the cut rectangle and q^* is the opposite pinning obtained by reversing the boundary orientation, its transport is defined by $q^* = gp$ in [9, Theorem 5.3, (198)]. Consequently, for $g = UhL$ with U upper unipotent, $h \in H$, and L lower unipotent, the four flags in cyclic order are

$$(B^+, B^-, gB^-, gB^+), \qquad gB^- = UB^-.$$

The two triangles cut out by the diagonal have the one-sign word coordinates. Amalgamating their diagonal coordinates multiplies matching coordinates and adds the exchange forms, giving the double-word product g; see [9, Theorems 5.3–5.4 and §13.1]. Gluing the remaining sides identifies their pinnings up to H, giving precisely (3.1). In particular, the holonomy in this convention is g itself. The root label at the right endpoint of a word is $i^* = n - i$; the boundary gluing reverses labels by the same involution [9, §§13.1.2–13.1.3, (503)–(506)]. It therefore joins the two endpoints of the same word-root i.

After this gluing, the two extremal root-i coordinates are replaced by their product; all internal coordinates remain. These are coordinates on an integral torus, as follows directly from the seam action. Its weights on the two extremal coordinates are $\alpha_i, -\alpha_i$, where $\alpha_i(\operatorname{diag}(h_1, \dots, h_n)) =$

h_i/h_{i+1}. The simple roots form a $\mathbb{Z}$-basis of $X^*(H)$ because G is adjoint. An invariant character has equal exponents at the two extremities of each root. Hence the endpoint products, together with the internal coordinates, form a basis of the invariant character lattice. In the resulting cyclic word there is one coordinate for each interval between successive occurrences of a root, including the interval across the cut. All these coordinates are mutable in the unpinned surface seed: the frozen coordinates on the two original boundary edges are deleted [9, §13, opening paragraphs]. Each of the two triangles has $n(n-1)/2 - r$ interior coordinates, and the diagonal and the glued seam each have r coordinates. Thus the full seed has $n(n-1) = r(r+1)$ mutable coordinates.

3.2. **Positive matrix parameters.** For prescribed reduced expressions of w_0 in the upper and lower blocks, every strictly totally positive real matrix has a factorization

$$A = U(\mathbf{u})\,\mathrm{diag}(t_1, \dots, t_n)L(\boldsymbol{\ell}), \qquad u_a, \ell_b, t_j > 0, \tag{3.2}$$

where the elementary factors in U, L occur in the prescribed orders [7, Theorem 23]. The cyclic coordinates of this factorization are obtained from the identities

$$I + ue_{i,i+1} = H_i(u)E_iH_i(u^{-1}), \qquad I + \ell e_{i+1,i} = H_i(\ell^{-1})F_iH_i(\ell),$$

and $\mathrm{diag}(t_1, \dots, t_n)$ represents $\prod_i H_i(t_i/t_{i+1})$ in G. Collecting the Cartan factors of root i in (3.2) gives, in cyclic order,

$$\frac{u_{\mathrm{next}}}{u_{\mathrm{previous}}}, \qquad \frac{t_i/t_{i+1}}{u_{\mathrm{last}}\ell_{\mathrm{first}}}, \qquad \frac{\ell_{\mathrm{previous}}}{\ell_{\mathrm{next}}}, \qquad o_i = \ell_{\mathrm{last}}u_{\mathrm{first}}. \tag{3.3}$$

The first and third expressions are taken over successive occurrences in the corresponding sign block. The initial factor $H_i(u_{\mathrm{first}})$ and the final factor $H_i(\ell_{\mathrm{last}})$ combine across the cyclic cut to give o_i. Thus every source coordinate is positive. We call o_i the *outer coordinate* of root i; the remaining coordinates are internal.

4. Reduction to a path

We next make the holonomy depend on $2r$ coordinates of a mutated annular seed. The other coordinates will remain mutable coordinates of the full seed. The reduction is a version of the cyclic word construction in [18, §§5.4–5.6]; we give the word order and the local identities that produce the path needed here.

4.1. **Local word transformations.** For a seed basis (e_v) write $\mathsf{B}_{uv} = \Omega(e_u, e_v)$. Our convention for X-mutation at k is

$$X_k' = X_k^{-1}, \qquad X_v' = X_v(1 + X_k^{-\operatorname{sgn}\mathsf{B}_{vk}})^{-\mathsf{B}_{vk}} \quad (v \neq k). \tag{4.1}$$

The exchange matrix is mutated by changing the signs in row and column k and, for $u, v \neq k$, by

$$\mathsf{B}'_{uv} = \mathsf{B}_{uv} + [\mathsf{B}_{uk}]_+[\mathsf{B}_{kv}]_+ - [-\mathsf{B}_{uk}]_+[-\mathsf{B}_{kv}]_+, \qquad [a]_+ = \max(a, 0).$$

The word seed is obtained by adding elementary tile forms. For a positive root-i tile, let L, R be its two root-i endpoints, and let j denote the interval of a neighboring root through the tile. Its nonzero entries, up to skew-symmetry, are

$$\mathsf{B}_{L,R} = -1, \qquad \mathsf{B}_{L,j} = \frac{1}{2}, \qquad \mathsf{B}_{R,j} = -\frac{1}{2}. \tag{4.2}$$

A negative tile has the opposite form. This is the elementary word form of [9, (416)]. The same rule defines its amalgamation on a cyclic word. For two cyclic intervals of neighboring roots,

the number of interval endpoints contributing to their pairing is even. Their half-integral tile contributions therefore sum to an integer.

We shall use commutations, braids, mixed swaps, and merges. Opposite signs of distinct roots commute, as do equal signs of nonadjacent roots. These are just commutations of the corresponding elementary matrices. The remaining transformations are single mutations, as the following identities show.

First, with $c = x/(1+x)$, a merge is given by

$$\begin{aligned} E_i H_i(x) E_i &= H_i(1+x) E_i H_i(c), \\ F_i H_i(x) F_i &= H_i(c) F_i H_i(1+x). \end{aligned} \tag{4.3}$$

The middle column has entries $(-1,+1)$ at the two exterior root-i intervals in the positive case, and the opposite entries in the negative case. Formula (4.1) gives the two Cartan factors on the right of (4.3). The mutated middle coordinate is x^{-1} and does not occur in the shorter product.

For a braid with $|i-j|=1$, set

$$a' = a(1+x), \quad d' = \frac{dx}{1+x}, \quad b' = \frac{bx}{1+x}, \quad c' = c(1+x).$$

Then

$$\begin{aligned} &H_i(a) H_j(b) E_i H_i(x) E_j H_j(c) E_i H_i(d) \\ &\quad = H_i(a') H_j(b') E_j H_j(x^{-1}) E_i H_i(d') E_j H_j(c'). \end{aligned} \tag{4.4}$$

The mutating column on (i_L, i_R, j_L, j_R) is $(-1,+1,+1,-1)$. Replacing the E's by F's negates this column; the same identity holds with

$$a' = \frac{ax}{1+x}, \quad d' = d(1+x), \quad b' = b(1+x), \quad c' = \frac{cx}{1+x}.$$

To check both identities, factor out $H_i(a)H_j(b)$ on the left and $H_i(d)H_j(c)$ on the right. The two positive products have the common 3×3 blocks

$$\begin{pmatrix} x & x+1 & 1 \\ 0 & 1 & 1 \\ 0 & 0 & 1 \end{pmatrix} \quad (j=i+1), \qquad \begin{pmatrix} x & x & x \\ 0 & x & x+1 \\ 0 & 0 & 1 \end{pmatrix} \quad (j=i-1).$$

The negative products give the transposes of these blocks. Above the affected block the common diagonal factor is x, and below it the common factor is 1.

Finally, define $H_0(t) = I$ and $H_n(t) = tI$. The two mixed identities, valid as identities of GL_n matrices, are

$$\begin{aligned} E_i H_i(x) F_i &= H_{i-1}(c) H_i(1+x) F_i H_i(x^{-1}) E_i H_i(1+x) H_{i+1}(c), \\ F_i H_i(x) E_i &= H_{i-1}(1+x) H_i(c) E_i H_i(x^{-1}) F_i H_i(c) H_{i+1}(1+x), \qquad c = \frac{x}{1+x}. \end{aligned} \tag{4.5}$$

Their affected 2×2 blocks are respectively

$$\begin{pmatrix} x+1 & 1 \\ 1 & 1 \end{pmatrix}, \qquad \begin{pmatrix} x & x \\ x & x+1 \end{pmatrix}.$$

Above the block both sides have diagonal factor x, since $c^2(1+x)^2/x = x$, and below it both have factor 1. For $E_iF_i \longmapsto F_iE_i$, the middle column is $(-1,-1)$ on the exterior root-i intervals and $+1$ on each neighboring-root interval. The reverse move negates this column. Thus (4.5) gives precisely (4.1). When $i=r$, the factor H_n is scalar and is omitted in G. The function F_λ is invariant under this scalar factor.

The exchange forms can be checked on the same local intervals. Write $b_u = \mathsf{B}_{uk}$ for the displayed mutating column. Before cyclic gluing its nonzero entries are 1 or -1, and for $u, v \neq k$ the change is

$$\mathsf{B}'_{uv} - \mathsf{B}_{uv} = \begin{cases} 1, & b_u = 1,\ b_v = -1, \\ -1, & b_u = -1,\ b_v = 1, \\ 0, & \text{otherwise.} \end{cases}$$

For a positive merge this changes the pairing of the two exterior root-i intervals by -1, giving a single positive tile. For a positive braid, in the order (i_L, i_R, j_L, j_R), the changed upper-triangular entries are $-1, -1, +1, +1$ on $(i_L, i_R), (i_L, j_L), (i_R, j_R), (j_L, j_R)$, respectively. For a mixed swap the change is -1 from each exterior root-i interval to each neighboring-root interval. Negating a local column negates these changes. Together with the sign reversal of the mutating row and column, these are the sums of the new tiles in (4.2). Entries outside the column's support are unchanged. At a merge the shorter word uses the principal submatrix on the remaining intervals.

4.2. **Cyclic gluing and retained coordinates.** The local transformations descend to cyclic words by the following elementary amalgamation observation. Replace coordinates in an identified fiber by their product and exchange entries by their sums. Mutation at a singleton fiber commutes with this operation provided that the entries from each other fiber to the mutating vertex have one weak sign. For the coordinates, multiply the factors in (4.1). For the exchange matrix, use $[\sum b_j]_+ = \sum [b_j]_+$ for numbers of one weak sign in the displayed mutation formula.

In all the words below, each root occurs with each sign. For a same-sign braid or merge, an occurrence of the opposite sign lies outside the local patch. It separates the two exterior intervals of that root. In a braid this holds for both roots, so exterior entries of opposite signs are not identified. In a mixed swap, the two exterior root-i intervals can coincide, but their entries have the same sign; the resulting entry is 2 or -2. The mutating interval is a singleton. The amalgamation observation therefore applies to every local transformation. Changing the cut of a cyclic word replaces a product AB by $BA = A^{-1}(AB)A$ and preserves its holonomy class.

To keep track of the full seed, call an interval *active* if it occurs in the current word evaluation. Initially all intervals are active. At a merge, perform the mutation in the full seed and then mark its middle interval inactive. Its coordinate still belongs to the full mutable seed. Inductively, the active principal submatrix is the current word matrix and the holonomy class depends only on active coordinates. Indeed, a later mutation at an active vertex updates active coordinates using only that coordinate and the active principal submatrix; matrix mutation restricts to the same submatrix. The inactive coordinates and all mixed entries are simultaneously updated by the full mutation. This proves the induction statement at every stage.

4.3. **The word reduction.** For integers a, b write $[a \downarrow b] = (a, a-1, \dots, b)$, with the empty word understood when $a < b$, and put

$$S_b = [r \downarrow b], \qquad W_r = S_1 S_2 \cdots S_r, \qquad C_r = (1, \dots, r), \qquad \rho(i) = r + 1 - i.$$

Both W_r and $\rho(W_r)$ are reduced words for w_0. We choose the source word

$$\mathsf{D}_0 = (\rho(W_r), -W_r), \tag{4.6}$$

where the minus sign changes each letter without reversing its order. Formula (3.2) applies to this prescribed source.

Lemma 4.1. *There is a finite sequence of mutations of the full annular seed* (4.6) *for which the active cyclic word is*

$$E_rE_{r-1}\cdots E_1F_1F_2\cdots F_r. \tag{4.7}$$

The holonomy conjugacy class is given by its cyclic evaluation. The full seed has $r(r+1)$ *mutable coordinates, of which* $2r$ *are active at the endpoint.*

Proof. First consider $(\mathbf{b}, W_r)$, where $\mathbf{b}$ is any all-negative word containing every root. We reduce its positive block to C_r, restoring the order of $\mathbf{b}$ after each step. The intermediate positive words are

$$T_{j,k} = (1,\dots,j)[r-k\downarrow j+1]S_{j+2}\cdots S_r, \qquad 0\le j\le r-1, \quad 0\le k\le r-j-1.$$

Here $T_{0,0}=W_r$, $T_{j,r-j-1}=T_{j+1,0}$ for $j<r-1$, and $T_{r-1,0}=C_r$. Suppose $k\le r-j-2$ and set $a=r-k$. Move the first displayed letter a to the left through $1,\dots,j$, which is possible because $a\ge j+2$. Move it through $\mathbf{b}$ by opposite-sign commutations and mixed equal-root swaps. Rotate this initial letter through the cut to the end of the cyclic word, and commute it left through $S_{a+2},\dots,S_r$. It remains to perform the absorption

$$S_aS_{a+1}a \longrightarrow S_aS_{a+1}. \tag{4.8}$$

This absorption follows by induction on $r-a$. For $a=r$ it is the merge $rr\longrightarrow r$. For $a<r$, commutations and a single braid give

$$\begin{aligned} S_aS_{a+1}a &= S_{a+1}aS_{a+2}(a+1)a \\ &\longleftrightarrow S_{a+1}S_{a+2}a(a+1)a \\ &\longleftrightarrow S_{a+1}S_{a+2}(a+1)a(a+1). \end{aligned}$$

Apply the induction hypothesis to the initial $S_{a+1}S_{a+2}(a+1)$, then commute a back through S_{a+2}. The result is S_aS_{a+1}. The whole step has therefore changed $(\mathbf{b},T_{j,k})$ into $(\mathbf{b},T_{j,k+1})$. The displayed finite set of indices proves termination at $(\mathbf{b},C_r)$.

Apply the Dynkin-reflected version of this procedure to $(-W_r,\rho(W_r))$, obtained from (4.6) by a change of cut. This gives $(-W_r,(r,\dots,1))$. Change the cut again and use the sign-reversed versions of the same local identities to reduce the negative block. The result is $(r,\dots,1,-1,\dots,-r)$, as required.

For clarity, the procedure also determines the mutation vertices. Attach to each signed letter the label of the interval of its root immediately following it. A commutation interchanges the two letter-label pairs. The other update rules are

$$\begin{aligned} (a,u),(-a,v) &\longmapsto (-a,u),(a,v), && \text{mixed swap},\\ (a,u),(b,v),(a,w) &\longmapsto (b,u),(a,w),(b,v), && \text{braid},\\ (a,u),(a,v) &\longmapsto (a,v), && \text{merge}. \end{aligned}$$

In each row the mutated interval is u; in the last row it is retained as an inactive mutable coordinate. These rules are exactly the coordinate placements in (4.3)–(4.5). Each root retains at least one occurrence of each sign throughout, so the cyclic amalgamation conditions continue to hold. The preceding full-seed induction lifts every step to an annular mutation. Finally, the initial word has $r(r+1)$ letters and the terminal word has $2r$; a cyclic word has one interval per letter. All initial coordinates remain in the full seed. □

4.4. **The terminal determinant and its lattice.** Denote the active interval following E_i by x_i and the one following F_i by y_i. Thus the terminal product is

$$E_rH_r(x_r)\cdots E_1H_1(x_1)F_1H_1(y_1)\cdots F_rH_r(y_r).$$

Set

$$d_n = 1, \qquad d_i = x_i y_i d_{i+1}, \qquad b_i = y_i d_{i+1} \quad (i < n). \tag{4.9}$$

Moving the Cartan factors to the middle gives the matrix

$$M = \left(I + \sum_{i=1}^{r} e_{i,i+1}\right) \operatorname{diag}(d_1, \ldots, d_n) \left(I + \sum_{i=1}^{r} y_i e_{i+1,i}\right). \tag{4.10}$$

The descending upper order and ascending lower order make the outer factors bidiagonal. In particular, M is tridiagonal, with diagonal entries $d_i + b_i$ for $i < n$ and d_n for $i = n$, and

$$M_{i,i+1} M_{i+1,i} = b_i d_{i+1}.$$

For a weighted graph P, let $\mathcal{E}_k(P)$ be the sum of the weight products of its independent sets of size k.

Lemma 4.2. *Let P be the path with successive weights $d_1, b_1, d_2, b_2, \ldots, b_r, d_n$. Then*

$$\det(I + tM) = \sum_{k=0}^{n} \mathcal{E}_k(P) t^k, \qquad \det M = \prod_{i=1}^{n} d_i. \tag{4.11}$$

In the active order $(x_1, y_1, \ldots, x_r, y_r)$ the exchange form is

$$\mathsf{B}_{j,j+1} = 2, \qquad \mathsf{B}_{j,j+2} = -1, \qquad \mathsf{B}_{j,k} = 0 \quad (k > j + 2). \tag{4.12}$$

Its determinant is n^2.

Proof. Let $I_i(t)$ be the independent-set generating polynomial of the tail path starting at d_i, and let $J_i(t)$ be that of the tail starting at b_i, with $J_n = I_{n+1} = 1$. Splitting at the first vertex gives

$$I_i = J_i + t d_i I_{i+1}, \qquad J_i = I_{i+1} + t b_i J_{i+1}.$$

Eliminating $J_{i+1} = I_{i+1} - t d_{i+1} I_{i+2}$ yields

$$I_i = (1 + t(d_i + b_i)) I_{i+1} - t^2 b_i d_{i+1} I_{i+2}, \qquad I_n = 1 + t d_n, \quad I_{n+1} = 1.$$

The trailing principal minors of $I + tM$ satisfy the same recurrence and initial conditions. This proves the first equality in (4.11). The unique independent set of size n consists of the n vertices d_i, giving the second equality.

For the exchange form, the positive and negative tiles of root i each contribute 1 to $\mathsf{B}(x_i, y_i)$. For consecutive roots, adding their endpoint contributions from (4.2) gives

(u, v)	(x_i, x_{i+1})	(y_i, y_{i+1})	(y_i, x_{i+1})	(x_i, y_{i+1})
$\mathsf{B}(u, v)$	-1	-1	2	$0.$

Nonadjacent roots have zero pairing. This is (4.12). Let p_r be its Pfaffian. For $r \geq 2$, the first row has entries $2, -1$ in columns $2, 3$. Deleting rows and columns $1, 2$ leaves the same band matrix of size $2r - 2$, contributing $2p_{r-1}$. After deleting $1, 3$, the first remaining vertex 2 has its sole nonzero entry $\mathsf{B}_{2,4} = -1$, so that minor has Pfaffian $-p_{r-2}$. The Pfaffian expansion therefore gives

$$p_r = 2p_{r-1} - p_{r-2}, \qquad p_0 = 1, \quad p_1 = 2.$$

Thus $p_r = r + 1 = n$ and $\det \mathsf{B} = p_r^2 = n^2$. □

The ratios of successive path weights are $x_i = d_i / b_i$ and $y_i = b_i / d_{i+1}$. We reverse the path, so that its forward ratios, in order, are $(y_r, x_r, \ldots, y_1, x_1)$. This is a permutation of the actual active coordinates.

Introduce the free vertex lattice

$$L = \mathbb{Z}D_1 \oplus \mathbb{Z}B_1 \oplus \mathbb{Z}D_2 \oplus \cdots \oplus \mathbb{Z}B_r \oplus \mathbb{Z}D_n.$$

We call D_i the anchor vectors and B_i the bridge vectors; their vertex weights are the characters X^{D_i} and X^{B_i}. They homogenize the reversed weights d_{n+1-i}, b_{n-i}: specializing $X^{D_1} = 1$ recovers (4.9). Give every vertex weight one, let $L_0 = \ker(\mathrm{wt} : L \to \mathbb{Z})$, and put

$$c_{2i-1} = B_i - D_i, \qquad c_{2i} = D_{i+1} - B_i \quad (1 \le i < n). \tag{4.13}$$

These forward differences form a $\mathbb{Z}$-basis of L_0. Together with D_1 they form a $\mathbb{Z}$-basis of L; in this auxiliary vertex seed the gaps are mutable and D_1 is frozen. More explicitly, if $V_1, \ldots, V_{2n-1}$ denote the vertices in the displayed order, then for $\sum_j a_j = 0$,

$$\sum_{j=1}^{2n-1} a_j V_j = -\sum_{j=1}^{2n-2} \left(\sum_{h=1}^{j} a_h \right) c_j.$$

In particular, a vector of L_0 belongs to $\sum_j \mathbb{N}c_j$ precisely when all its proper vertex prefix sums are nonpositive.

Define Ω on L, extending by skew-symmetry, by $\Omega(V_{j+1}, V_j) = 1$ and zero on nonadjacent vertices. Its restriction to the ordered gaps in (4.13) has first superdiagonal -2 and second superdiagonal 1. It therefore agrees with (4.12) after the preceding permutation. The character map

$$\iota(y_{n-i}) = X^{c_{2i-1}}, \qquad \iota(x_{n-i}) = X^{c_{2i}} \tag{4.14}$$

identifies the actual active character lattice with L_0 integrally. The vector

$$\delta = D_1 + \cdots + D_n \tag{4.15}$$

is radical: each bridge pairs with its two neighboring anchors with opposite signs, and the anchors pair trivially with each other. Since the restriction to L_0 is nondegenerate, $\mathbb{R}\delta$ is the entire radical. Since $\mathrm{wt}(\delta) = n$, the sublattice $L_0 + \mathbb{Z}\delta$ has index n in L. We retain L_0 as the active integral lattice.

From now on P denotes the reversed path with the formal weights $X^{D_1}, X^{B_1}, \ldots, X^{B_r}, X^{D_n}$. Finally, let Y denote its formal spectral alphabet, defined by $e_k(Y) = \mathcal{E}_k(P)$ for $1 \le k \le n$, where e_k is the kth elementary symmetric polynomial. The map (4.14) satisfies

$$\iota(d_{n+1-i}) = X^{D_i - D_1}, \qquad \iota(b_{n-i}) = X^{B_i - D_1}.$$

Thus the common factor X^{-D_1} cancels between the degree-n numerator and denominator of F_λ. Since $\mathcal{E}_n(P) = X^\delta$, (4.11) gives

$$\iota(F_\lambda) = X^{-\delta} m_\lambda(Y). \tag{4.16}$$

Both sides have vertex weight zero, so this is an identity in the active Laurent ring.

5. Theta functions and finite Newton polytopes

We use the classical theta functions of Gross–Hacking–Keel–Kontsevich [12], with the X-realization and positivity results of Davison–Mandel [3]. We first specify the realization, since its chamber convention and its coefficient grading will both enter the proof.

5.1. **The chamber convention.** Let M be the character lattice of a full unpinned seed, with basis $e_1, \ldots, e_d$ and integral skew form Ω. We write the matrix of the form in this basis also as

Ω, so that $\Omega(v,w) = v^{\mathsf{T}}\Omega w$. Put

$$C = \sum_{i=1}^{d} \mathbb{Z}_{\geq 0} e_i.$$

The positive root degree is recorded by variables $y_1, \ldots, y_d$. The monomial belonging to a root $q \in C$ is $x^q y^q$; in particular, the root map on the character lattice is the identity.

It is useful to allow the form Ω to be degenerate. For this purpose use the position lattice

$$\widehat{M} = M \oplus \mathbb{Z}^d, \qquad \widehat{P}(v) = (v,0), \qquad \widehat{Q}(v) = (\Omega^{\mathsf{T}} v, v) \in \widehat{M}^*, \tag{5.1}$$

where coordinate bases identify the displayed dual lattices. The initial wall of root e_i has normal $\widehat{Q}(e_i)$ and factor $1 + x^{e_i} y_i$. Labels and root increments lie in $E = M \oplus 0$. The second summand of $\widehat{M}$ is used for wall positions; the variables y_i separately record root degrees. The initial negative chamber is nonempty, since $(0, -\mathbf{1})$ has negative pairing with each initial normal. We always choose an endpoint Q in this chamber such that

$$Q \notin H \quad \text{for every proper rational linear hyperplane } H \subset \widehat{M}_{\mathbb{R}}. \tag{5.2}$$

Such points are dense in every nonempty open set.

Here this negative chamber gives the usual seed expression of an X-theta function. To check this integrally, consider

$$T(m,a) = (\Omega^{\mathsf{T}} m - a, m), \qquad T^{-1}(u,v) = (v, \Omega^{\mathsf{T}} v - u). \tag{5.3}$$

The standard principal A-datum has root and normal maps

$$P_A(v) = (\Omega^{\mathsf{T}} v, v), \qquad Q_A(v) = (v,0).$$

Directly,

$$T\widehat{P} = P_A, \qquad T^*(-Q_A) = \widehat{Q}.$$

Reversing a normal reverses both its crossing orientation and the sign of its pairing in the wall action, so it leaves that action unchanged. Thus our negative chamber corresponds to the positive chamber for the principal A-datum. Moreover,

$$T(\nu, 0) = (\Omega^{\mathsf{T}} \nu, \nu)$$

is precisely the integral embedding which defines the X-theta functions in [3, §§4.3–4.4 and 4.7.2]. This integral comparison fixes the labels and the chamber convention also for a degenerate full form. The construction in [3, §4.7.2] permits generic endpoints in the full principal position space and identifies its chamber charts with the usual X-cluster charts.

We recall the properties of positive scattering used below. Its wall supports are rational cones in central hyperplanes, and each joint cone has a rational linear span of codimension at least two. Modulo any fixed positive root degree, only finitely many walls are nontrivial. A broken line is a piecewise linear path $\Gamma : (-\infty, 0] \to \widehat{M}_{\mathbb{R}}$, ending at Q, whose velocity on a segment with attached exponent $m \in E$ is $-m$. It has finitely many nontrivial bends. At a joint-free crossing, the next monomial is chosen from the wall action on the preceding monomial. The theta function sums the resulting terminal monomials with their coefficients.

The positive scattering theorem gives a representative in which these coefficients are nonnegative integers. In the chamber just fixed, the principal expression is pointed:

$$\Theta_\nu(x,y) = x^\nu \sum_{q \in C} a_{\nu,q} x^q y^q, \qquad a_{\nu,q} \in \mathbb{Z}_{\geq 0}, \qquad a_{\nu,0} = 1. \tag{5.4}$$

Strong positivity also gives, coefficientwise in the principal completion,

$$\Theta_{k\nu} \;\le\; \Theta_\nu^k \qquad (k \in \mathbb{Z}_{>0}). \tag{5.5}$$

Give $x^m y^q$ the degree $m - q \in M$. Then Θ_ν is homogeneous of degree ν, so a product term indexed by $\nu + \mu + r$ carries the factor y^r. The positive product expansion is

$$\Theta_\nu \Theta_\mu = \Theta_{\nu+\mu} + \sum_{0 \neq r \in C} \alpha_r y^r \Theta_{\nu+\mu+r}, \qquad \alpha_r \in \mathbb{N}.$$

Iterating its leading term proves (5.5). We use the classical specializations of [3, Theorems 3.9–3.10]; their compatibility with the X-construction follows from [3, Lemma 4.6 and Proposition 4.8].

5.2. **Finite principal expressions.** Suppose that the actual X-theta function ϑ_ν is a Laurent polynomial in the chosen seed. Then Θ_ν in (5.4) is finite as well. To see this, remove the coefficient monomial y^q on the displayed support graph. Distinct q have distinct character exponents $\nu + q$, so this is an injective coefficientwise map to the completion in the cone C. By the preceding principal comparison, its image is the formal seed expansion of ϑ_ν. Laurentness of the latter therefore forces all but finitely many $a_{\nu,q}$ to vanish. After this observation one can write

$$\vartheta_\nu(x) = \Theta_\nu(x, 1)$$

as an ordinary substitution in a finite polynomial.

For a nonzero Laurent polynomial f write $\mathrm{Newt}(f)$ for the convex hull of its character exponents, and put

$$v_w(f) = \min_{a \in \mathrm{Supp}(f)} \langle w, a \rangle \qquad (w \in M^*_{\mathbb{R}}).$$

The following finite form of valuative independence will be needed. It is a special case of [1, Theorem 3.10]. We give a proof from positive scattering and (5.5).

Lemma 5.1. *Let ν_s, $s \in S$, be a nonempty finite set of distinct theta labels whose seed expressions ϑ_{ν_s} are Laurent polynomials. If $c_s \in \mathbb{C}^\times$, then $\sum_{s\in S} c_s \vartheta_{\nu_s} \neq 0$ and*

$$\mathrm{Newt}\left(\sum_{s\in S} c_s \vartheta_{\nu_s}\right) = \mathrm{conv} \bigcup_{s\in S} \mathrm{Newt}(\vartheta_{\nu_s}). \tag{5.6}$$

Equivalently, for every $w \in M^*_{\mathbb{R}}$,

$$v_w\left(\sum_{s\in S} c_s \vartheta_{\nu_s}\right) = \min_{s\in S} v_w(\vartheta_{\nu_s}). \tag{5.7}$$

Proof. All principal expressions in the statement are finite by the preceding argument. We work at the single endpoint Q chosen in (5.2). We extend each covector on $M_{\mathbb{R}}$ to $\widehat{M}_{\mathbb{R}}$ by zero on the auxiliary summand.

Rational broken lines. At a joint-free crossing, combine the wall factors having the same supporting hyperplane. Choose its primitive normal ℓ so that $\ell(m) > 0$ for the incoming exponent m. If an individual normal has absolute multiplicity b relative to ℓ, include its factor to the power b. The resulting positive factor f, with constant term one, has crossing rule

$$x^m \longmapsto x^m f^{\ell(m)}.$$

Every exponent in f lies in $\ker \ell$. Coefficient degrees are understood throughout. A *rational bend* is a choice

$$m' = m + \ell(m) r, \qquad r \in \mathrm{conv}_{\mathbb{Q}}\, \mathrm{Supp}_x(f), \tag{5.8}$$

where $\mathrm{Supp}_x(f)$ is the set of character exponents occurring in f, and $\mathrm{conv}_{\mathbb{Q}}$ is the set of their finite rational convex combinations. Thus $\ell(r) = 0$, and the bend is the shear

$$T_r = I + r \otimes \ell, \qquad T_r^{-1} = I - r \otimes \ell.$$

It fixes its supporting hyperplane pointwise. An ordinary bend is a rational bend, since an exponent of $f^{\ell(m)}$ is a sum of $\ell(m)$ exponents of f.

Consider a rational broken line with initial exponent ν and finitely many nontrivial bends. At each bend choose a finite expression

$$r_i = \sum_j b_{ij} q_{ij}, \qquad b_{ij} \in \mathbb{Q}_{\geq 0}, \quad \sum_j b_{ij} = 1, \quad q_{ij} \in \mathrm{Supp}_x(f_i).$$

Choose a positive integer k clearing all denominators of the slopes and of the numbers $\ell_i(m_{i-1})b_{ij}$. The positive expansion of $f_i^{k\ell_i(m_{i-1})}$ then contains the required increment $k\ell_i(m_{i-1})r_i$. Multiplying all slopes by k and rescaling time turns the path into an ordinary broken line with initial exponent $k\nu$. If m was its original terminal exponent, it follows that

$$km \in \mathrm{Supp}(\vartheta_{k\nu}) \subseteq \mathrm{Supp}(\vartheta_\nu^k), \qquad m \in \mathrm{Newt}(\vartheta_\nu). \tag{5.9}$$

Here (5.5) first makes $\Theta_{k\nu}$ finite, because Θ_ν^k is finite. Consequently all the supports used in (5.9) are finite. Positivity shows that the support of ϑ_ν^k consists of sums of k members of $\mathrm{Supp}(\vartheta_\nu)$, which proves the last inclusion.

If f_i is an infinite product of parallel wall factors, each monomial used in the displayed convex combinations requires only finitely many factors, by local finiteness in each root degree. The remaining factors contribute their constant term. Thus the same denominator-clearing argument applies. In particular, an ordinary line whose terminal exponent minimizes a covector on ϑ_ν also minimizes it among all rational lines with the same initial label and endpoint.

Deformation with a fixed endpoint. Let a minimizing line have successive exponents $m_0, \ldots, m_\ell$ and normalized bends r_i, with

$$T_i = I + r_i \otimes \ell_i, \qquad m_i = T_i m_{i-1}, \qquad \ell_i(m_{i-1}) > 0.$$

When testing a crossing at which the line is straight, include it in this finite list with $r_i = 0$. Put

$$A_i = T_\ell \cdots T_{i+1}, \qquad A_\ell = I.$$

Suppose that at crossing j an admissible r_j' has a smaller pairing than r_j with the transported covector $A_j^* w$. Replace r_j by $r_j(s) = (1-s)r_j + sr_j'$, where $s > 0$ is rational, and leave all other normalized bends unchanged. Starting with the same m_0, compute the new exponents forward. Starting with the same endpoint, compute the bend positions backward:

$$\begin{aligned} z_{\ell+1}(s) &= Q, \\ \tau_i(s) &= -\frac{\ell_i(z_{i+1}(s))}{\ell_i(m_i(s))}, \\ z_i(s) &= z_{i+1}(s) + \tau_i(s) m_i(s) \qquad (i = \ell, \ldots, 1). \end{aligned} \tag{5.10}$$

At $s = 0$ the denominators are nonzero, the segment lengths τ_i are positive, and the specified wall factors contain the bend points in the relative interiors of their supports. For all sufficiently small s these finitely many conditions persist. Choose beforehand the finitely many factors witnessing the convex combinations for the r_i and r_j'. Their support conditions persist as well. Extra factors and all unprescribed crossings are traversed using constant terms. The initial geometric ray may move; its exponent remains m_0.

It remains to check admissibility for the entire scattering diagram, which can have infinitely many joints. For rational s, the shears

$$A_i(s) = T_\ell(s) \cdots T_{i+1}(s)$$

are rational invertible maps. They unfold the spatial segments to one straight line. More precisely, with time zero at the endpoint, on the ith segment

$$A_i(s)\Gamma_s(t) = Q - tm_\ell(s). \tag{5.11}$$

Indeed, all transformed velocities equal $-m_\ell(s)$, and the transformed segments join because $T_i(s)$ fixes the ith bend hyperplane pointwise.

If a segment met a joint cone J, then

$$Q \in A_i(s)\,\mathrm{span}_{\mathbb{R}}(J) + \mathbb{R}m_\ell(s).$$

The right side is a proper rational linear subspace of $\widehat{M}_{\mathbb{R}}$: $\mathrm{span}_{\mathbb{R}}(J)$ has codimension at least two, and only one direction has been added. This contradicts (5.2). If a segment lay in a wall hyperplane H, its direction would satisfy $m_i(s) \in H$, and hence $m_\ell(s) = A_i(s)m_i(s) \in A_i(s)H$. Equation (5.11) would then imply $Q \in A_i(s)H$, again a proper rational hyperplane. Thus all crossings are transverse and the perturbed path avoids every joint of the full diagram for every sufficiently small rational s.

The later shears and the incoming exponent at crossing j were fixed. Hence the terminal exponent changes by

$$m_\ell(s) - m_\ell = s\,\ell_j(m_{j-1})A_j(r'_j - r_j).$$

Its pairing with w is strictly negative, contradicting (5.9). We have proved that a minimizing line makes a minimizing normalized choice at every crossing, including its straight crossings. Call a line with this property *taut*.

Backward uniqueness. Choose w so that its restriction to E is nonzero on every nonzero rational vector of E. Every rational invertible shear preserves this property on E. A local minimum of a transported covector on $\mathrm{conv}_{\mathbb{Q}}\,\mathrm{Supp}_x(f)$, when attained, is therefore attained at a unique point r_*. For a taut line, the outgoing exponent m and the outgoing transported covector determine the preceding exponent:

$$m_{\mathrm{in}} = m - \ell(m)r_*.$$

The normal component is unchanged by the bend, so its positive value $\ell(m)$ is known. The shear $I + r_* \otimes \ell$, and hence the preceding transported covector, are known as well.

Consequently two taut lines having the same endpoint and terminal exponent must have the same initial label. For a precise backward comparison, take the finite union of their actual nontrivial bends. Starting from their common terminal segment, take the latest remaining bend. At this crossing both lines are taut, even if one of them was originally straight there. The preceding formula forces the same incoming segment for both. Repetition over the finite union proves the assertion. If the common terminal exponent is zero, invertibility of the bend shears forces the initial label to be zero; the constant theta function is handled by the same argument.

The vertices cannot cancel. Let

$$\mathcal{P} = \mathrm{conv} \bigcup_{s \in S} \mathrm{Newt}(\vartheta_{\nu_s}).$$

Fix a vertex v of this finite polytope. There is a nonempty open cone of covectors exposing v as the unique minimum on $\mathcal{P}$. Choose w in that cone satisfying the genericity used in the preceding paragraph. If v occurred in two theta functions, choose an ordinary broken line contributing it in

each. Both lines minimize w, and hence are taut. Backward uniqueness would make their initial labels equal. Thus v belongs to exactly one theta function in the given family. Its coefficient in the proposed sum is a nonzero scalar c_s times a positive integer, and cannot vanish. Every vertex of $\mathcal{P}$ survives, while no support exponent of the sum lies outside $\mathcal{P}$. This proves (5.6), including nonvanishing of the sum. Taking minima of linear forms gives (5.7). □

6. A finite expansion on the active lattice

Let $\mathcal{X}$ be the framed PGL_n-space of the annulus from Section 3. All coordinates of a full unpinned seed remain present, including the inactive mutable coordinates produced by Lemma 4.1. We denote by $\mathcal{O}_{\mathrm{cl}}(\mathcal{X})$ the intersection of the Laurent rings of all mutation-equivalent seeds.

6.1. **The full annular expansion.**

Lemma 6.1. *For $\lambda \vdash n$, the annular holonomy function*

$$F_\lambda(g) = \frac{m_\lambda(\operatorname{eig} g)}{\det g}$$

belongs to $\mathcal{O}_{\mathrm{cl}}(\mathcal{X})$. In every full seed it has an expansion

$$F_\lambda = \sum_{s\in S} c_s \vartheta^{\mathrm{full}}_{\nu_s}, \qquad |S| < \infty, \tag{6.1}$$

where the labels ν_s are distinct, $c_s \neq 0$, and each $\vartheta^{\mathrm{full}}_{\nu_s}$ is a regular theta function.

Proof. By (2.1), F_λ is a regular class function on PGL_n. Its pullback by holonomy to the space with boundary pinnings is therefore regular.

Shen's regularity theorem [19, Theorem 1.1 and Definition 2.1] identifies the regular functions on the pinned space with the intersection of the Laurent rings over all its mutation-equivalent seeds. Forgetting the boundary pinnings deletes their frozen coordinates and gives precisely the framed cluster X-structure [9, Theorem 8.3, §8.2.3, §20.1.3, (704), (708)]. Every mutation word of an unpinned seed lifts to the same word in the pinned seed: the updates of the mutable coordinates and of their principal exchange submatrix do not use the deleted frozen coordinates.

Write x for the unpinned coordinates and z for those boundary frozen coordinates in a lifted chart. The pullback of F_λ is Laurent in x, z and belongs to $\mathbb{C}(x)$. Algebraic independence gives

$$\mathbb{C}[x^{\pm1}, z^{\pm1}] \cap \mathbb{C}(x) = \mathbb{C}[x^{\pm1}]. \tag{6.2}$$

This proves Laurentness in every unpinned seed, including seeds obtained by arbitrary mutation words.

The annulus with one marked point on each boundary is an admissible surface in [8, Definition 1.1]. The cluster Donaldson–Thomas transformation of [8, Theorem 1.3] supplies both signs of the initial basis rays in the associated $\mathcal{A}$ cluster complex. Their positive span is the whole space, so this complex is contained in no proper linear half-space. The full theta-basis criterion [12, Definition 8.27 and Proposition 8.28] therefore applies to the actual unpinned X-seed. It identifies its upper algebra with the vector space having the regular theta functions as a basis. This is an ordinary vector-space basis, and hence gives the finite expansion (6.1); see also [8, Theorem 1.4(i)]. The construction uses the character lattice of this seed throughout. □

6.2. **Eliminating the retained directions.**

At the terminal full seed, write

$$M = M_I \oplus M_J, \qquad \Omega = \begin{pmatrix} B & K \\ -K^{\mathsf{T}} & \Omega_J \end{pmatrix}. \tag{6.3}$$

Here I denotes the active gaps and J the retained inactive vertices, and Ω_J is the form on M_J. The matrix B is the active form computed in (4.12); in particular, $\det B = n^2 \neq 0$. The following descent also applies in any full seed satisfying the same hypotheses.

Lemma 6.2. *Suppose a nonzero finite linear combination of distinct regular full theta functions depends only on the active variables x_I, and the active form B is nonsingular. Then each summand is a Laurent polynomial in x_I alone. Its label lies in M_I, and its seed expression equals the theta function of the active datum with that same integral label.*

Proof. Let

$$F = \sum_{s \in S} c_s \vartheta^{\mathrm{full}}_{\nu_s}, \qquad c_s \neq 0.$$

The ring intersection (6.2), now with x_I, x_J, makes F a Laurent polynomial in x_I. For every $j \in J$ both $v_{e_j^*}(F)$ and $v_{-e_j^*}(F)$ equal zero. Lemma 5.1 therefore makes the corresponding two valuations of each summand nonnegative. Every monomial in every summand has its jth exponent both nonnegative and nonpositive, and hence equal to zero.

Each principal refinement is finite and has the form (5.4). Its leading coefficient is one, so $(\nu_s)_J = 0$. For every nonzero coefficient in that refinement the character exponent is $\nu_s + q$; its retained part is zero as well. Thus $q_J = 0$ for every term. These conclusions hold for each theta function separately.

Now take the continuous quotient of the root-degree completion obtained by setting

$$y_j = 0 \qquad (j \in J).$$

The monoid $\mathbb{Z}^I_{\geq 0} \times 0$ is a saturated face of the full free positive monoid. Degrees add under Lie brackets and under consecutive bends, so a term with positive retained degree cannot contribute to degree zero in the retained variables. The quotient diagram is therefore the consistent completion with only the active initial roots. Their root and normal maps, before the auxiliary enlargement, are

$$P_0(v) = \begin{pmatrix} v \\ 0 \end{pmatrix}, \qquad Q_0(v) = \begin{pmatrix} B^{\mathsf{T}} v \\ K^{\mathsf{T}} v \end{pmatrix}. \tag{6.4}$$

In particular, the mixed block K remains in the normals.

Theta functions of this quotient are the retained-degree-zero parts of the full theta functions. This can also be seen directly from broken lines: a line of total retained degree zero has only active nontrivial bends, and deleted walls are crossed using their constant terms. Conversely, a line in the quotient can be regarded in the full diagram in this way. At the fixed endpoint (5.2), the unfolding argument (5.11) excludes the additional joints of the full diagram as well. Since each summand was already proved to have only retained degree zero, it equals its quotient theta function.

It remains to identify the latter with the active theta function. The comparison takes place in real position spaces. Without auxiliary directions it is

$$\pi(u, z) = u + B^{-1} K z, \qquad u \in (M_I)_{\mathbb{R}}, \quad z \in (M_J)_{\mathbb{R}}.$$

In the enlargement (5.1), write the extra position as $a = (a_I, a_J)$ and use

$$\pi(u, z, a) = u + B^{-1}(Kz + a_I). \tag{6.5}$$

For every active root v,

$$\langle \widehat{Q}((v, 0)), (u, z, a) \rangle = v^{\mathsf{T}}(Bu + Kz + a_I) = v^{\mathsf{T}} B \pi(u, z, a).$$

Moreover, π fixes the active root directions. The pullback of the active scattering diagram consequently has exactly the initial normals (6.4), with the auxiliary components included. Its wall Lie brackets are the same, since the pairing of two active roots is $v^{\mathsf{T}}Bw$. Pullback preserves consistency and outgoingness; uniqueness of consistent completion identifies it with the quotient diagram.

A broken line with active initial exponent has only active exponents, so its retained and auxiliary positions z_0, a_0 are constant. Applying π sends it to an active broken line. The inverse lift is

$$\gamma(t) \longmapsto (\gamma(t) - B^{-1}(Kz_0 + (a_0)_I),\, z_0,\, a_0).$$

The normal identity preserves each crossing and its coefficient, and identifies the negative chambers. A strongly generic endpoint projects to one which is strongly generic in the active position space, because π is a rational surjection. This proves the claimed equality of theta expressions.

Only real positions have been changed. Roots, labels and character exponents remain in the original integral active lattice. Since the retained components of all labels are zero, their distinctness is preserved by the descent. □

Apply Lemma 6.2 to (6.1). Lemma 4.1 makes F_λ an active rational function, and Lemma 6.1 gives the required full regularity. Thus every summand in the terminal chart is an active theta Laurent polynomial.

6.3. **The vertex lattice.** We now use the integral identification

$$\iota : M_I \xrightarrow{\sim} L_0$$

of (4.14). Recall that L has basis $D_1, B_1, \dots, B_{n-1}, D_n$, that every vertex has weight one, and that

$$L_0 = \ker(\mathrm{wt}), \qquad \delta = \sum_{i=1}^{n} D_i, \qquad \mathrm{wt}(\delta) = n.$$

The forward gaps $c_1, \dots, c_{2n-2}$ form an integral basis of L_0. Together with the frozen vector D_1 they form an integral basis of L. We use the vertex seed with these mutable roots and with frozen vector D_1.

By (4.15), δ is in the radical of the vertex form. The active and vertex scattering diagrams compare by the real projection

$$\rho : L_{\mathbb{R}} \longrightarrow (L_0)_{\mathbb{R}}, \qquad \rho(z) = z - \frac{\mathrm{wt}(z)}{n}\delta. \tag{6.6}$$

Indeed, ρ fixes all root directions c_j, and for every such root

$$\Omega(c_j, \rho(z)) = \Omega(c_j, z).$$

The initial walls, their graded wall groups, and their outgoing completions therefore agree under pullback. For a weight-zero label, all exponents of a broken line have weight zero. The entire path has constant position weight, so (6.6) is just a translation of it in the radical direction. Translation back by the prescribed multiple of δ gives the inverse comparison and preserves all coefficients. Hence the active theta function of a label $\nu \in L_0$ equals the vertex theta function ϑ^L_ν.

The projection acts on real positions. The character lattice remains L_0, and $L_0 + \mathbb{Z}\delta$ has index n in L, since weight identifies $L/(L_0 + \mathbb{Z}\delta)$ with $\mathbb{Z}/n\mathbb{Z}$.

Radical translation gives the individual identity

$$X^{\delta}\vartheta^{L}_{\nu} = \vartheta^{L}_{\nu+\delta}. \tag{6.7}$$

A broken line $\Gamma(t)$ for ν is changed to $\Gamma(t) - t\delta$ for $\nu + \delta$. Wall supports are invariant under translation by δ, and all their normal pairings with δ vanish. Consequently the crossings, their times, and their coefficients are unchanged. The terminal exponent increases by δ, proving (6.7); this is also [3, Proposition 3.18]. All theta expressions are homogeneous for the weight grading, since every root has weight zero.

Let Y denote the spectral variables of the path, so $e_k(Y) = \mathcal{E}_k(P)$ as in Section 4. The homogenized identity (4.16) is

$$\iota(F_\lambda) = X^{-\delta} m_\lambda(Y).$$

It incorporates the normalization $d_n = 1$ of the terminal matrix. Multiplying the descended finite expansion by X^{δ} and using (6.7) gives

$$m_\lambda(Y) = \sum_{s\in S} c_s \vartheta^{L}_{m_s}, \qquad m_s = \iota(\nu_s) + \delta, \qquad \mathrm{wt}(m_s) = n. \tag{6.8}$$

Here ν_s is identified with its active component, the labels remain distinct, and each summand is a finite Laurent polynomial in the present chart. The shift is made only in the vertex lattice, where δ is radical. The next section identifies all coefficients in this finite expansion.

7. Rotations and the spectral theta function

We now identify the summands in the finite expansion (6.8). Throughout this section $n \geq 2$ and $\lambda \vdash n$. We use the vertex lattice and its form from Section 4. For a graph Γ on a vertex basis of L, write

$$\mathcal{E}_k(\Gamma) = \sum_{\substack{S \text{ independent in } \Gamma \\ |S|=k}} X^{\sum_{v\in S} v}.$$

The notation $f(Y_\Gamma)$ means substitution of $\mathcal{E}_k(\Gamma)$ for e_k in a symmetric polynomial f in n variables. The rotations below have one bridge on each level. They specialize the interval-graph rotations and monomial-section identity of [16, §§5.1–5.2, Theorem 5.8]. We give the finite coefficient argument needed for the annular comparison.

7.1. **The local rotation.** Place D_i at time i and B_i at time $i+\eta_i$, where the phases $\eta_1, \dots, \eta_{n-1}$ are distinct elements of $(0,1)$. Join two vertices when the difference of their times has absolute value less than one. The skew form is specified by

$$\Omega(D_i, D_j) = 0, \qquad \Omega(B_i, D_j) = \begin{cases} 1, & j = i, \\ -1, & j = i+1, \\ 0, & \text{otherwise}, \end{cases}$$
$$\Omega(B_{i+1}, B_i) = \mathbf{1}_{\eta_i > \eta_{i+1}}, \qquad \Omega(B_j, B_i) = 0 \quad (j > i+1). \tag{7.1}$$

Here $\mathbf{1}_{\mathcal{C}}$ is 1 when the condition $\mathcal{C}$ holds and 0 otherwise. The reverse pairings follow by skew-symmetry. Initially take $\eta_i = \theta_i$, where $0 < \theta_1 < \cdots < \theta_{n-1} < 1$; the graph is then the path of Lemma 4.2.

Suppose that η_i is the smallest phase. Put

$$p_i = B_i - D_i, \qquad q_i = D_{i+1} - B_i, \qquad \delta_i = D_{i+1} - D_i = p_i + q_i. \tag{7.2}$$

Replace this phase by a phase larger than all the others and replace the corresponding vertex vector by

$$B_i' = B_i + \delta_i. \tag{7.3}$$

All other vertex vectors are unchanged. Only the relative order of the phases matters.

Lemma 7.1. *The operation* (7.3) *is mutation at the actual gap p_i, followed by an interchange of the two gap indices on level i. If Γ' is the resulting graph, then*

$$\mathcal{E}_k(\Gamma) = \mathcal{A}_{p_i}(\mathcal{E}_k(\Gamma')), \qquad \mathcal{A}_p(X^a) = X^a(1+X^p)^{\Omega(p,a)}. \tag{7.4}$$

Consequently the same identity holds with $\mathcal{E}_k$ replaced by $f(Y_\Gamma)$ for any symmetric polynomial f. Individual theta labels transform by

$$T_p(m) = m + [-\Omega(p,m)]_+ p. \tag{7.5}$$

Proof. The smallest-phase condition and (7.1) give

$$\Omega(p_i, D_i) = 1, \qquad \Omega(p_i, D_{i+1}) = -1, \qquad \Omega(p_i, B_i) = 1,$$

and zero pairing with every other vertex. In the integral vertex seed with mutable basis $(p_1, q_1, \ldots, p_{n-1}, q_{n-1})$ and frozen vector D_1, the nonzero entries in the column of p_i are

$$\Omega(q_i, p_i) = 2, \quad \Omega(p_{i+1}, p_i) = -1, \quad \Omega(q_{i-1}, p_i) = -1, \quad \Omega(D_1, p_i) = -\mathbf{1}_{i=1}, \tag{7.6}$$

with nonexistent indices omitted. The basis mutation $e \mapsto e + [\Omega(e, p_i)]_+ p_i$ for $e \neq p_i$, and $p_i \mapsto -p_i$, therefore changes the two gaps on this level to $(-p_i, q_i + 2p_i)$ and fixes the frozen vector. Interchanging these two indices gives

$$p_i' = B_i' - D_i = 2p_i + q_i, \qquad q_i' = D_{i+1} - B_i' = -p_i.$$

The other gaps are unchanged. The only changed bridge pairings are $\Omega(B_i, B_{i-1}) = 1$ becoming 0, and $\Omega(B_{i+1}, B_i) = 0$ becoming 1. These are precisely the changes obtained by moving the smallest phase to the largest. Thus (7.1) holds in the new vertex basis. By the full-seed retention in Lemma 4.1, this mutation is also a mutation of the full annular seed; all retained coordinates and mixed exchange entries undergo that full mutation.

To prove the polynomial identity, fix an independent set S outside $\{D_i, B_i, D_{i+1}\}$. Set

$$a = \mathbf{1}_{B_{i-1} \notin S}, \qquad b = \mathbf{1}_{B_{i+1} \notin S},$$

taking the corresponding value to be 1 at a boundary. Use $x = X^{D_i}$, $y = X^{B_i}$, $z = X^{D_{i+1}}$, and $y' = X^{B_i'} = yz/x$. If u records cardinality, the possible extensions of S in the old and new triples contribute, respectively,

$$1 + u\{a(x+y) + bz\} + u^2 abxz,$$
$$1 + u\{ax + b(z+y')\} + u^2 abxz.$$

Indeed, the old bridge has the same outside neighbors as D_i, and the new bridge has the same outside neighbors as D_{i+1}. Both anchors can be chosen together. The automorphism $\mathcal{A}_{p_i}$ fixes all outside monomials, and

$$\mathcal{A}_{p_i}(x) = x + y, \qquad \mathcal{A}_{p_i}(z + y') = z, \qquad \mathcal{A}_{p_i}(xz) = xz.$$

It takes the second extension polynomial to the first. Multiplying by $X^{\sum_{v \in S} v}$ and summing over S proves (7.4). Since the elementary symmetric functions generate the symmetric polynomial ring, it proves the assertion for f as well. Formula (7.5) is the individual X-theta mutation formula in the negative-chamber normalization of Section 5; see [3, Appendix A.3]. □

We apply these rotations in the repeated order $1, 2, \dots, n-1$. This is the successive smallest-phase order. One such sequence is called a period. After a period the phase order is the original one, so the original numerical phases may be used again. After k periods the vertex vectors are

$$D_i, \qquad B_i^{(k)} = B_i + k\delta_i. \tag{7.7}$$

Every intermediate graph contains the anchor–bridge path and has the anchors as its unique independent set of cardinality n. In particular,

$$\mathcal{E}_n(\Gamma) = X^\delta, \qquad \delta = D_1 + \cdots + D_n, \tag{7.8}$$

throughout the sequence.

7.2. **The labels are pure and dominant.** Each summand in (6.8) comes from a regular theta function of the full annular variety. After an active mutation its rational expression still involves only the active coordinates: their transformation is autonomous. Its full expression is Laurent by regularity, so the Laurent-ring intersection used in Lemma 6.2 makes it a finite active Laurent polynomial in the new chart. An active basis change has an integral matrix U with $\det U = \pm 1$, and hence $\det(U^{\mathsf{T}} B U) = \det B = n^2$. Apply that lemma anew in this full seed. It identifies the summand with the active theta having its transported label and retains the full mixed normals in making the comparison.

Index the charts by the number $t \geq 0$ of individual rotations, with $\Gamma_0 = P$. Write $f_s^{(t)}$ for these full theta expressions in a rotated chart, and ι_t for its integral active-to-vertex embedding. The individual central shift gives finite Laurent polynomials

$$\Psi_s^{(t)} = X^\delta \iota_t(f_s^{(t)}), \qquad P_t = m_\lambda(Y_{\Gamma_t}) = \sum_s c_s \Psi_s^{(t)}.$$

The second equality follows from Lemma 7.1 and (7.8). For a real covector on L, extend its pullback along ι_t to the full seed lattice. Applying Lemma 5.1 there and then making the common affine shift by δ gives

$$\operatorname{Newt}(P_t) = \operatorname{conv} \bigcup_s \operatorname{Newt}\big(\Psi_s^{(t)}\big). \tag{7.9}$$

Each $\mathcal{E}_k(\Gamma_t)$ is an ordinary homogeneous vertex polynomial of degree k. Hence P_t is an ordinary homogeneous vertex polynomial of degree n. Its support is contained in the simplex

$$\Delta_n = \left\{ \sum_v a_v v : \ a_v \geq 0, \ \sum_v a_v = n \right\}.$$

Equation (7.9) puts the support of each $\Psi_s^{(t)}$ in this simplex. In particular, its pointed label has nonnegative integral vertex coordinates in every rotated chart.

Lemma 7.2. *Every label in* (6.8) *is of the form*

$$m = \sum_{i=1}^n \mu_i D_i, \qquad \mu_1 \geq \cdots \geq \mu_n \geq 0, \qquad \sum_i \mu_i = n. \tag{7.10}$$

Such a label is fixed by every subsequent rotation.

Proof. Write a current label as $m = \sum_j \nu_j D_j + \sum_j \kappa_j B_j$. At a level-$i$ rotation put $\epsilon = [-\Omega(p_i, m)]_+$. Reexpressing (7.5) in the basis with $B_i' = B_i + D_{i+1} - D_i$ gives

$$\kappa_i' = \kappa_i + \epsilon, \qquad \nu_i' = \nu_i + \kappa_i, \qquad \nu_{i+1}' = \nu_{i+1} - \kappa_i'. \tag{7.11}$$

All other coordinates stay fixed. Thus each bridge coefficient is nondecreasing. Define

$$\tau_i = n - \sum_{j\le i}(\nu_j + \kappa_j).$$

The simplex inclusion gives $0 \le \tau_i \le n$. An event at another level leaves τ_i unchanged, while an event at level i decreases it by κ_i'. After $n+1$ periods,

$$0 \le \tau_i^{\mathrm{end}} = \tau_i^{\mathrm{start}} - \sum_{a=1}^{n+1} \kappa_i^{(a)} \le n - (n+1)\kappa_i^{(1)},$$

where $\kappa_i^{(a)}$ is the coefficient after the ath event at that level. Nonnegativity and integrality imply $\kappa_i^{(1)} = 0$. Its initial value and its first increment are therefore both zero. This holds at every level. Consequently all events in the first period fixed the original label, which was already pure. Their inequalities $\epsilon = 0$ say $\mu_i - \mu_{i+1} \ge 0$, proving (7.10). Finally,

$$\Omega(p_i + a\delta_i, m) = \mu_i - \mu_{i+1} \ge 0 \qquad (a \ge 0),$$

so (7.5) fixes m at every later event. □

7.3. **A finite coefficient extraction.** We next recover the scalar coefficient of each pure label. The argument concerns only finitely many root degrees. In the original vertex seed put

$$c_{2i-1} = p_i, \qquad c_{2i} = q_i, \qquad C^+ = \sum_{j=1}^{2n-2} \mathbb{N}c_j.$$

If m is as in (7.10), then

$$h = m - nD_1 = \sum_{j=1}^{2n-2} h_j c_j = \sum_{i=1}^{n-1} \left(\sum_{a>i} \mu_a \right) (c_{2i-1} + c_{2i}), \qquad 0 \le h_j \le n. \tag{7.12}$$

Every pure nonnegative label ν of weight n also belongs to $nD_1 + C^+$. A broken line from ν with final exponent m has total increment $m - \nu \in C^+$. Each of its nonzero bend increments therefore belongs to the finite set

$$R_m = \left\{ \sum_{j=1}^{2n-2} a_j c_j : \ a_j \in \mathbb{N},\ a_j \le h_j,\ \sum_j a_j > 0 \right\}. \tag{7.13}$$

For wall positions it is convenient to use covectors on L. Set $V(u) = \Omega(\,\cdot\,, u)$. A wall of degree r is contained in $r^\perp$, and a broken-line segment carrying exponent m has covector velocity $-V(m)$. Since the radical is $\mathbb{R}\delta$ and the form on $(L_0)_{\mathbb{R}}$ is nondegenerate, V maps $L_{\mathbb{R}}$ onto $\delta^\perp$ and restricts to an isomorphism $(L_0)_{\mathbb{R}} \to \delta^\perp$. Wall data are constant along the radical in the exponent-position realization. For $Q = V(u)$, we write $\vartheta^L_{\nu,Q}$ for the section defined using any lift u. Two lifts differ by a real multiple of δ; translating each broken line by that difference preserves all coefficients.

Put $\ell = n+1$ and retain the initial phases $0 < \theta_1 < \cdots < \theta_{n-1} < 1$. Define a covector by

$$Q_\ell(D_i) = \frac{n+1}{2} - i, \qquad Q_\ell(p_i) = \ell - \theta_i. \tag{7.14}$$

Then $Q_\ell(\delta) = 0$ and $Q_\ell(\delta_i) = -1$. The final gaps from (7.7) satisfy

$$Q_\ell(p_i + \ell\delta_i) = -\theta_i < 0, \qquad Q_\ell(q_i - \ell\delta_i) = -1 + \theta_i < 0. \tag{7.15}$$

Thus this covector lies in the final negative chamber. Every earlier cut has the form $p_i + a\delta_i$, $0 \le a < \ell$, and

$$\begin{aligned} Q_\ell(p_i + a\delta_i) &= \ell - \theta_i - a > 0, \\ \Omega(p_i + a\delta_i, m) &= \mu_i - \mu_{i+1} \ge 0. \end{aligned} \tag{7.16}$$

Consequently the whole ray $Q_\ell + sV(m)$, $s \ge 0$, lies on the identity side of every geometric mutation. Indeed, the geometric $\mathcal{X}$-map on positions is $u \mapsto u + [\Omega(u,p)]_+ p$ [3, Appendix A]. Since $\Omega(u,p) = -Q(p)$, it is the identity when $Q(p) > 0$, and the same applies to the whole lifted ray $u + sm$.

Lemma 7.3. *There is a generic endpoint Q in the final negative chamber, arbitrarily close to Q_ℓ, for which*

$$[X^m]\vartheta^L_{\nu,Q} = \delta_{\nu,m} \tag{7.17}$$

for every pair of pure dominant labels m, ν of weight n.

Proof. We first work with the displayed covector Q_ℓ. Suppose that a wall with a relevant nonzero degree $r \in R_m$ could meet the ray $Q_\ell + sV(m)$. In the mutation of scattering diagrams, [3, Proposition A.1 and Lemma A.2] split each non-cut wall into its two half-space pieces, transport each piece and its wall function by the corresponding linear map, and place each transported wall function in the new positive-root group. We apply this construction to the descended active diagram and pull it back along (6.6). Choose the representatives recursively by this half-space construction along the rotation sequence, as in [12, Construction 1.30]. Thus every transported wall function belongs to the positive-root group of the current seed. By (7.16), every local map near the hypothetical intersection is the identity, and no mutation cut meets the ray. The actual wall piece and its nonzero degree r are therefore unchanged there through the sequence. The same degree must belong to the final positive gap cone

$$C^+_\ell = \sum_{i=1}^{n-1} \big(\mathbb{N}(p_i + \ell\delta_i) + \mathbb{N}(q_i - \ell\delta_i)\big).$$

This statement applies to each transported wall piece. A bend increment that is a positive multiple of its wall direction also belongs to C^+_ℓ; we use that actual increment as r.

Write $r = \sum_j a_j c_j \in R_m$ and put $\kappa_i(r) = a_{2i-1} - a_{2i}$. In the original ordered vertex basis the prefix sum of the coefficients of r through B_i is $-a_{2i}$. Expressing r in the final vertex basis (7.7) changes this prefix sum to

$$-a_{2i} + \ell\kappa_i(r).$$

Every proper vertex prefix of a vector in a forward-gap cone is nonpositive: if the vector is $\sum_j b_j(v_{j+1} - v_j)$ with $b_j \ge 0$, its prefix through v_j is $-b_j$. Thus $r \in C^+_\ell$ gives

$$-a_{2i} + \ell\kappa_i(r) \le 0. \tag{7.18}$$

Since $a_{2i} \le n < \ell$ and $\kappa_i(r)$ is integral, it follows that $\kappa_i(r) \le 0$ for every i. We obtain

$$Q_\ell(r) = \sum_i (\ell - \theta_i)\kappa_i(r) - \sum_i a_{2i} < 0, \tag{7.19}$$

$$\Omega(r, m) = \sum_i \kappa_i(r)(\mu_i - \mu_{i+1}) \le 0. \tag{7.20}$$

The first inequality is strict: either some $\kappa_i(r) < 0$, or all are zero and the nonzero vector r has some $a_{2i} > 0$. These two inequalities imply $(Q_\ell + sV(m))(r) < 0$ for every $s \ge 0$. The supporting hyperplane $r^\perp$ cannot meet the ray, a contradiction.

We can make the endpoint generic while preserving this argument. There are finitely many earlier cuts, finitely many final gaps, and finitely many nonzero lattice vectors $r = \sum_j a_j c_j$ with $0 \le a_j \le n$. Preserve the strict inequalities in (7.15), the first line of (7.16), and (7.19) for those vectors belonging to C_ℓ^+. They hold on a neighborhood of Q_ℓ in $\delta^\perp$. Choose in its inverse image under V an exponent-position endpoint outside every proper rational linear hyperplane, and let Q be its image. Such endpoints exist by surjectivity of V. This perturbation changes only the real endpoint; all labels remain in L. Inequality (7.20) is independent of the endpoint. The same exclusion of relevant walls consequently holds for $Q + sV(m)$, simultaneously for every pure dominant target of weight n.

Any broken line with final exponent m has its backward terminal segment on this ray. A last nontrivial bend would have an increment in R_m and would lie on one of the excluded walls. There is therefore no nontrivial bend. The straight line has coefficient one and initial label m, which proves (7.17). If m is radical, the projected ray is a point and the same inequalities still exclude every possible last bend. Equivalently, a bend of degree r into m would have input pairing $\Omega(r, m-r) = \Omega(r, m) = 0$, so its wall fixes that monomial. □

7.4. **Identification of the spectral function.**

Lemma 7.4. *In the terminal full annular seed, the function F_λ of* (2.1) *is the regular theta function*

$$F_\lambda = \vartheta^{\mathrm{full}}_{\mu_\lambda}, \qquad \mu_\lambda = \sum_{i=1}^{n} (\lambda_i - 1) D_i. \tag{7.21}$$

Here λ is padded by zeros, and $\mu_\lambda \in L_0$ is interpreted through the actual active gap lattice, with zero component on every retained vertex.

Proof. By Lemma 7.2, the finite vertex expansion consists of pure dominant labels of weight n. Denote the scalar coefficient of a label m by c_m, setting it to zero if the label is absent. Choose the endpoint Q of Lemma 7.3. Write $\vartheta^{L,0}_{\nu,Q}$ and $\vartheta^{L,\mathrm{fin}}_{\nu,Q}$ for the sections in the original and final scattering diagrams, respectively. The geometric mutation maps fix a lift of Q and are the identity near it by (7.16), so their induced linear maps on attached monomials are also the identity. The label maps fix every pure dominant label by Lemma 7.2. Individual theta mutation compatibility [3, Proposition A.3, equation (108), and Appendix A.3] therefore gives

$$\vartheta^{L,0}_{\nu,Q} = \vartheta^{L,\mathrm{fin}}_{\nu,Q}.$$

The right side is the final negative-chamber expression. There are finitely many summands. Ordinary chart compatibility and Lemma 7.1 therefore identify their sum with the finite polynomial $m_\lambda(Y_{\Gamma_{\mathrm{fin}}})$ in the final vertex variables. It follows from (7.17) that

$$c_m = [X^m] m_\lambda(Y_{\Gamma_{\mathrm{fin}}}).$$

The final vertices form an integral basis of L, since (7.7) adds only integral combinations of anchors to bridges. A monomial with a positive bridge coordinate cannot equal the pure-anchor monomial X^m. We may consequently set all final bridge variables to zero when computing this coefficient. The remaining graph consists of the independent anchors, so, for $m = \sum_i \mu_i D_i$,

$$c_m = [z_1^{\mu_1} \cdots z_n^{\mu_n}] m_\lambda(z_1, \ldots, z_n). \tag{7.22}$$

Among weakly decreasing nonnegative tuples μ of sum n, this coefficient is one precisely when $\mu = \lambda$, and is zero otherwise. Thus the finite vertex expansion is the single theta function $\vartheta^{L}_{\sum_i \lambda_i D_i}$. Undoing the shift by the radical vector δ and using the individual descent of Lemma 6.2

gives (7.21). Since it is one of the original full theta summands, it is regular on the full annular cluster variety. □

8. Proof of the immanant inequality

Proof of Theorem 1.1. For $n = 1$ the assertion is $a_{11} \geq 0$. Suppose $n \geq 2$ and first assume that every minor of A is strictly positive. The prescribed double reduced word in Section 4 gives a factorization

$$A = U(\mathbf{u}) \operatorname{diag}(t_1, \ldots, t_n) L(\boldsymbol{\ell})$$

with every parameter positive [7, 6]. Consequently all the annular source coordinates in (3.3) are positive. By Lemma 7.4 and theta positivity, F_λ is a finite Laurent polynomial with nonnegative integral coefficients in these coordinates.

For $Z = \operatorname{diag}(z_1, \ldots, z_n)$ we have

$$ZA = (ZUZ^{-1})(Z \operatorname{diag}(t_1, \ldots, t_n))L.$$

Every upper parameter of root i is multiplied by z_i/z_{i+1}, each t_i by z_i, and every lower parameter is unchanged. Substituting in (3.3) shows that all internal gaps are unchanged, whereas the outer gap satisfies

$$o_i \longmapsto (z_i/z_{i+1})o_i \qquad (1 \leq i < n). \tag{8.1}$$

The characters z_i/z_{i+1}, $1 \leq i < n$, are multiplicatively independent. Hence a source Laurent monomial has zero z-weight if and only if its exponent on every o_i is zero. Combining (2.3) and (8.1), we obtain

$$\operatorname{Imm}_{\phi^\lambda}(A) = \det(A)\ \operatorname{CT}_{o_1,\ldots,o_{n-1}} F_\lambda. \tag{8.2}$$

The constant term on the right is a Laurent polynomial with nonnegative coefficients in the internal source coordinates. Its value is nonnegative because those coordinates are positive; also $\det(A) > 0$. This proves the desired inequality for strictly totally positive A.

The strictly totally positive real matrices are dense in the totally nonnegative matrices, including the singular ones [7, Theorem 11 and the paragraph following it]. The immanant is a polynomial in the entries of A. Passing to the limit proves the inequality for every totally nonnegative A. □

References

[1] M.-W. Cheung, T. Magee, T. Mandel, and G. Muller, *Valuative independence and cluster theta reciprocity*, preprint (2026), arXiv:2505.09585v2.

[2] S. Clearman, B. Shelton, and M. Skandera, *Path tableaux and combinatorial interpretations of immanants for class functions on S_n*, Discrete Math. Theor. Comput. Sci. Proc. **AO** (2011), 233–244; published version.

[3] B. Davison and T. Mandel, *Strong positivity for quantum theta bases of quantum cluster algebras*, Invent. Math. **226** (2021), 725–843; arXiv:1910.12915v3.

[4] X. Dong, T. Wu, and X. Gao, *On the nonnegativity of monomial immanants for hook partitions*, preprint (2026), arXiv:2609.36665v1.

[5] V. V. Fock and A. B. Goncharov, *Cluster $\mathcal{X}$-varieties, amalgamation, and Poisson–Lie groups*, in *Algebraic Geometry and Number Theory*, Progr. Math. **253**, Birkhäuser, 2006, 27–68; arXiv:math/0508408v2.

[6] S. Fomin and A. Zelevinsky, *Double Bruhat cells and total positivity*, J. Amer. Math. Soc. **12** (1999), 335–380; arXiv:math/9802056.

[7] S. Fomin and A. Zelevinsky, *Total positivity: tests and parametrizations*, Math. Intelligencer **22** (2000), no. 1, 23–33; arXiv:math/9912128.

[8] A. Goncharov and L. Shen, *Donaldson–Thomas transformations of moduli spaces of G-local systems*, Adv. Math. **327** (2018), 225–348; arXiv:1602.06479v2.

[9] A. Goncharov and L. Shen, *Quantum geometry of moduli spaces of local systems and representation theory*, arXiv:1904.10491v4 (2024).

[10] I. P. Goulden and D. M. Jackson, *Immanants of combinatorial matrices*, J. Algebra **148** (1992), no. 2, 305–324, doi:10.1016/0021-8693(92)90196-S.

[11] C. Greene, *Proof of a conjecture on immanants of the Jacobi–Trudi matrix*, Linear Algebra Appl. **171** (1992), 65–79, doi:10.1016/0024-3795(92)90250-E.

[12] M. Gross, P. Hacking, S. Keel, and M. Kontsevich, *Canonical bases for cluster algebras*, J. Amer. Math. Soc. **31** (2018), 497–608; arXiv:1411.1394v2.

[13] M. Haiman, *Hecke algebra characters and immanant conjectures*, J. Amer. Math. Soc. **6** (1993), no. 3, 569–595; author's version.

[14] K. McNulty, *Hook monomial immanants are Kazhdan–Lusztig positive*, preprint (2026), arXiv:2610.39021v2.

[15] K. McNulty, *Haiman's conjecture for $n \leq 10$: Kazhdan–Lusztig positivity of monomial immanants*, preprint (2026), arXiv:2610.41555v2.

[16] OpenAI, *Elementary positivity of chromatic quasisymmetric functions*, manuscript, September 24, 2026.

[17] B. Rhoades and M. Skandera, *Kazhdan–Lusztig immanants and products of matrix minors*, J. Algebra **304** (2006), no. 2, 793–811, doi:10.1016/j.jalgebra.2005.07.017.

[18] G. Schrader and A. Shapiro, *Continuous tensor categories from quantum groups I: algebraic aspects*, preprint (2017), arXiv:1708.08107.

[19] L. Shen, *Duals of semisimple Poisson–Lie groups and cluster theory of moduli spaces of G-local systems*, Int. Math. Res. Not. IMRN **2022**, no. 18, 14295–14318; arXiv:2003.07901v2.

[20] J. R. Stembridge, *Immanants of totally positive matrices are nonnegative*, Bull. London Math. Soc. **23** (1991), no. 5, 422–428, doi:10.1112/blms/23.5.422.

[21] J. R. Stembridge, *Some conjectures for immanants*, Canad. J. Math. **44** (1992), no. 5, 1079–1099, doi:10.4153/CJM-1992-066-1.

Zhili College, Tsinghua University, Beijing 100084, China
Email address: tangqy24@mails.tsinghua.edu.cn