

Supplemental Material: “Gluons and Knots” by Alfredo Guevara

Abstract

We define an integer coefficient for plane polygons, establish its local wall and soft-insertion laws, and give equivalent tree and knot-polynomial constructions. The comparison rests on uniqueness: fixed-count transport reduces the difference to soft anchors and triangle values. Triple-wall invariance of the state sum is supplied by the separate proof and its bridge to the conventions used here. Definitions and results come first; short proofs remain beside their statements and longer proofs are collected in Part II. The final part records the literature inputs and the scope of the accompanying Lean formalization. The comparison on the generic locus is distinguished from its continuation to weakly generic polygons, and from the later star, decay, relation and census results.

Contents

Introduction and mathematical overview	4
The formalized proof	6
Reader’s guide	8
 Part I. Definitions and mathematical statements	 10
1 Polygons, chirotopes, chambers, and walls	10
1.1 Conventions	10
1.2 Polygons and chirotopes	10
1.3 Generic polygons and chambers	11
1.4 The Gauss word, interlacement, and the visible signature	13
1.5 Rotation number	14
1.6 Cyclic shift and reversal	15
1.7 Walls	16
1.8 Relative general position in the regular locus	20
 2 The tree coefficient A_g	 20
2.1 Roots, boundary words, gates	21
2.2 The coefficient	21
2.3 Chamber constancy and the resummation	22
2.4 The laws	23
2.5 Continuation through silent configurations	24
2.6 The soft theorem	26
2.7 Reversal and shift	27

3	The corner state sum C	28
3.1	Decompositions and subpolygons	28
3.2	The positive lift and the corner coefficient	29
3.3	From decorated crossing records to knots	29
3.4	The local source construction and its algebraic consequences	32
3.5	Scalar equality from named records	33
3.6	Marked products and mixed rows	34
3.7	The polynomial bound for an individual front	38
3.7.1	Ordinary certificates for the local front moves	39
3.7.2	The two oriented cusp-skein branches	41
3.7.3	The finite geometric reduction	41
3.7.4	The degree induction	42
3.8	Contact conventions and a bound for each representative	42
3.9	Exact regular rounding of the cusp endpoint	44
3.10	Contact-path polynomial transport and the representative bound	45
3.11	Tangent lifts and the full corner construction	47
3.12	Actual carrier blocks and elementary corner values	51
4	Laws for the corner state sum	52
4.1	The triple wall	54
5	Fixed-count transport	55
5.1	The circular stars	55
5.2	Connectivity of the regular fibres	55
5.3	Transport through simple walls	56
5.4	Anchors	56
6	Root independence, uniqueness, and the comparison theorem	57
6.1	Root independence of the tree coefficient	57
6.2	Uniqueness	59
6.3	The comparison theorem	61
7	Values at the circular stars and on convex polygons	61
7.1	The tree coefficient of a star	61
7.2	The state sum of a star	62
7.3	Finite computation for torus knots	62
7.3.1	Finite braid-word certificates	63
7.3.2	Parameters and the selected finite source	65
7.3.3	Finite spanning and the continuous representation family	65
7.3.4	The finite product and its two branching identities	66
7.3.5	Construction and uniqueness of the full Markov trace	67
7.3.6	Cycle characters and projector ranks	68
7.3.7	Uniform full-twist words and the torus trace	68
7.3.8	Identification with the local diagram polynomial	68
7.4	Convex polygons	70
8	The decay region	70

9	Linear relations: $U(1)$ decoupling and Kleiss–Kuijf	71
9.1	Common edge data and its block walls	72
9.2	The point-sign interface for the common family	72
9.3	Wall terms in the shuffle sum	73
9.4	Soft normalization and the shuffle identity	74
9.5	The word-algebra passage to the standard relation	74
9.6	The stated linear relations	75
10	Computations	75
10.1	Enumeration by carrier arrangements	76
10.2	Stored certificates and their scope	77
10.3	The pentagon atlas	77
	 Part II. Longer proofs	 79
	Proofs for Section 1	80
	Proofs for Section 2	92
	Proofs for Section 3	107
	Proofs for Section 4	150
	Proofs for Section 5	174
	Proofs for Section 6	179
	Proofs for Section 7	181
	Proofs for Section 8	196
	Proofs for Section 9	199
	 Part III. Inputs and formalization	 207
11	Assumptions and formalization	207
11.1	Literature inputs	207
11.2	Proof map	208
11.3	Correspondence with the letter	208
11.4	The accompanying Lean development	208
11.5	Coverage by chapter	210

Introduction and mathematical overview

This supplement develops the mathematical connection between two ways of assigning an integer to an oriented plane polygon: a combinatorial sum built from determinant signs, and a state sum built from knot polynomials. The organizing objects are the chambers of polygon space and the local walls between them. The value is constant within a chamber, while its changes across walls are governed by polygons with fewer vertices. The aim is to show that these local rules determine a global function and identify the two constructions.




Figure 1: Two constructions of the same integer on the generic locus. Determinant signs determine the rooted tree sum; uniform admissible smoothings determine a signed sum of products of positive-lift HOMFLY-PT coefficients. Root independence and comparison identify the outputs (Theorems 6.2 and 6.11).

Comparison. The tree coefficient $A_g(P)$ is a finite sum over rooted planar trees, with gates determined by the polygon’s chirotope. The corner state sum $C(P)$ sums over uniform admissible crossing smoothings, using specified HOMFLY-PT coefficients of the resulting positive lifts. On the generic locus $\mathcal{U}_n$ of Definition 1.5, root independence and comparison give

$$A_g(P) = A(P) = C(P), \quad P \in \mathcal{U}_n, \quad n \geq 3. \quad (\text{I.1})$$

These are Theorems 6.2 and 6.11. On the wider weakly generic locus $\mathcal{W}_n$, the corresponding identity uses the continuation A^{cont} across silent configurations (Corollary 6.13).

Local laws and uniqueness. Theorem 6.7 characterizes the common value among integer-valued functions on generic polygons of every arity. Its data are chamber constancy, invariance across silent and triple walls, the flat and vertex-edge laws, the soft-insertion theorem, and the two triangle values. The proof uses transport at fixed vertex count and soft insertions to reduce to smaller polygons. The cusp law is an additional consequence, with the deletion hypotheses in Corollary 6.12. Triple-wall invariance of the corner sum is supplied by the separate R proof and its bridge (Theorem 4.10).




Figure 2: Wall-crossing jumps. The upper row shows the polygon before (pale blue) and after (black) a vertex–edge event (W1) or a flat event (W2). The lower row shows the oriented smoothing or deletion controlling the jump; the numbers give the chamber values on the two sides (Theorems 2.14 and 2.12).

The knot-theoretic bound. Theorem 3.64 controls the lowest a -degree of the positive-lift HOMFLY–PT polynomial by the carrier’s crossing count and absolute rotation, and establishes nonnegative even powers of z . It applies to decomposition carriers admitting an orientation with all turns left, or exactly one right and all others left. This estimate supplies the degree control used in the corner-coefficient wall laws; it does not assert that the corner coefficient is nonzero.

Evaluations and further consequences. For the circular stars, the common value is a signed Catalan number:

$$A(K_r) = C(K_r) = (-1)^r \frac{1}{r+1} \binom{2r}{r}, \quad r \geq 1. \quad (\text{I.2})$$

These evaluations are consequences, not extra uniqueness premises (Section 7). Further results give the values on convex polygons, the bound $|C| \leq 1$ in the specified decay region, and the $U(1)$ and Kleiss–Kuijff identities under their generic-family hypotheses. Section 10 records the finite chamber census and its computational scope. The accompanying formalization and its dependency graphs are described below.

Figure 3 gives concrete examples.



Figure 3: Examples of the common coefficient: the triangle K_1 , bow-tie K_0 , and circular stars K_2, K_3 , with values $-1, -1, 2, -5$. Crossing decorations depict positive lifts. Higher-star values are consequences; only the two oriented triangle values enter uniqueness (Section 7).

The formalized proof

The comparison between the tree coefficient and the corner state sum is also supported by a computer-checked proof in Lean 4 [35], available in the [gluon-knot-lean repository](#). Lean requires each definition and each step of the argument to be written in a precise language, then checks that the conclusions follow from the stated assumptions. For an argument with several mathematical ingredients, this makes explicit how the pieces fit together and where each result is used.

We use the wall labels of the main paper: **W1** for a vertex-edge event, **W2** for a flat event, **W3** for triple-wall invariance, and **W4** for a cusp event. The W1–W3 laws for the corner state sum, together with soft insertion and the other local properties, lead through transport and uniqueness to the comparison with the tree coefficient. The comparison then transfers W4 from the tree coefficient to the corner state sum, with the stated deletion hypotheses. The separate triple-wall argument, called R, supplies W3; the bridge translates it into this supplement’s conventions. Figure 4 gives a map of this part of the proof.

Like the mathematical proof, the Lean development takes some results from the literature as starting points. Lean checks the deductions from these inputs; relating its precise definitions to the geometric language of the paper is a further part of the mathematical work. Section 11 explains these points in detail. The formalized comparison concerns generic polygons. The extension to weakly generic polygons and the later results on families, relations and the census are treated separately in this supplement.




Figure 4: The W1 and W2 laws for C feed into the comparison $C = A$, which transfers the W4 cusp law from A to C . The W3 bridge supplies triple-wall invariance. Arrows show selected dependencies, as explained below. Adapted from the repository’s [proof map](#) [35].

How to read a DAG. A **DAG** is a **directed acyclic graph**. Here, boxes are mathematical statements and arrows record dependencies: an arrow from one box to another means that the second uses the first. *Directed* means the arrows have an orientation; *acyclic* means that following them can never bring you back to your starting box. Read from the inputs toward the conclusions, with prerequisites before the results that use them. Several incoming arrows bring ingredients together; branches show one result being used in different arguments. These diagrams show selected connections, so an unshown arrow does not mean that no dependency exists.

Figure 4 shows this structure in the comparison and the W1–W4 laws. Figure 5 follows the argument for W3: knot-polynomial and contact-geometric inputs give bounds for the carriers, then a case analysis and the bridge give the triple-wall law for the polygon state sum. The repository’s [interactive blueprint](#) links statements to their Lean counterparts; the [complete blueprint drawing](#) shows the wider dependency graph. The interactive HTML can be opened locally from the repository.




Figure 5: The argument for W3 (triple-wall invariance). Purple boxes mark results taken from the literature, while green boxes mark results proved in the development. Bounds for the carriers feed into the case analysis, and the bridge brings the conclusion back to the polygon state sum. Selected connections from the repository’s [proof map](#) [35].

Reader’s guide

The document has two constructions at its centre: the tree coefficient of Section 2 and the corner state sum of Section 3. Definitions and statements come first; longer proofs are collected in Part II, with links beside their statements. Short proofs remain nearby. The following routes locate their definitions, the laws they satisfy, and the comparison argument. All theorem numbers below refer to the full statements and proofs that follow.

Definitions and conventions. Begin with the polygon and its chirotope (Definitions 1.1 and 1.3), and distinguish the generic locus of Definition 1.5 from the weakly generic locus of Definition 1.15. The tree construction is given in Definitions 2.2 and 2.4. For the state sum, read Definitions 3.1, 3.6, 3.8 and 3.28; Theorem 3.64 is the corner bound supporting that construction.

Laws and comparison. Sections 2 and 4 establish the local laws of the two constructions. Section 5 provides the transport used in Section 6, which contains root independence, the uniqueness theorem (Theorem 6.7) and the comparison theorem (Theorem 6.11). The triple-wall input is Theorem 4.10. The weakly generic continuation and its comparison are stated separately in Definition 6.5 and Corollary 6.13.

Examples and further consequences. Section 7 treats circular stars and convex polygons; Section 8, the decay region; and Section 9, the linear relations. The census is specified by Definition 10.1 and recorded in Computation 10.2. The hypotheses and literature inputs are collected in Section 11.

Law names used in the letter.

Letter name	Tree coefficient	Corner state sum
W1: vertex–edge	Theorem 2.14	Theorem 4.6
W2: flat	Theorem 2.12	Theorem 4.4
W3: triple	Theorem 2.15(i)	Theorem 4.10
W4: cusp	Theorem 2.13	Corollary 6.12
Soft insertion	Theorem 2.24	Theorem 4.8

Each entry refers to its full statement, with the domain and literature inputs specified there. The triple-wall proof is cited in Theorem 4.10. The letter’s generic domain corresponds to the weakly generic locus of Definition 1.15. The generic locus in Definition 1.5 is narrower, and the continuation to the wider locus is stated separately.

Notation and formal declarations. The Lean names below are in the SM namespace and identify the corresponding definition rows.

Letter object	Companion definition	Lean declaration
Polygon and chirotope	1.1 , 1.3	<code>polygonData</code> , <code>chirotopeData</code>
Generic polygons of the letter	$\mathcal{W}_n$, 1.15	<code>weak_definition</code>
Auxiliary generic locus	$\mathcal{U}_n$, 1.5	<code>Generic</code>
Root, gates, and rooted tree sum	2.1 , 2.2 , 2.4	<code>rootData</code> , <code>gatesData</code> , <code>treesumData</code>
Decompositions and uniformity	3.1 , 3.6	<code>decomposition_definition</code> , <code>uniform_definition</code>
Positive lift and corner sum	3.8 , 3.28	<code>positive_lift_definition</code> , <code>corner_state_sum_definition</code>
Soft insertion	2.22	<code>softInsertion_definition</code>

At silent configurations the letter's amplitude is read as A^{cont} , defined in [6.5](#). Formalizing the definition of the weak locus does not establish the continuation theorem. The letter's decay region $\mathcal{R}_1$ is $\mathcal{D}_1$ of Definition [8.1](#).

How the proof fits together. The gates and corner coefficients lead to their respective chamber and wall laws. Fixed-count transport and soft anchors give uniqueness from triangle values. Applying uniqueness to the corner sum, with the triple-wall theorem, gives comparison. The calculations at higher circular stars are consequences and are not premises of uniqueness. The correspondence with the letter and the formalization is collected in Section [11](#).

Part I. Definitions and mathematical statements

1 Polygons, chirotopes, chambers, and walls

1.1 Conventions

For $u = (u_1, u_2)$, $v = (v_1, v_2) \in \mathbb{R}^2$ put $\det(u, v) = u_1 v_2 - u_2 v_1$. The sign function is $\text{sgn}: \mathbb{R} \rightarrow \{-1, 0, 1\}$. The step function is

$$\Theta: \mathbb{R} \setminus \{0\} \rightarrow \{0, 1\}, \quad \Theta(x) = 1 \ (x > 0), \quad \Theta(x) = 0 \ (x < 0);$$

it is never evaluated at 0 in this document, and every statement that consults it is made on a domain where its argument is nonzero (Lemma 2.3). For $n \geq 3$ indices of vertices and edges are elements of $\mathbb{Z}/n$; $i + 1$, $i - 1$ are read modulo n . For comparison with the letter, set $[u, v] = u_2 v_1 - u_1 v_2 = -\det(u, v)$; this document uses $\det$ only. Here $E_i = [\mu_i, \mu_{i+1}]$ and $\tilde{\lambda}_i = \mu_{i+1} - \mu_i$, as in the letter.

1.2 Polygons and chirotopes

Definition 1.1 (labelled tuples and polygons). Let $n \geq 3$. A *labelled tuple* is an element $P = (\mu_1, \dots, \mu_n) \in (\mathbb{R}^2)^n$, indices in $\mathbb{Z}/n$ with $\mu_{n+1} = \mu_1$. The *cyclic shift* is $(\sigma P)_i = \mu_{i+1}$. A *polygon* on n vertices is an orbit $[P] = \{\sigma^k P : k \in \mathbb{Z}/n\}$ of labelled tuples; a labelled tuple in the orbit is a *representative*, and the space of polygons is the quotient $(\mathbb{R}^2)^n / (\mathbb{Z}/n)$. Every notion below is introduced on a representative and is invariant under σ (up to the induced relabelling of indices, which is recorded where it matters); it therefore defines a notion on polygons. A polygon has no distinguished vertex and no distinguished edge; when a construction needs one (a root, a base vertex), that choice is part of the construction's data.

For a representative, the *edge vectors* are

$$\tilde{\lambda}_i = \mu_{i+1} - \mu_i \quad (i \in \mathbb{Z}/n),$$

so that $\sum_i \tilde{\lambda}_i = 0$, and the *edge segments* are $E_i = [\mu_i, \mu_{i+1}] = \{\mu_i + t\tilde{\lambda}_i : 0 \leq t \leq 1\}$. Vertex i is *incident* to E_{i-1} and E_i . Two edges E_i, E_j are *adjacent* if $j - i \in \{-1, 0, 1\}$ and *remote* otherwise; an edge E_j is *remote to the vertex* i if $j \notin \{i - 1, i\}$, and *remote to the vertex and its two edges* if $j \notin \{i - 2, i - 1, i, i + 1\}$. No condition is imposed on the tuple by this definition; in particular coincident vertices and zero edges are allowed at this level.

Remark 1.2 (formalization). The quotient by the finite group $\mathbb{Z}/n$ is the only quotient used. On the generic locus the action is free (distinct vertices), so the quotient is a $2n$ -manifold and its connected components are the orbits of the components of the labelled generic locus. A formalization may work with labelled tuples throughout and record each construction's shift-equivariance as a lemma; the statements below are written so that both readings agree.

Definition 1.3 (chirotope and turns). For $i, j, k \in \mathbb{Z}/n$ put

$$\chi_{ijk}(P) = \text{sgn} \det(\mu_j - \mu_i, \mu_k - \mu_i) \in \{-1, 0, 1\}.$$

The *turn* at vertex i is $\tau_i(P) = \chi_{i-1, i, i+1}(P)$. The number of *left turns* is $\ell(P) = \#\{i : \tau_i(P) = 1\}$.

Lemma 1.4 (elementary properties of χ). *For all i, j, k :*

- (i) χ_{ijk} is alternating: it changes sign under a transposition of two of its indices and is unchanged under a cyclic permutation; it is 0 when two indices coincide.
- (ii) $\chi_{ijk} = \text{sgn}([\mu_j - \mu_i, \mu_j - \mu_k])$ with the bracket convention of Section 1.1, $[u, v] = u_2v_1 - u_1v_2$.
- (iii) $\chi_{i,i+1,k} = \text{sgn} \det(\tilde{\lambda}_i, \mu_k - \mu_i)$; when $\tilde{\lambda}_i \neq 0$ we say (a naming convention) that μ_k lies strictly to the left of the directed line through E_i if this sign is +1 and strictly to the right if it is -1.
- (iv) $\tau_i = \text{sgn} \det(\tilde{\lambda}_{i-1}, \tilde{\lambda}_i)$; $\tau_i = +1$ is a left turn.

Proof. (i) $\det$ is alternating and $\det(\mu_j - \mu_i, \mu_k - \mu_i)$ is the standard area form of the ordered triple; a cyclic permutation of the triple is even. (ii) $[\mu_j - \mu_i, \mu_j - \mu_k] = -\det(\mu_j - \mu_i, \mu_j - \mu_k) = \det(\mu_j - \mu_i, \mu_k - \mu_j) = \det(\mu_j - \mu_i, \mu_k - \mu_i)$, the last step by $\det(u, v - u) = \det(u, v)$. (iii) is the case $j = i + 1$ of the definition. (iv) $\det(\mu_i - \mu_{i-1}, \mu_{i+1} - \mu_{i-1}) = \det(\tilde{\lambda}_{i-1}, \tilde{\lambda}_{i-1} + \tilde{\lambda}_i) = \det(\tilde{\lambda}_{i-1}, \tilde{\lambda}_i)$. $\square$

1.3 Generic polygons and chambers

Definition 1.5 (generic polygon). A polygon P is *generic* if

- (G1) $\chi_{ijk}(P) \neq 0$ for all pairwise distinct $i, j, k \in \mathbb{Z}/n$;
- (G2) there is no point of $\mathbb{R}^2$ lying in the relative interiors of three distinct edge segments.

The set of generic polygons on n vertices is $\mathcal{U}_n \subseteq (\mathbb{R}^2)^n$.

Lemma 1.6 (what (G1) gives). *Let P satisfy (G1). Then:*

- (i) all vertices are distinct and every edge vector is nonzero;
- (ii) every turn τ_i is ± 1 ; consecutive edge vectors are linearly independent, so no edge vector is a positive or negative multiple (“flat” or “kink”) of its predecessor;
- (iii) no vertex lies on the line spanned by an edge not incident to it; in particular no vertex lies on a non-incident closed edge segment;
- (iv) two distinct adjacent edges meet exactly in their common vertex;
- (v) two remote edges are either disjoint or meet in exactly one point, which lies in the relative interior of both and at which the two edge directions are linearly independent.

Proof. (i) If $\mu_i = \mu_j$ with $i \neq j$, then $\chi_{ijk} = 0$ for any third index k (there is one since $n \geq 3$). (ii) $\tau_i = \text{sgn} \det(\tilde{\lambda}_{i-1}, \tilde{\lambda}_i) \neq 0$. (iii) If μ_k lies on the line through μ_i, μ_{i+1} with $k \notin \{i, i+1\}$ then $\chi_{i,i+1,k} = 0$. (iv) By (ii) the two edges at a vertex span the plane, so their lines meet only at the vertex. (v) Let E_i, E_j be remote and $x \in E_i \cap E_j$. By (iii) x is not an endpoint of either. If $\tilde{\lambda}_i, \tilde{\lambda}_j$ were parallel, the two segments would lie on one line, and then an endpoint of one would lie on the line of the other, contrary to (iii). Two segments on non-parallel lines meet in at most one point. $\square$

Definition 1.7 (crossings). Let P satisfy (G1). The *crossing set* $X(P)$ is the set of unordered pairs $\{i, j\}$ of remote edge indices with $E_i \cap E_j \neq \emptyset$. For $\{i, j\} \in X(P)$ the *crossing point* x_{ij} is the unique point of $E_i \cap E_j$, and its *parameters* $t_{ij} \in (0, 1)$ on E_i and $t_{ji} \in (0, 1)$ on E_j are defined by $x_{ij} = \mu_i + t_{ij}\tilde{\lambda}_i = \mu_j + t_{ji}\tilde{\lambda}_j$. The *crossing sign* of $\{i, j\}$ read from E_i is $\text{sgn} \det(\tilde{\lambda}_i, \tilde{\lambda}_j)$.

Lemma 1.8 (crossings through the chirotope). *Let P satisfy (G1) and let E_i, E_j be remote. Then $\{i, j\} \in X(P)$ if and only if*

$$\chi_{i,i+1,j} \chi_{i,i+1,j+1} = -1 \quad \text{and} \quad \chi_{j,j+1,i} \chi_{j,j+1,i+1} = -1.$$

If P is generic, distinct crossings have distinct crossing points, and $X(P)$ is finite.

Proof. By Lemma 1.6(i),(iii) the edge vectors $\tilde{\lambda}_i, \tilde{\lambda}_j$ are nonzero and none of the four endpoints of E_i, E_j lies on the line of the other segment, so the four chirotope entries in the display are ± 1 . Separation criterion: along E_j the signed distance from the line of E_i , $f(t) = \det(\tilde{\lambda}_i, \mu_j + t\tilde{\lambda}_j - \mu_i)$, is affine in t with $f(0), f(1) \neq 0$, so E_j meets the line of E_i exactly when $f(0)f(1) < 0$, that is, when $\chi_{i,i+1,j} \chi_{i,i+1,j+1} = -1$ (Lemma 1.4(iii)); then f has nonzero slope $\det(\tilde{\lambda}_i, \tilde{\lambda}_j)$, the two lines meet in one point, and that point lies in the relative interior of E_j . Symmetrically, E_i meets the line of E_j exactly when the second condition holds. If both conditions hold, the common point of the two lines lies in the relative interior of both segments, so $E_i \cap E_j \neq \emptyset$; conversely, a point of $E_i \cap E_j$ is a point of E_j on the line of E_i and a point of E_i on the line of E_j , so both conditions hold. If two distinct crossings had the same point, that point would lie in the relative interior of three distinct edges (the two crossings share at most one edge), contrary to (G2). Finiteness is clear. $\square$

Lemma 1.9 (stability of separated segment intersections). *Consider finitely many continuously varying nonzero oriented segments. For a specified pair, suppose that at the centre the segments are either disjoint or meet transversely in both relative interiors. In the first case they stay disjoint nearby. In the second case their unique intersection persists, remains transverse and interior, and its two parameters vary continuously. For finitely many such persistent crossings on a segment, every pair with distinct central parameters keeps its parameter order nearby.*

Proof. Disjoint compact segments have positive distance. Moving their endpoints by at most η moves every convex combination of the endpoints by at most η , so a sufficiently small perturbation preserves their disjointness. For the second case write the segments as $A + uD$ and $B + vH$, where $0 \leq u, v \leq 1$. Solving $A + uD = B + vH$ by taking determinants gives

$$u = \frac{\det(B - A, H)}{\det(D, H)}, \quad v = \frac{\det(B - A, D)}{\det(D, H)}. \quad (1)$$

The central denominator is nonzero. These quotients are therefore continuous near the centre, and their central values in $(0, 1)$ remain in that interval. A nonzero continuous difference of two central parameters has constant sign after shrinking. There are only finitely many requirements, so one common interval satisfies all of them. $\square$

Definition 1.10 (chambers). A *chamber* is a connected component of the space of generic polygons $\mathcal{U}_n/(\mathbb{Z}/n)$; its preimage in $\mathcal{U}_n$ is a union of at most n components of $\mathcal{U}_n$ permuted by σ , the *labelled chambers*.

Proposition 1.11 (chambers). *$\mathcal{U}_n$ is open in $(\mathbb{R}^2)^n$, and every chamber is open and path connected. Along any path in $\mathcal{U}_n$ the following data are constant: the chirotope (χ_{ijk}) ; the crossing set $X(P)$; and, for every edge E_i , the order of the parameters t_{ij} , $\{i, j\} \in X(P)$, along E_i .*

Proof. The finitely many determinants $\det(\mu_j - \mu_i, \mu_k - \mu_i)$ are continuous and nonzero on $\mathcal{U}_n$, so (G1) is open. Condition (G2) is open as well: if three edge interiors had a common point at a limit of polygons, the three segments would meet in the limit; we argue directly. Fix $P \in \mathcal{U}_n$.

By Lemma 1.6(v) each pair of remote edges is either disjoint or meets transversally in one interior point; the first condition persists under small perturbation, and in the second the crossing persists and varies continuously (Lemma 1.9). The finitely many crossing points of P are pairwise distinct (Lemma 1.8), hence remain distinct nearby; a common point of three edge interiors nearby would be a coincidence of two crossing points (any two of the three edges are remote, since an adjacent pair meets only at its vertex by Lemma 1.6(iv), and a vertex on a remote edge is excluded by (G1) nearby). So $\mathcal{U}_n$ is open. An open subset of $\mathbb{R}^{2n}$ is locally path connected, so its components are open and path connected.

Along a path in $\mathcal{U}_n$ each χ_{ijk} is a continuous nonzero integer-valued function, hence constant. The crossing set is a function of the chirotope (Lemma 1.8). Two crossings on a common edge E_i change their order only if their parameters coincide at some time; at that time the two crossing points coincide, which puts a common point in three edge interiors, contrary to (G2) — the two other edges are remote to E_i and, if they were adjacent to each other, their only common point is their shared vertex, which cannot lie on E_i by (G1). $\square$

1.4 The Gauss word, interlacement, and the visible signature

Definition 1.12 (traversal circle and Gauss word). Let P be generic. The *traversal circle* $\Gamma(P)$ is the set of pairs (i, t) , $i \in \mathbb{Z}/n$, $t \in [0, 1)$, with the cyclic order in which $(i, t) < (i, t')$ for $t < t'$ and $(i, t) < (i + 1, t')$ for all t, t' ; the point of the plane at (i, t) is $\mu_i + t\tilde{\lambda}_i$. Each crossing $\{i, j\} \in X(P)$ has two *visits* (i, t_{ij}) and (j, t_{ji}) on $\Gamma(P)$. The *Gauss word* of P is the cyclic sequence of the $2|X(P)|$ visits in the order of $\Gamma(P)$, each visit labelled by its crossing.

Definition 1.13 (interlacement). Two distinct crossings $x, y \in X(P)$ *interlace* if the four visits x_0, y_0, x_1, y_1 occur in the cyclic order $x_0 < y_0 < x_1 < y_1$ (up to rotation), that is, exactly one visit of y lies between the two visits of x . The *interlacement graph* G_P has vertex set $X(P)$ and an edge between each interlacing pair. $\text{Ind}(G_P)$ denotes the set of independent subsets of $X(P)$, including $\emptyset$; for $S \subseteq X(P)$ let $N(S)$ be the set of crossings interlacing some element of S , and let $U(S) = X(P) \setminus (S \cup N(S))$.

Definition 1.14 (visible signature). The *visible signature* of a generic P is the triple

$$((\tau_i)_{i \in \mathbb{Z}/n}, X(P), \text{ the Gauss word of } P).$$

By Proposition 1.11 it is constant on chambers.

Definition 1.15 (weakly generic polygons and the silent locus). A polygon P is *weakly generic* if every edge vector is nonzero, every turn τ_i is nonzero, no vertex lies on a non-incident closed edge segment, any two remote edges meet in at most one point and there transversally, and (G2) holds. This is the letter's “generic”. Every generic polygon is weakly generic; the *silent locus* is $\mathcal{S}_n = \mathcal{W}_n \setminus \mathcal{U}_n$, where $\mathcal{W}_n$ is the set of weakly generic polygons. A *visible chamber* is a connected component of $\mathcal{W}_n$.

Lemma 1.16 (openness of the weak locus). *Read on labelled representatives, $\mathcal{W}_n$ is open in $(\mathbb{R}^2)^n$. Every connected component is open and polygonally connected. In particular, every labelled weakly generic tuple has an open ball contained in its visible chamber.*

Proof in Part II, page 80.

Lemma 1.17 (the silent locus). *If P is weakly generic and $\chi_{ijk}(P) = 0$ for pairwise distinct i, j, k , then either no two of i, j, k are consecutive (a pure chord), or exactly one pair is consecutive, say*

$\{i, i+1\}$, and μ_k lies on the line of E_i outside the closed segment E_i (an exterior extension). The visible signature (Definition 1.14) is defined and locally constant on $\mathcal{W}_n$, hence constant on visible chambers, and $\mathcal{U}_n$ is open and dense in $\mathcal{W}_n$.

Proof. Two consecutive pairs in $\{i, j, k\}$ make it a turn triple, excluded. One consecutive pair $\{i, i+1\}$ with μ_k on the line of E_i : μ_k is not on the closed segment by hypothesis. Crossings are created or destroyed only by a vertex passing through a remote closed segment, and their order along an edge changes only through a coincidence of crossing points; both are excluded on $\mathcal{W}_n$ exactly as in Proposition 1.11, so the visible signature is locally constant. Density: $\mathcal{W}_n$ is open (Lemma 1.16) and $\mathcal{U}_n$ is its subset on which every χ_{ijk} is nonzero; the zero set of the nonzero polynomial $\prod \det(\mu_j - \mu_i, \mu_k - \mu_i)$ over the triples has empty interior in $(\mathbb{R}^2)^n$, so it contains no open subset of $\mathcal{W}_n$, and every point of $\mathcal{W}_n$ is a limit of points of $\mathcal{U}_n$. $\square$

1.5 Rotation number

Definition 1.18 (regular locus and principal turns). For a polygon P and an index i with $\tilde{\lambda}_{i-1}, \tilde{\lambda}_i \neq 0$ and $\tilde{\lambda}_i$ not a negative real multiple of $\tilde{\lambda}_{i-1}$, the *principal turn* $\vartheta_i(P) \in (-\pi, \pi)$ is the unique angle with $\cos \vartheta_i = \langle \tilde{\lambda}_{i-1}, \tilde{\lambda}_i \rangle / (|\tilde{\lambda}_{i-1}| |\tilde{\lambda}_i|)$ and $\text{sgn } \vartheta_i = \text{sgn } \det(\tilde{\lambda}_{i-1}, \tilde{\lambda}_i)$. The *regular locus* $\mathcal{R}_n$ is the set of polygons with all $\lambda_i \neq 0$ and no $\tilde{\lambda}_i$ a negative multiple of $\tilde{\lambda}_{i-1}$; equivalently, all principal turns exist. Every generic polygon lies in $\mathcal{R}_n$ (Lemma 1.6(ii)), and $\vartheta_i(P) = 0$ exactly when $\tilde{\lambda}_i$ is a positive multiple of $\tilde{\lambda}_{i-1}$.

Definition 1.19 (reversal). For a representative $P = (\mu_1, \dots, \mu_n)$ the *reversal* is $(\bar{P})_i = \mu_{2-i}$ ($i \in \mathbb{Z}/n$), which fixes μ_1 and reverses the traversal. Reversal commutes with the shift up to a shift, so it is defined on polygons; it is not a relabelling but a change of orientation, and no law below is asserted to be invariant under it except where stated. Write $\rho(P) = \bar{P}$ for this reversal map.

Lemma 1.20 (rotation number). For $P \in \mathcal{R}_n$ put $\text{rot}(P) = \frac{1}{2\pi} \sum_{i \in \mathbb{Z}/n} \vartheta_i(P)$. Then:

- (i) $\text{rot}(P) \in \mathbb{Z}$;
- (ii) rot is constant along every continuous path in $\mathcal{R}_n$;
- (iii) inserting a vertex in the relative interior of an edge does not change rot ; reversing the traversal negates it;
- (iv) if $n = 3$ then $\text{rot}(P) = \tau_1(P) = \tau_2(P) = \tau_3(P) \in \{\pm 1\}$ and in particular $\text{rot}(P) \neq 0$;
- (v) $2|\text{rot}(P)| < n$.

Proof. Identify $\mathbb{R}^2$ with $\mathbb{C}$ and put $q_i = \tilde{\lambda}_i / |\tilde{\lambda}_i|$. The defining property of the principal turn is $q_i = e^{i\vartheta_i} q_{i-1}$. Multiplying around the cycle gives $\exp(i \sum_i \vartheta_i) = 1$, so the sum is in $2\pi\mathbb{Z}$: (i). The principal turn is a continuous function of $(\tilde{\lambda}_{i-1}, \tilde{\lambda}_i)$ on the open set where both are nonzero and not antiparallel, so rot is a continuous integer-valued function on $\mathcal{R}_n$: (ii). Inserting a vertex on an edge adds one zero turn and changes no other turn; reversal reverses the cyclic order of the pairs and replaces each $(\tilde{\lambda}_{i-1}, \tilde{\lambda}_i)$ by $(-\tilde{\lambda}_i, -\tilde{\lambda}_{i-1})$, negating every principal turn: (iii). For (iv), all three turn determinants of a triangle equal $\det(\mu_2 - \mu_1, \mu_3 - \mu_1)$, which is nonzero for $P \in \mathcal{R}_3$ (three collinear points with nonzero edges have a doubled-back consecutive pair); so the three turns have one sign, their sum lies strictly between 0 and 3π in absolute value, and is a multiple of 2π . For (v), $2\pi|\text{rot}(P)| \leq \sum_i |\vartheta_i| < n\pi$. $\square$

Lemma 1.21 (rotation of uniform and one-dissent polygons). *Let $P \in \mathcal{R}_n$ have all turns nonzero.*

- (i) *If all turns are left, then $\text{rot}(P) \geq 1$; if all are right, then $\text{rot}(P) \leq -1$.*
- (ii) *If exactly one turn is right and all others are left, then $\text{rot}(P) \geq 1$; symmetrically with left and right exchanged, $\text{rot}(P) \leq -1$.*

Proof. (i) All $\vartheta_i > 0$ and there are at least three of them, so the sum is positive and $\text{rot}(P)$ is a positive integer. (ii) Let $\alpha \in (0, \pi)$ be the magnitude of the right turn and Π the sum of the left turns, so $\Pi - \alpha = 2\pi \text{rot}(P)$. If $\text{rot}(P) \leq 0$ then $\Pi \leq \alpha < \pi$. Cut the traversal after the right-turn vertex; the edge directions from there on are obtained by successive positive turns summing to $\Pi < \pi$, so all edge vectors lie in a closed angular sector of width $< \pi$, whose dual cone is nonempty: there is u with $\langle u, \tilde{\lambda}_i \rangle > 0$ for all i . Summing contradicts $\sum_i \tilde{\lambda}_i = 0$.

Both negative cases follow by reversing the traversal: this exchanges left and right turns and negates rotation by Lemma 1.20(iii). Applying the respective positive case to the reversed polygon proves each claim. $\square$

Remark 1.22 (the chirotope formula at a hull vertex). For comparison with Whitney's formula [31], let P be generic and suppose μ_1 is a vertex of the convex hull of $\{\mu_1, \dots, \mu_n\}$. Then

$$\text{rot}(P) = \chi_{n,1,2} + \sum_{\substack{\{i,j\} \in X(P) \\ 1 \leq i < j \leq n}} \chi_{i,i+1,j},$$

where the crossing pairs are written with $i < j$ as integers in $\{1, \dots, n\}$ — equivalently E_i is visited before E_j when the traversal starts at μ_1 . This is explanatory background only; no proof in this document uses this remark.

Remark 1.23 (the hull hypothesis cannot be dropped). Without the hull hypothesis the formula can fail: for the embedded dart $P = ((2, 1), (2, 4), (0, 0), (4, 0))$, traversed counterclockwise, $\text{rot}(P) = 1$, there are no crossings, and the turn at the reflex base vertex $\mu_1 = (2, 1)$ is -1 ; the formula returns -1 . With any of the three hull vertices as base it returns $+1$. The general formula with an arbitrary base vertex adds the winding numbers of the two regions adjacent to the base vertex in place of the turn sign [22]; at a hull vertex one adjacent region is the unbounded one, with winding number 0, and the general form reduces to the displayed one. The hull hypothesis is sufficient, not necessary: at a convex vertex of an embedded polygon the two adjacent winding numbers are 0 and τ_1 whether or not that vertex lies on the hull.

1.6 Cyclic shift and reversal

Lemma 1.24 (behaviour under shift and reversal). (i) *The chirotope identities are*

$$\chi_{ijk}(\sigma P) = \chi_{i+1,j+1,k+1}(P), \quad \chi_{ijk}(\bar{P}) = \chi_{2-i,2-j,2-k}(P).$$

The edge $E_i(\sigma P)$ is $E_{i+1}(P)$, and the edge $E_i(\bar{P})$ is $E_{1-i}(P)$ traversed backwards.

- (ii) σ and ρ map $\mathcal{U}_n$ onto itself and map chambers onto chambers; they preserve $X(P)$ up to relabelling of edges.

- (iii) $\tau_i(\sigma P) = \tau_{i+1}(P)$ and $\tau_i(\bar{P}) = -\tau_{2-i}(P)$; hence $\ell(\bar{P}) = n - \ell(P) - z(P)$, where $z(P) = \#\{i : \tau_i(P) = 0\}$; in particular $\ell(\bar{P}) = n - \ell(P)$ for $P \in \mathcal{U}_n$ (Lemma 1.6(ii)).

- (iv) for $P \in \mathcal{R}_n$: σP and $\bar{P}$ lie in $\mathcal{R}_n$, $\text{rot}(\sigma P) = \text{rot}(P)$ and $\text{rot}(\bar{P}) = -\text{rot}(P)$.

Proof. (i) is the definition; for the reversal, $\det(\mu_{2-j} - \mu_{2-i}, \mu_{2-k} - \mu_{2-i})$ is the determinant of the triple $(2-i, 2-j, 2-k)$ of P . (ii): (G1) and (G2) are conditions on the set of vertices and edges, which are preserved; both maps are linear homeomorphisms of $(\mathbb{R}^2)^n$. (iii): $\tau_i(\bar{P}) = \chi_{i-1, i+1}(\bar{P}) = \chi_{3-i, 2-i, 1-i}(P) = -\chi_{1-i, 2-i, 3-i}(P) = -\tau_{2-i}(P)$ by Lemma 1.4(i). Since $i \mapsto 2-i$ is a bijection of $\mathbb{Z}/n$, the turn multiset of $\bar{P}$ is the negative of that of P : the left turns of $\bar{P}$ are the right turns of P , of which there are $n - \ell(P) - z(P)$; on $\mathcal{U}_n$ every turn is ± 1 (Lemma 1.6(ii)), so $z(P) = 0$ there. (iv): σ permutes the edge vectors cyclically, $\tilde{\lambda}_i(\sigma P) = \tilde{\lambda}_{i+1}(P)$, and ρ negates and reverses them, $\lambda_i(\bar{P}) = -\tilde{\lambda}_{1-i}(P)$, so the consecutive pairs $(\lambda_{i-1}, \lambda_i)$ of σP are those of P and the consecutive pairs of $\bar{P}$ are the pairs $(-\tilde{\lambda}_k, -\tilde{\lambda}_{k-1})$ of P ; the condition of Definition 1.18 (every edge vector nonzero, no $\tilde{\lambda}_k$ a negative multiple of $\tilde{\lambda}_{k-1}$) is symmetric in the pair and invariant under negation, so it is preserved and $\sigma P, \bar{P} \in \mathcal{R}_n$. The identity for ρ is then Lemma 1.20(iii), and for σ the multiset of principal turns is unchanged. $\square$

Definition 1.25 (admissible pairs). A pair $(n, r) \in \mathbb{Z}^2$ is *admissible* if $n \geq 3$, $2|r| < n$ and $(n, r) \neq (3, 0)$. It is *minimal* if moreover $n = 2|r| + 1$ (for $r \neq 0$) or $(n, r) = (4, 0)$. For an admissible pair the *fibre* is $\mathcal{U}_{n,r} = \{P \in \mathcal{U}_n : \text{rot}(P) = r\}$; a vertex of a fibre is a generic polygon, a point of the quotient.

Lemma 1.26 (nonempty fibres). *For integers n, r , a generic polygon with n vertices and rotation r exists exactly when $n \geq 3$, $2|r| < n$, and $(n, r) \neq (3, 0)$. Here genericity means that every determinant of three distinct vertices is nonzero and no three distinct edge interiors meet. Moreover, for $n \geq 3$ the generic polygons are dense in $(\mathbb{R}^2)^n$: every nonempty open set of labelled n -tuples contains a generic polygon.*

Proof in Part II, page 80.




Figure 6: The admissible (n, r) lattice at the $n = 8$ induction step; the shaded rows are the completed induction data. The boundary portraits display the star values evaluated in Section 7; their higher Catalan values are not premises of Theorem 6.7.

1.7 Walls

Definition 1.27 (wall germ; sides). A *wall germ* is a continuous map $P: (-\varepsilon, \varepsilon) \rightarrow (\mathbb{R}^2)^n$ such that $P(t) \in \mathcal{U}_n$ for $t \neq 0$ and $P(0) \notin \mathcal{U}_n$. Its *sides* P_- and P_+ are the chambers containing $P((-\varepsilon, 0))$ and $P((0, \varepsilon))$; each of these two sets is connected and generic, hence lies in one chamber. For a

function F that is constant on chambers, $F(P_\pm)$ denotes its value on the respective side. At the centre $P(0)$ put

$$\begin{aligned} Z_{\text{pt}} &= \{\{i, j, k\} \text{ pairwise distinct} : \chi_{ijk}(P(0)) = 0\}, \\ Z_c &= \{\{e, f, g\} \text{ pairwise remote} : \text{relint } E_e \cap \text{relint } E_f \cap \text{relint } E_g \neq \emptyset \text{ at } t = 0\}. \end{aligned}$$

A real function ϕ of the germ *changes sign at 0* if there is $\delta > 0$ with $\phi(P(t))\phi(P(-t)) < 0$ for $0 < t < \delta$.

Lemma 1.28 (sides of a triple wall). *Let a wall germ have $Z_{\text{pt}} = \emptyset$ and $Z_c = \{\{e, f, g\}\}$. Then for small $t \neq 0$ the three pairs $\{e, f\}, \{e, g\}, \{f, g\}$ are crossings, their three crossing points are distinct, and on each of E_e, E_f, E_g its two crossings occupy adjacent positions among the crossing visits on that edge.*

Proof. At the centre (G1) holds. In particular all vertices and edges are nonzero, adjacent edges meet only at their common endpoint, and no vertex lies on the line of a nonincident edge. The specified common point q lies in all three relative interiors. Two of their directions cannot be parallel: their lines would then coincide, forcing an endpoint of one remote edge onto the line of the other, contrary to (G1). Thus all three central intersections are transverse. Lemma 1.9 makes them persistent interior crossings. On either punctured side (G2) makes the three points distinct.

In fact every remote edge pair at the centre is either disjoint or a transverse interior pair: endpoint incidence or overlapping collinear remote edges would contradict (G1). Thus the same stability lemma applies to all remote pairs and gives a fixed finite crossing set with continuous parameters near the centre.

Suppose, contrary to the required adjacency on E_e , that for a sequence $t_\nu \rightarrow 0$ another crossing lies between x_{ef} and x_{eg} on E_e . There are finitely many crossing labels, so after taking a subsequence this crossing has one fixed other edge $h \notin \{e, f, g\}$. The two bounding parameters converge to the parameter of q ; the intervening parameter must have the same limit. Continuity of its intersection parameters puts q in the relative interiors of E_e and E_h . It is also interior to E_f . The edges E_h, E_f cannot be adjacent: under (G1) adjacent edges meet only at their shared endpoint, whereas q is interior to both. Hence e, f, h are pairwise remote and define an element of Z_c distinct from $\{e, f, g\}$, a contradiction. The same argument on E_f and E_g proves all three adjacency assertions on one common punctured interval. No sign-change hypothesis was used for this lemma. $\square$

Definition 1.29 (the named walls). A wall germ is *simple of one of the following types*; the types are mutually exclusive.

- (F) *Flat at j ($n \geq 4$):* $Z_{\text{pt}} = \{\{j-1, j, j+1\}\}$, $Z_c = \emptyset$, $\mu_j(0)$ lies strictly between $\mu_{j-1}(0)$ and $\mu_{j+1}(0)$, and τ_j changes sign at 0. The *right side* is the side with $\tau_j = -1$ and the *left side* the one with $\tau_j = +1$.
- (K) *Cusp at j ($n \geq 4$):* $Z_{\text{pt}} = \{\{j-1, j, j+1\}\}$, $Z_c = \emptyset$, the three points $\mu_{j-1}(0), \mu_j(0), \mu_{j+1}(0)$ are collinear with $\mu_j(0)$ outside the closed segment $[\mu_{j-1}(0), \mu_{j+1}(0)]$, and τ_j changes sign at 0. The *newborn pair* is $\{j-1, j+1\}$ if $\mu_{j+1}(0)$ lies between $\mu_{j-1}(0)$ and $\mu_j(0)$, and $\{j-2, j\}$ if $\mu_{j-1}(0)$ lies between $\mu_j(0)$ and $\mu_{j+1}(0)$. The *loop side* P_{loop} is the side on which the newborn pair is a crossing and the *no-loop side* P_{no} the other. The cusp is *empty* if on the loop side the two visits of the newborn crossing are cyclically adjacent in the Gauss word.
- (V) *Vertex-edge at $(M; a)$:* $M \notin \{a-1, a, a+1, a+2\}$, $Z_{\text{pt}} = \{\{a, a+1, M\}\}$, $Z_c = \emptyset$, $\mu_M(0)$ lies in the relative interior of $E_a(0)$, and $\chi_{a, a+1, M}$ changes sign at 0. The germ is of *bigon* type if $\chi_{a, a+1, M-1}(P(0)) = \chi_{a, a+1, M+1}(P(0))$ and of *sliding* type otherwise. The *contact sign* is $s = \chi_{a, a+1, M}(P_-)$.

- (T) *Triple at $\{e, f, g\}$* : $Z_{\text{pt}} = \emptyset$, $Z_c = \{\{e, f, g\}\}$, and on each of the three edges the two crossing parameters of the other two edges (which exist on both sides by Lemma 1.28) have a difference that changes sign at 0.
- (E) *Exterior extension at $(M; a)$* : $M \notin \{a-1, a, a+1, a+2\}$, $Z_{\text{pt}} = \{\{a, a+1, M\}\}$, $Z_c = \emptyset$, $\mu_M(0)$ lies on the line of $E_a(0)$ outside the closed segment $E_a(0)$, and $\chi_{a, a+1, M}$ changes sign at 0.
- (C) *Pure cut at $\{i, j, k\}$* : no two of i, j, k are consecutive modulo n , $Z_{\text{pt}} = \{\{i, j, k\}\}$, $Z_c = \emptyset$, and χ_{ijk} changes sign at 0.

The walls (E) and (C) are called *silent*. The type is determined by the centre; the sign-change conditions exclude tangential germs.

Lemma 1.30 (flat-side geometry). *Let $n \geq 4$ and let $P(t) = (\mu_1(t), \dots, \mu_n(t))$ be continuous for $|t| < \varepsilon$. Indices are cyclic modulo n ; put $E_i = [\mu_i, \mu_{i+1}]$ and $d_i = \mu_{i+1} - \mu_i$. For $t \neq 0$ assume that no three distinct vertices are collinear and no three edge interiors concur. At $t = 0$ suppose that the only collinear triple is $\{j-1, j, j+1\}$, that no three pairwise remote edge interiors concur, and that μ_j is strictly between μ_{j-1} and μ_{j+1} . Suppose also that $\tau_j(t) = \text{sgn det}(d_{j-1}(t), d_j(t))$ changes sign at zero. At a transverse crossing of E_e, E_f , the positive over/under data mean that E_e is over E_f exactly when $\text{det}(d_e, d_f) > 0$. Then, after shrinking the interval:*

- (i) *The central polygon has nonzero edges and no antiparallel consecutive directions. Its only zero turn is at j , where the two directions are positive multiples. Every other point-triple determinant has constant nonzero sign on the whole interval.*
- (ii) *The set of remote edge pairs that cross is constant, including at zero. Each crossing is interior and transverse, no crossing is a vertex, all crossing points are distinct, and the crossing visits have constant order on every edge. Hence their cyclic Gauss word, pairing, interlacement graph, and positive over/under data are constant. Here the central crossing data are read from the segments of $P(0)$ exactly as for a generic polygon: a crossing is a transverse intersection of the relative interiors of two remote edge segments, its two visits are ordered along each edge by the intersection parameters, and the Gauss word, pairing, interlacement graph and positive over/under data are formed from these visits as in Definitions 1.7, 1.12 and 1.13; no function on the generic locus is evaluated at the flat centre.*
- (iii) *The deletion $Q = P(0) \setminus j$ is generic. The directions $u = d_{j-1}(0)$ and $v = d_j(0)$ fuse to $w = \mu_{j+1}(0) - \mu_{j-1}(0)$, and there is $0 < \lambda < 1$ with*

$$u = \lambda w, \quad v = (1 - \lambda)w. \quad (2)$$

Replacing either incident edge in a crossing label by the fused edge gives a bijection with the crossings of Q , preserving their cyclic visit order, pairing, determinant signs, and positive over/under data. The deletions $P(t) \setminus j$ are generic on a smaller interval and lie in the chamber of Q .

Proof in Part II, page 82.

Lemma 1.31 (the sides of the named walls). *Let the germ be simple of type (F), (V), (T), (E) or (C). Then $P(0) \in \mathcal{R}_n$. After shrinking the interval, every chirotope outside Z_{pt} is nonzero with constant sign. Every crossing not involved in the stated local degeneration persists, and the order of any two persistent visits with distinct central parameters is constant. Moreover:*

(F) The crossing set and Gauss word are the same on both sides and at the centre.

(V) In the bigon branch,

$$X(P_+) \triangle X(P_-) = \{\{a, M-1\}, \{a, M\}\},$$

with both pairs crossing on one side and neither on the other. In the sliding branch exactly one of these two pairs crosses on each side, and the two sides exchange them. No other crossing is created or lost, and no order changes outside the contact neighbourhood.

(T) The crossing set is the same on both sides. The only visit-order changes are the exchanges of the two triangle crossings on each of E_e, E_f, E_g .

(E),(C) The crossing set and Gauss word are the same on both sides and at the centre.

Here, when crossing data are mentioned at a nongeneric centre, they mean the transverse intersections of the relative interiors of the actual segments and their traversal order. This does not evaluate a function defined only on generic polygons at that centre.

Proof in Part II, page 83.

Lemma 1.32 (sides of a cusp). *Let $n \geq 4$ and let $P(t) = (\mu_1(t), \dots, \mu_n(t))$, $|t| < \varepsilon$, be continuous. Indices are read modulo n , and $E_i = [\mu_i, \mu_{i+1}]$, $\tilde{\lambda}_i = \mu_{i+1} - \mu_i$ and $\tau_i = \text{sgn det}(\tilde{\lambda}_{i-1}, \tilde{\lambda}_i)$. Assume that $P(t)$ is generic for $t \neq 0$: no three distinct vertices are collinear and no three edge interiors concur. At $t = 0$ assume that the only collinear triple of distinct vertex indices is $\{j-1, j, j+1\}$, that no three pairwise remote edge interiors concur, that μ_j lies outside the closed segment $[\mu_{j-1}, \mu_{j+1}]$, and that τ_j changes sign. These are the simple-cusp hypotheses in the present conventions.*

Use the following data, with the betweenness conditions imposed at $t = 0$:

case	middle point on the line	(f, g)	(c_1, c_2)	$\Delta(t)$
A	$\mu_{j+1} \in (\mu_{j-1}, \mu_j)$	$(j-1, j+1)$	$(j, j+1)$	$\text{det}(\tilde{\lambda}_{j-1}, \tilde{\lambda}_{j+1})$
B	$\mu_{j-1} \in (\mu_j, \mu_{j+1})$	$(j-2, j)$	$(j-1, j)$	$\text{det}(\tilde{\lambda}_{j-2}, \tilde{\lambda}_j)$

(3)

Exactly one case applies. After shrinking the parameter interval:

- (i) Δ is nonzero with constant sign. The remote pair $\{f, g\}$ crosses precisely on the side with $\tau_j = -\text{sgn } \Delta$, called the loop side; it does not cross on the side with $\tau_j = \text{sgn } \Delta$, called the no-loop side. No other crossing pair changes.
- (ii) On the no-loop side the turns at c_1, c_2 have opposite signs. On the loop side both have sign τ_j .
- (iii) On the loop side, one oriented traversal arc between the two visits of the newborn crossing consists of a terminal portion of E_f , the complete edge between c_1 and c_2 , and an initial portion of E_g . Its only polygon vertices are μ_{c_1}, μ_{c_2} . This arc is allowed to contain other crossing visits. If the two newborn visits are cyclically adjacent in the Gauss word, this particular arc contains no other crossing visit. In that case the complete edge between c_1, c_2 carries no crossing on either side of the wall.
- (iv) With rotation defined by the sum of principal turns,

$$\text{rot}(P_{\text{loop}}) - \text{rot}(P_{\text{no}}) = \tau_j(P_{\text{loop}}) \in \{-1, 1\}. \quad (4)$$

Proof in Part II, page 85.

Definition 1.33 (deletion and halves). For $n \geq 4$ and $j \in \mathbb{Z}/n$, the *deletion* $P \setminus j$ is the $(n-1)$ -tuple obtained by omitting μ_j , with the induced cyclic order. At a simple vertex-edge wall at $(M; a)$ with $b = a + 1$, the *halves* at the centre $P^0 = P(0)$ are the polygons

$$\lambda_1 = (\mu_M, \mu_{M+1}, \dots, \mu_a), \quad \lambda_2 = (\mu_M, \mu_b, \mu_{b+1}, \dots, \mu_{M-1}),$$

all vertices taken at $t = 0$: λ_1 follows the edges $E_M, \dots, E_{a-1}$ of P^0 and closes with the segment $[\mu_a, \mu_M] \subset E_a$; λ_2 opens with $[\mu_M, \mu_b] \subset E_a$, follows $E_b, \dots, E_{M-2}$ and closes with E_{M-1} .

Lemma 1.34 (deletions and halves are generic). (i) *At a simple flat wall at j with $n \geq 4$, $P(0) \setminus j$ is generic, and for small $t \neq 0$ the deletions $P(t) \setminus j$ lie in the chamber of $P(0) \setminus j$.*

(ii) *At a simple vertex-edge wall at $(M; a)$, both halves of Definition 1.33 are generic and have between 3 and $n - 2$ vertices. The central parent is automatically regular.*

Proof in Part II, page 87.

1.8 Relative general position in the regular locus

Lemma 1.35 (the polynomial controls for SM genericity). *Write the scalar vertex coordinates as $x_1, \dots, x_{2n}$. For each unordered triple of distinct vertex indices use one ordered representative of*

$$\Delta_{ijk} = \det(\mu_j - \mu_i, \mu_k - \mu_i). \quad (5)$$

For each unordered triple of pairwise remote edges use one ordered representative of

$$T_{efg} = \det \begin{pmatrix} \ell_e \\ \ell_f \\ \ell_g \end{pmatrix}, \quad \ell_i = (-d_{i,y}, d_{i,x}, \det(\mu_i, d_i)), \quad d_i = \mu_{i+1} - \mu_i. \quad (6)$$

Let $\mathcal{F}$ be this finite family. Every member is nonzero and irreducible over $\mathbb{R}$, and distinct named members are nonassociate. Each is affine in every scalar coordinate separately. If F depends on x_j , write $F = ax_j + b$; then a is a nonzero polynomial in the other coordinates. If distinct $F = ax_j + b$ and $G = cx_j + d$ both depend on x_j , the polynomial $ad - bc$ is nonzero. At a specialization with $a \neq 0$, the root of F in x_j is simple; at a specialization with $ad - bc \neq 0$, the two polynomials have no common root.

Proof in Part II, page 87.

Theorem 1.36 (relative general position and the wall dictionary). *Let $\gamma : [0, 1] \rightarrow \mathcal{R}_n$ be continuous with generic endpoints, and let $\delta > 0$. There is a piecewise-affine path in $\mathcal{R}_n$ with the same labelled endpoints, uniformly within δ of γ , which is collision-free and whose nongeneric parameters are finitely many isolated simple wall germs, each of exactly one of the types (F), (V) bigon, (V) sliding, (T), (E), (C). No cusp occurs. Every occurrence label is retained.*

Proof in Part II, page 89.

2 The tree coefficient A_g

Throughout this section P is a polygon satisfying (G1); (G2) is not used. The object is the finite sum in the rooted vertex expansion and tree formula of the letter, written here as a recursion over interval compositions, which is the form in which every law is proved.

2.1 Roots, boundary words, gates

Definition 2.1 (root and boundary word). A *root* of a polygon is one of its edges; on a representative it is an edge index $g \in \mathbb{Z}/n$, the root edge being E_g from μ_g to μ_{g+1} , and the shift carries the root g of P to the root $g - 1$ of σP . Everything in this section is a function of the pair (polygon, root). The *boundary word* of (P, g) is the sequence of vertices

$$a_0 = \mu_{g+1}, a_1 = \mu_{g+2}, \dots, a_N = \mu_{g+n} = \mu_g, \quad N = n - 1,$$

so that consecutive boundary vertices a_{k-1}, a_k are joined by the *leaf* edge E_{g+k} , $1 \leq k \leq N$; the leaves are the $n - 1$ edges other than the root, in traversal order. An *interval* is a pair $[i, j]$ of integers $0 \leq i < j \leq N$; it has $j - i$ leaves. A *composition* of $[i, j]$ is a sequence $\pi = (r_0 < r_1 < \dots < r_s)$ with $r_0 = i$, $r_s = j$, $s \geq 1$; its *parts* are the intervals $[r_{k-1}, r_k]$, and $|\pi| = s$.

Definition 2.2 (gates). Let $\pi = (r_0 < \dots < r_s)$ be a composition of $[i, j]$. At each interior cut r_k , $1 \leq k \leq s - 1$, put

$$d_k(\pi) = \chi(a_{r_{k+1}}, a_{r_k}, a_{r_{k-1}}), \quad h_k(\pi) = \chi(a_j, a_{r_k}, a_i),$$

where $\chi(x, y, z) = \text{sgn det}(y - x, z - x)$ for points x, y, z ; d_k is the *near sign* and h_k the *far sign*. The *ordinary weight* and the *root weight* of π are

$$V^+(\pi) = \prod_{k=1}^{s-1} d_k \Theta(-h_k d_k), \quad V^-(\pi) = \prod_{k=1}^{s-1} d_k \Theta(h_k d_k),$$

with empty products equal to 1 (so $V^\pm(\pi) = 1$ when $s = 1$).

Lemma 2.3 (gates are defined on (G1)). *If P satisfies (G1) then every near and far sign is ± 1 and every Θ in Definition 2.2 is evaluated at ± 1 . On that domain*

$$d \Theta(-hd) = \frac{d - h}{2}, \quad d \Theta(hd) = \frac{d + h}{2} \quad (d, h \in \{\pm 1\}),$$

and $V^+(\pi) = 0$ for every composition with exactly two parts.

Proof. The three boundary vertices entering a gate are pairwise distinct vertices of P (the cut positions are distinct and $a_0 \neq a_N$ by (G1)), so the signs are nonzero. The two identities are checked on the four sign pairs. For $s = 2$ the only cut is r_1 and $d_1 = \chi(a_j, a_{r_1}, a_i) = h_1$, so $\Theta(-h_1 d_1) = \Theta(-1) = 0$. $\square$

2.2 The coefficient

Definition 2.4 (open sums and the tree coefficient). Let P satisfy (G1) and fix a root g . Define $b_{[i,j]} \in \mathbb{Z}$ for every interval by recursion on the number of leaves:

$$b_{[i,i+1]} = 1, \quad b_{[i,j]} = - \sum_{\substack{\pi \text{ composition of } [i,j] \\ |\pi| \geq 2}} V^+(\pi) \prod_{[r_{k-1}, r_k] \in \pi} b_{[r_{k-1}, r_k]} \quad (j - i \geq 2).$$

The *tree coefficient* of P at the root g is

$$A_g(P) = \sum_{\pi \text{ composition of } [0, N]} V^-(\pi) \prod_{[r_{k-1}, r_k] \in \pi} b_{[r_{k-1}, r_k]},$$

the sum including the one-part composition, whose term is $b_{[0, N]}$. The letter's $A(P)$ is $A_n(P)$, rooted at the last edge E_n from μ_n to μ_1 .

Lemma 2.5 (equivalence with the plane-tree sum). *Let $\mathfrak{T}_g(P)$ be the set of rooted plane trees whose leaves are the $n - 1$ leaf edges in traversal order, whose root vertex has one or more children, and whose other internal vertices (the ordinary vertices) have at least two children each; every internal vertex is labelled by the composition of its leaf interval into the leaf intervals of its children. Then*

$$A_g(P) = \sum_{T \in \mathfrak{T}_g(P)} (-1)^{N_+(T)} V^-(\pi_{\text{root}}) \prod_{v \text{ ordinary}} V^+(\pi_v),$$

where $N_+(T)$ is the number of ordinary vertices. The identity is formal: the bijection of the proof matches the terms of the recursion of Definition 2.4 with the trees term by term, so it holds with the gate values $V^\pm(\pi)$ replaced by independent formal variables, as an identity in the polynomial ring over $\mathbb{Z}$ in those variables, and specializes to every evaluation of them. This is the tree formula in the letter, whose vertex weights are Definition 2.2 with $m = s$, a_k the boundary vertices of the cuts, and whose sign $(-1)^{N_+}$ is the one minus per ordinary vertex.

Proof. Removing the top vertex of an ordinary tree on $[i, j]$ leaves one subtree per part of its composition; grafting is the inverse. The minus sign of the top vertex is the minus sign in the recursion for b . The root vertex carries V^- , no sign, and may have one child, which is the one-part composition. $\square$

Remark 2.6. Ordinary binary vertices contribute 0 (Lemma 2.3), which is why the Berends–Giele recursion in the letter sums over $m \geq 3$. The letter’s “ V^+ propagators are discarded” is the requirement that ordinary vertices have at least two children.

2.3 Chamber constancy and the resummation

Proposition 2.7 (chirotope chambers). *If two polygons satisfying (G1) have the same chirotope, then for every root g every composition weight, every $b_{[i,j]}$ and A_g agree. In particular A_g is constant on every labelled chamber (the root fixed), and constant along any path of (G1)-polygons with constant chirotope.*

Proof. Every quantity in Definitions 2.2 and 2.4 is a function of the signs $\chi(a_x, a_y, a_z)$, which are chirotope entries. A labelled chamber is a component of $\mathcal{U}_n$, an open set (Proposition 1.11), hence path connected, and the chirotope is constant along every path in $\mathcal{U}_n$ (Proposition 1.11); so A_g is constant on each labelled chamber. $\square$

Definition 2.8 (the near–far arrays). Over a commutative ring in which 2 is invertible, let $D(x, y, z)$ and $H(x, y, z)$ be arbitrary scalars indexed by increasing triples $0 \leq x < y < z \leq N$. For an array $X = (X_{ij})_{0 \leq i < j \leq N}$ put

$$\Phi_{D,H}(X)_{ij} = \sum_{\pi=(i=r_0<\dots<r_s=j)} \left(\prod_{k=1}^{s-1} \frac{D(r_{k-1}, r_k, r_{k+1}) - H(i, r_k, j)}{2} \right) \prod_{k=1}^s X_{r_{k-1}r_k},$$

$F_H = \Phi_{0,H}$, $G_D = \Phi_{D,0}$, and $E_{ij} = 1$ if $j = i + 1$ and 0 otherwise. The *geometric arrays* are $D(x, y, z) = \chi(a_z, a_y, a_x)$ and $H(x, y, z) = \chi(a_z, a_y, a_x)$ with the boundary vertices of (P, g) .

Lemma 2.9 (factorization and far-only output). *(i) Over any commutative ring in which 2 is invertible, $\Phi_{D,H}$ has a two-sided polynomial inverse and*

$$\Phi_{D,H} = F_H \circ G_D. \tag{7}$$

(ii) With the geometric arrays, $\Phi_{D,H}(b) = E$ and $A_g(P) = \Phi_{D,-H}(b)_{0N}$. Consequently, with $c = F_H^{-1}(E)$,

$$A_g(P) = F_{-H}(c)_{0N}. \quad (8)$$

The complete output is independent of D when H is held fixed.

(iii) For an interval $J = [i, j]$ put $\mathcal{E}_J = F_H(c)_J$ and $\mathcal{B}_J = F_{-H}(c)_J$. Then $\mathcal{E}_J = 0$ if $j - i \geq 2$, and $\mathcal{E}_J = \mathcal{B}_J = 1$ if $j - i = 1$. For $j - i \geq 2$, if $(a_i, \dots, a_j)$ satisfies (G1), $\mathcal{B}_J$ is its tree coefficient rooted at its closing edge from a_j to a_i .

Proof in Part II, page 92.

Theorem 2.10 (Response to one collinear vertex triple). Let $t \mapsto P(t)$ be a wall germ with $Z_c = \emptyset$, $Z_{pt} = \{K\}$ a single triple of pairwise distinct vertices, and χ_K changing sign at zero. This includes types (F), (K), (V), (E), and (C). Cut at a fixed physical root g and let $x < y < z$ be the critical boundary positions. Put $I_* = [x, z]$, $L = [x, y]$, $R = [y, z]$, $F = [0, N]$, and

$$\delta = \frac{H^+(x, y, z) - H^-(x, y, z)}{2} \in \{-1, 1\}. \quad (9)$$

Write the three wall points as $a_v = p_* + \xi_v \omega$ with $\omega \neq 0$ and distinct ξ_v . Set

$$\epsilon_L = \operatorname{sgn} \frac{\xi_y - \xi_x}{\xi_z - \xi_x}, \quad \epsilon_R = \operatorname{sgn} \frac{\xi_z - \xi_y}{\xi_z - \xi_x}. \quad (10)$$

For $J = L, R$ use the unchanged interval values of Lemma 2.9, and put

$$\mathcal{U}_J = \begin{cases} \mathcal{E}_J, & \epsilon_J = 1, \\ \mathcal{B}_J, & \epsilon_J = -1, \end{cases} \quad \mathcal{V}_J = \begin{cases} \mathcal{B}_J, & \epsilon_J = 1, \\ \mathcal{E}_J, & \epsilon_J = -1. \end{cases} \quad (11)$$

If $I_* \neq F$, let Q be the polygon obtained at $t = 0$ by replacing the boundary arc $a_x, \dots, a_z$ with the single edge from a_x to a_z . Then Q satisfies (G1), has at least three vertices, retains the physical root g , and

$$A_g(P_+) - A_g(P_-) = \delta \mathcal{U}_L \mathcal{U}_R A_g(Q). \quad (12)$$

If $I_* = F$, instead

$$A_g(P_+) - A_g(P_-) = \delta (\mathcal{U}_L \mathcal{U}_R + \mathcal{V}_L \mathcal{V}_R). \quad (13)$$

Here $P_{\pm}$ are on sufficiently close punctured sides of the germ; all gap and contracted coefficients are evaluated at their nondegenerate wall data, equivalently at the same signs on either sufficiently close side.

Proof in Part II, page 92.

2.4 The laws

Definition 2.11 (induced roots). At a simple flat or cusp wall at j , the *deletion map* on roots is

$$D_j(g) = \begin{cases} \text{the fused edge } [\mu_{j-1}, \mu_{j+1}] \text{ of } P \setminus j, & g \in \{j-1, j\}, \\ \text{the edge of } P \setminus j \text{ with the same endpoints as } E_g, & \text{otherwise.} \end{cases}$$

At a simple vertex-edge wall at $(M; a)$, with $d_1 = [\mu_a, \mu_M]$ the closing edge of λ_1 and $d_2 = [\mu_M, \mu_{a+1}]$ the opening edge of λ_2 , the *half map* on roots is

$$H(g) = \begin{cases} (d_1, d_2), & g = a, \\ (E_g \text{ in } \lambda_1, d_2), & g \in \{M, M+1, \dots, a-1\}, \\ (d_1, E_g \text{ in } \lambda_2), & g \in \{a+1, \dots, M-1\}. \end{cases}$$

Theorem 2.12 (flat law for A_g ; W2). *At a simple flat wall at j with $n \geq 4$, for every root g ,*

$$A_g(P_{\text{right}}) - A_g(P_{\text{left}}) = A_{D_j(g)}(P(0) \setminus j), \quad (14)$$

where P_{right} is the side with $\tau_j = -1$, and D_j is the physical deletion map of Definition 2.11.

Proof in Part II, page 94.

Theorem 2.13 (cusp law for A_g ; W4). *At a simple cusp wall at j with $n \geq 4$ whose deletion $Q = P(0) \setminus j$ satisfies (G1), for every root g ,*

$$A_g(P_{\text{loop}}) - A_g(P_{\text{no}}) = -\kappa A_{D_j(g)}(Q), \quad \kappa = \text{rot}(P_{\text{loop}}) - \text{rot}(P_{\text{no}}) \in \{-1, 1\}. \quad (15)$$

Proof in Part II, page 94.

Theorem 2.14 (vertex-edge law for A_g ; W1). *At a simple vertex-edge wall at $(M; a)$, of bigon or sliding type, let $s = \chi_{a, a+1, M}(P_-)$ and let λ_1, λ_2 be the halves in Definition 1.33. For every physical root g with $H(g) = (h_1, h_2)$ as in Definition 2.11,*

$$A_g(P_+) - A_g(P_-) = s A_{h_1}(\lambda_1) A_{h_2}(\lambda_2). \quad (16)$$

Proof in Part II, page 95.

Theorem 2.15 (triple law and silence for A_g ; W3 and silent walls). *(i) At a simple triple wall every composition weight, every $b_{[i,j]}$ and A_g are identical on the two sides, for every root g .*

(ii) At a simple exterior-extension wall or a simple pure cut, $A_g(P_+) = A_g(P_-)$ for every root g .

Proof in Part II, page 96.

Figure 2 in the introduction illustrates the vertex-edge and flat wall laws.

Figures 2, 7 and 8 illustrate W1–W4. The lower portraits in the first plate are the smoothed halves or deletion entering the laws. The cusp plate includes a spectator strand and retains the nondegenerate-deletion hypothesis of Theorem 2.13.



Figure 7: The two chambers adjacent to a W3 wall; the circled triangle collapses at the wall.

2.5 Continuation through silent configurations

Theorem 2.16 (continuation across silent walls). *For every root g there is a unique function $\tilde{A}_g$ on the weakly generic locus $\mathcal{W}_n$ that is constant on visible chambers and agrees with A_g on $\mathcal{U}_n$.*

Proof in Part II, page 96.




Figure 8: A spectator strand passing through the W4 cusp region.

Remark 2.17 (Amplitude on the silent locus). On the silent locus $\mathcal{S}_n = \mathcal{W}_n \setminus \mathcal{U}_n$ the tree sum of Definition 2.4 is not evaluated: its gates are step functions of chirotope entries that vanish there, and no value assigned to a gate at zero makes the literal sum independent of the root. An exact six-gon with one exterior collinearity has literal root values $(-1, -2t, -t, -t, -1, -1)$ when every gate at zero is given the value t , while the continuation of Theorem 2.16 takes the value -1 at every root. The letter's A on its generic locus, which is $\mathcal{W}_n$, is therefore *defined* on $\mathcal{S}_n$ as that continuation, in Section 6, once root independence of the continuation is available; nothing in this section assigns a value to a gate at zero.

Lemma 2.18 (multi-affine form). *Fix n and a root index g on labelled representatives. Replace each cut factor in the finite tree sum by $(d - h)/2$ at an ordinary vertex and $(d + h)/2$ at the root, interpreting each ordered chirotope as the signed canonical variable X_{ijk} , $1 \leq i < j < k \leq n$. The resulting polynomial $A_g \in \mathbb{Q}[X_{ijk} : i < j < k]$ is affine in each variable and satisfies $A_g(\chi(P)) = A_g(P)$ on $(G1)$ -polygons.*

Within a tree, each unordered vertex triple occurs in at most one cut factor. At a binary vertex its near and far entries coincide: the ordinary factor is zero, whereas the root factor is that single entry. Distinct internal vertices can have the same leaf interval only when they are the unary root and its sole child; the unary root contributes an empty product.

[Proof in Part II, page 97.](#)

Lemma 2.19 (A nontrivial gap on a silent line). *Let $a_0, \dots, a_N$ be the boundary word of a weakly generic polygon. Suppose $r_0 < \dots < r_k$, $k \geq 2$, index collinear vertices. Write their distinct line coordinates as $t_0, \dots, t_k$ and put*

$$\epsilon_j = \operatorname{sgn} \frac{t_j - t_{j-1}}{t_k - t_0}, \quad 1 \leq j \leq k. \quad (17)$$

At least one gap has $\epsilon_j = 1$ and $r_j - r_{j-1} \geq 2$. If $r_0 = 0$ and $r_k = N$, at least one gap also has $\epsilon_j = -1$ and $r_j - r_{j-1} \geq 2$.

[Proof in Part II, page 98.](#)

Lemma 2.20 (Formal cancellation of all silent entries). *For the boundary word of any $P \in \mathcal{W}_n$, let H^0 be its geometric far array, with its zero entries equal to zero. Form H by replacing every zero entry of H^0 by its own indeterminate, keeping all nonzero entries fixed. Over the polynomial ring in those indeterminates, put $c = F_H^{-1}(E)$ and $c^0 = F_{H^0}^{-1}(E)$. Then each coordinate c_{ij} equals the corresponding constant c_{ij}^0 , and*

$$F_{-H}(c)_{0N} = F_{-H^0}(c^0)_{0N}. \quad (18)$$

No independence of the zero-determinant differentials is assumed.

[Proof in Part II, page 98.](#)

Corollary 2.21 (The continuation is the polynomial form). *On all of $\mathcal{W}_n$, for every root g ,*

$$\tilde{A}_g(P) = A_g(\chi(P)). \quad (19)$$

After the nonzero chirotope entries are fixed at their values at P , the polynomial is constant in all remaining entries. Thus its nonconstant terms in these silent variables cancel, and no value of $\Theta(0)$ has to be chosen.

Proof in Part II, page 99.

2.6 The soft theorem

Definition 2.22 (soft insertion). Let P satisfy (G1), fix a vertex index j and a vector $q \neq 0$. The *soft insertion* of q at j is the family of $(n+1)$ -gons

$$P_\varepsilon^{(j,q)} = (\mu_1, \dots, \mu_j, \mu_j + \varepsilon q, \mu_{j+1}, \dots, \mu_n), \quad \varepsilon > 0,$$

with the new vertex $\mu_* = \mu_j + \varepsilon q$ inserted immediately after μ_j ; its edges are $\tilde{\lambda}_{j-1}$, the *soft edge* εq , and the *return edge* $\tilde{\lambda}_j - \varepsilon q$, the others unchanged. The *attachment signs* are

$$\chi_- = \chi_{*,j,j-1}(P_\varepsilon) = -\operatorname{sgn} \det(\tilde{\lambda}_{j-1}, q), \quad \chi_+ = \chi_{*,j,j+1}(P_\varepsilon) = -\operatorname{sgn} \det(q, \tilde{\lambda}_j),$$

the two equalities holding for every $\varepsilon > 0$ (they are the turns at the two ends of the soft edge, with the sign reversed). The vector q is *admissible* if $\det(\tilde{\lambda}_{j-1}, q) \neq 0$, $\det(q, \tilde{\lambda}_j) \neq 0$, and $\det(q, \mu_k - \mu_j) \neq 0$ for all $k \neq j$.

Lemma 2.23 (the soft family is generic). *Let P be generic, j a vertex, and q admissible in Definition 2.22. Put $M = \mu_j$, $A = \mu_{j-1}$, $B = \mu_{j+1}$, $u = M - A$, $v = B - M$, and $\tau = \tau_j(P)$. Write $M_\varepsilon = M + \varepsilon q$ and $D_\varepsilon = v - \varepsilon q$. There is $\varepsilon_0 > 0$ such that, for every $0 < \varepsilon < \varepsilon_0$:*

- (i) $P_\varepsilon = P_\varepsilon^{(j,q)}$ is generic, and all these polygons lie in one chamber. The soft edge meets other edges only at its incident endpoints.
- (ii) Every original vertex other than M retains its original turn sign. The turns at M and M_ε are $-\chi_-$ and $-\chi_+$, respectively. The unchanged edge directions are fixed, $D_\varepsilon \rightarrow v$, and $\operatorname{sgn} \det(u, D_\varepsilon) = \tau$, while $\operatorname{sgn} \det(D_\varepsilon, u) = -\tau$.
- (iii) Identify each unchanged edge with its parent edge and the return edge with the parent edge E_j . Every parent crossing persists. Its point and its parameters on the corresponding edges converge to the parent values. Inherited visits retain their order along each directed edge, and the determinant sign of the ordered edge directions at each inherited crossing is unchanged.
- (iv) If $\chi_- = \chi_+ = -\tau$ (same-sign sector), or $\chi_- \neq \chi_+$ (mixed sector), these are all the crossings and the Gauss word is the parent's Gauss word. If $\chi_- = \chi_+ = \tau$ (loop sector), there is exactly one additional crossing y , between the return edge and E_{j-1} . It tends to M . Its incoming-edge visit a and return-edge visit b bound the oriented traversal arc from a through M, M_ε to b , containing no other crossing visit. Thus the two newborn visits are cyclically adjacent; deleting them gives the parent's Gauss word.

Proof in Part II, page 100.




Figure 9: Rotating the soft edge through the three local sectors.

Theorem 2.24 (soft theorem for A_g). *Let P be generic, q admissible, and ε_0 as in Lemma 2.23. For every root g of P_ε other than the soft edge, with g' the corresponding root of P (the return edge corresponds to E_j), there is $\varepsilon_1 \in (0, \varepsilon_0]$ with*

$$A_g(P_\varepsilon) = \frac{\chi_- + \chi_+}{2} A_{g'}(P) \quad (0 < \varepsilon < \varepsilon_1). \quad (20)$$

Equivalently, in the same-sign, mixed and loop sectors the multipliers are $-\tau_j(P)$, 0 and $\tau_j(P)$, respectively.

[Proof in Part II, page 101.](#)

Lemma 2.25 (the local vertex identity). *Let P be generic, q admissible and $0 < \varepsilon < \varepsilon_0$ as in Lemma 2.23; in the soft insertion, let π be a composition of an interval of the boundary word of (P_ε, g) whose first part is the single soft leaf, $\pi = (\mu_j, \mu_*, a_2, \dots, a_m)$ in the letter's notation $(1, 2, a_2, \dots, n)$. Then, for every such ε ,*

$$V^\pm(\mu_j, \mu_*, a_2, \dots, a_m) = \frac{\chi(\mu_*, \mu_j, a_2) \mp \chi(\mu_*, \mu_j, a_m)}{2} V^\pm(\mu_j, a_2, \dots, a_m),$$

the unlabelled local vertex identity following the soft-insertion formula in the letter.

Proof. By Lemma 2.23(i) P_ε satisfies (G1), so Lemma 2.3 applies to it: the gate at the cut μ_* has near sign $d = \chi(a_2, \mu_*, \mu_j) = \chi(\mu_*, \mu_j, a_2)$ and far sign $h = \chi(a_m, \mu_*, \mu_j) = \chi(\mu_*, \mu_j, a_m)$ (cyclic permutations), and its ordinary and root values are $(d \mp h)/2$. The right side is read in the boundary word of (P, g) , identified with that of (P_ε, g) with μ_* deleted through the identification of the edges of P_ε with those of P in Lemma 2.23(iii). Every other gate of π reads a triple in which μ_* occurs only as the neighbour a_1 of the cut a_2 , and for $0 < \varepsilon < \varepsilon_0$ the determinant of the triple (a_3, a_2, μ_*) is continuous in ε and nonzero by (G1) of P_ε (Lemma 2.23(i)), hence of constant sign on that connected interval, equal to its nonzero limit $\chi(a_3, a_2, \mu_j)$ as $\varepsilon \rightarrow 0$; the far signs of π do not involve μ_* . $\square$

Figure 9 shows the attachment sectors of Definition 2.22; flat and doubling transitions connect them.

2.7 Reversal and shift

Proposition 2.26 (reversal and cyclic shift). *For every polygon satisfying (G1) and every root g :*

(i) $A_g(\sigma P) = A_{g+1}(P).$

(ii) With $\bar{P}_i = \mu_{2-i}$ and $g^\vee = 1 - g$,

$$A_{g^\vee}(\bar{P}) = (-1)^n A_g(P). \quad (21)$$

All vertex and edge indices in this statement are read modulo n .

[Proof in Part II, page 105.](#)

Remark 2.27 (rooted scope). The construction and the local laws in this section retain the specified root. None of their proofs assumes root independence on a fixed polygon.

3 The corner state sum C

Except where an explicit larger domain is stated, P is generic throughout this section. The definitions follow the letter's Section VI and use the positive lifts of the actual carrier diagrams.

3.1 Decompositions and subpolygons

Definition 3.1 (decomposition). A *decomposition* of P is an independent set $S \in \text{Ind}(G_P)$ of crossings (Definition 1.13): no two crossings of S interlace.

Definition 3.2 (oriented smoothing and subpolygons). Let $S \in \text{Ind}(G_P)$. Cutting the traversal circle $\Gamma(P)$ at the two visits of each $x \in S$ and reconnecting the arc arriving at one visit of x to the arc leaving the other visit of x (for both visits) partitions $\Gamma(P)$ into closed cycles; the *subpolygons* (or *carriers*) of S are the closed polygonal curves in $\mathbb{R}^2$ traced by these cycles. Each subpolygon Q is a closed polygon whose corners are the vertices of P it passes through and the crossing points of S it turns at (its *smoothing corners*); at a smoothing corner between E_i and E_j it arrives along one of $\tilde{\lambda}_i, \tilde{\lambda}_j$ and leaves along the other. The *crossings of Q* are the crossings $x \in X(P) \setminus S$ both of whose visits lie on Q ; their number is m_Q . A crossing in $N(S)$ has its two visits on different subpolygons and is a crossing of no subpolygon.

Convention 3.3 (ownership of a selected visit). At a selected crossing with traversal visits a, b , reconnect the incoming arc at a to the outgoing arc at b , and the incoming arc at b to the outgoing arc at a . Assign the original visit a to the resulting cycle containing the incoming one-sided arc at a ; assign b by the analogous rule. Unselected visits retain their ordinary traversal ownership. Thus a selected visit is assigned to one abstract carrier, even though the two carriers through that smoothing site have the same point of the plane in their polygonal images. This is the *incoming-visit convention*. It specifies endpoint ownership only; the geometric reconnection is unchanged.

Remark 3.4 (why the convention exists). Clause (iv) of the lemma that follows quantifies over *all* crossing visits, including the two visits of a selected crossing, and both carriers through a smoothing site pass through the same point of the plane; without an ownership rule for those two visits the clause is not well posed. The convention supplies that rule and nothing else: it changes no support, turn or retained-crossing datum.

Lemma 3.5 (subpolygons and their visit assignment). *Let $P = (\mu_1, \dots, \mu_n)$, $n \geq 3$, be a generic polygon: for every three distinct indices the vertex determinant is nonzero, and no three edge interiors concur. Use one-based cyclic indices, directed edges $E_i = [\mu_i, \mu_{i+1}]$, directions $d_i = \mu_{i+1} - \mu_i$, and turns $\tau_i = \text{sgn det}(d_{i-1}, d_i)$. Let $X(P)$ be its set of transverse crossings of remote edges. The traversal circle has two visits for each crossing, in the order of the oriented traversal. Crossings interlace when their two pairs of visits alternate on that circle. Let $S \subseteq X(P)$ be independent: no two elements of S interlace. Perform the above oriented reconnections, tracing the resulting cycles by inherited straight subsegments, with corners at original vertices and selected smoothing sites. Call these polygonal cycles the carriers, and use the incoming-visit convention. Then:*

- (i) *There are exactly $|S| + 1$ carriers, independent of the order of reconnections. Each carrier traverses the visits assigned to it in their cyclic order inherited from the original traversal circle. The two visits of a selected crossing belong to different carriers.*
- (ii) *Every carrier has nonzero edges, at least three corners, and no antiparallel consecutive directions. All its corner turns are nonzero. At an original vertex i the turn sign is τ_i . At a selected crossing of E_i, E_j , the two smoothing corners have signs $\text{sgn det}(d_i, d_j)$ and $\text{sgn det}(d_j, d_i)$, one left and one right.*

- (iii) A carrier's self-intersections are exactly the unselected crossings both of whose visits are assigned to it. They are transverse, none is a corner, and no carrier has a triple point. An unselected crossing interlacing some element of S has its visits on different carriers; an unselected crossing interlacing no element of S has both visits on one carrier.
- (iv) The assignment of all crossing visits is noncrossing. There are no four distinct visits u_1, u_2, u_3, u_4 in that cyclic order with u_1, u_3 on one carrier and u_2, u_4 on a different carrier. This includes selected visits under the stated convention.

Proof in Part II, page 107.

Definition 3.6 (uniform subpolygons; retained data). A subpolygon Q is *uniform* if all its turns have one sign and *mixed* otherwise. A decomposition S is *uniform* if all its subpolygons are. For a subpolygon Q write $r_Q = \text{rot}(Q)$ (Lemma 1.20, applicable by Lemma 3.5), m_Q for its number of crossings, and ℓ_Q for its number of left turns.

Lemma 3.7 (Carriers on the weakly generic locus). *Let $P \in \mathcal{W}_n$, $n \geq 3$, be weakly generic (Definition 1.15). Form its Gauss word and interlacement graph from its actual transverse crossings and traversal visits, and apply Definitions 3.1, 3.2 and 3.6 with Convention 3.3. These constructions are well defined, and every conclusion (i)–(iv) of Lemma 3.5 holds for every independent set S of crossings of P . In particular every carrier lies in the regular locus, so r_Q , m_Q and the uniform/mixed classification of Definition 3.6 are defined.*

Proof in Part II, page 109.

3.2 The positive lift and the corner coefficient

Definition 3.8 (positive crossing; positive lift). An oriented link diagram in the plane is a finite collection of closed oriented curves with finitely many transverse double points, no triple points, and at each double point a choice of the *over* strand. A double point with over-strand direction u_o and under-strand direction u_u is *positive* if $\det(u_o, u_u) > 0$ and *negative* otherwise. The *positive lift* of a subpolygon Q is the oriented knot diagram whose curve is Q , whose double points are the crossings of Q , and in which at every double point the over strand is chosen so that the crossing is positive. Its writhe (the sum of crossing signs) is m_Q .




Figure 10: A positive-lift portrait of a polygon and the oriented reconnection at a selected crossing. This is the smoothing of Definition 3.2; the positive lift supplies crossing data for the carrier knot.

3.3 From decorated crossing records to knots

We use the plane orientation and crossing-sign convention of Definition 3.8. A diagram here is a finite polygonal immersion, or a regular smooth immersion with finitely many transverse double

points and possible finitely many corners away from crossings. The corners used below admit clean rounding discs; in particular all SM subpolygon diagrams and their regular clean modifications are included. The parametrizing circle is oriented. To retain the outer region in the construction, compactify the oriented plane to an oriented sphere. All topological extensions needed for the bridge are proved below. The named record defined next (crossing occurrences, successor, pairing, over/under bits and signs) is what this section and Section 4 call the *crossing record* of a diagram.

Definition 3.9. Let $\gamma : C \rightarrow S^2$ be an actual generic immersion of an oriented circle in an oriented sphere, with over/under choices at its transverse double points. Its finite crossing-occurrence set M has forward successor s , pairing involution τ without fixed points, an over/under bit at each occurrence, and crossing signs

$$\sigma(c) = \text{sgn det}(u_{c,O}, u_{c,U}). \quad (22)$$

A named record isomorphism is a bijection $\Phi : M \rightarrow M'$ preserving successor, pairing, over/under bits and these signs. The parametrizing circles are oriented; the finite bijection must preserve their cyclic orders. It does not prescribe a map at every unmarked parameter and does not permit traversal reversal. After a finite subdivision we choose an orientation-preserving piecewise-linear circle map $\bar{\Phi} : C \rightarrow C'$ extending Φ : on each interval between successive marked points use the positive affine map in oriented interval coordinates. If M is empty, choose any positive circle parametrization.

Lemma 3.10 (Finite planar models). *The finite polygonal diagrams and regular smooth modifications used here admit finite piecewise-linear shadow models by orientation-preserving ambient isotopies of the page. The crossing names, the four-ray orders, the traversal direction and the over/under designations are retained. No preassigned homeomorphism at unmarked circle parameters is required.*

[Proof in Part II, page 109.](#)

The same finite-disc and strip proof applies to finitely many components: each mixed crossing still has two transverse local branches, and the remaining compact arc collection is finite. All components are retained.

Lemma 3.11 (Constructive polygonal two-disc theorem). *A simple polygonal circle in the oriented sphere has exactly two complementary regions, and both closures are PL discs. Finite prescribed positive PL boundary maps between such discs extend to positive PL disc maps. A continuous positive boundary map also extends to a topological disc map. The statement includes the exterior region.*

[Proof in Part II, page 110.](#)

Lemma 3.12 (PL sphere isotopy). *Every positive PL self-homeomorphism of the sphere is isotopic to the identity through positive topological homeomorphisms. Only a relocation of two sufficiently small triangles is needed in the construction.*

[Proof in Part II, page 112.](#)

Lemma 3.13 (Compatible heights have a fibrewise ambient isotopy). *For one actual generic finite diagram, two continuous height assignments with the same strict over/under order at every double point represent ambient-isotopic oriented links. This holds for any finite number of components and requires no differentiability of the height functions.*

[Proof in Part II, page 113.](#)

Theorem 3.14 (Equivalence from decorated crossing records). *Let D, D' be oriented one-component generic diagrams in the finite polygonal or regular smooth class specified above, in the oriented plane. Suppose a bijection of crossing visits preserves their oriented cyclic order, crossing pairing, over/under designation and crossing sign $\text{sgn det}(u_o, u_u)$. Then the two diagrams present the same oriented knot: their compatible lifts in S^3 are joined by an ambient isotopy preserving the traversal orientation. The assertion includes diagrams with no crossings. No preservation of the unbounded face or of the point at infinity is assumed or asserted.*

Proof in Part II, page 113.

Definition 3.15 (carriers of the flat configurations). Under the hypotheses of Lemma 1.30, identify the crossing visits of the two sides, of the centre $P(0)$ and of the deletion $P(0) \setminus j$ through the common cyclic Gauss word of that lemma's clauses (ii) and (iii), and fix an independent set S in their common interlacement graph. In each of these four configurations mark the traversal circle at every original vertex and every crossing visit, exchange the two outgoing successors at the two visits of every selected crossing — the reconnection of Definition 3.2 with Convention 3.3 — and trace the resulting oriented closed cycles by the inherited straight subsegments. These oriented closed polygonal cycles are the *carriers* of S in that configuration. On the two generic sides they are the carriers of Definition 3.2; at the centre and on the deletion the same words define them, no genericity being assumed.

Corollary 3.16 (flat carrier data). *Under the hypotheses of Lemma 1.30, fix an independent set S in the common interlacement graph and form the carriers of S on both sides, at the centre and on the deletion (Definition 3.15). Then:*

- (i) *The carriers correspond under their named traversal arcs. Exactly one contains μ_j . The central copy of that carrier differs from its deletion copy only by the positive-flat subdivision at μ_j ; every other central carrier is unchanged by deletion. Every carrier has nonzero segments and no antiparallel corner; except for that one central zero turn, all corner turns are nonzero.*
- (ii) *Corresponding carriers have the same retained self-crossing visits, pairing, signs and positive over/under bits. They have the same signed rotation and hence the same absolute rotation. Every corresponding corner away from μ_j has the same turn sign on both sides and in the deletion. The extra corner at μ_j is right on the right side and left on the left side.*
- (iii) *Define a carrier's selector to be 1 if all its turns are right, $(-1)^c$ if all its c turns are left, and 0 if its turns are mixed. For the carrier through μ_j , let W_{right} , W_{left} , and W_{del} be these selectors on the two sides and on the deletion. Then*

$$W_{\text{right}} - W_{\text{left}} = W_{\text{del}}. \quad (23)$$

All other corresponding carrier selectors agree.

These assertions concern geometric and combinatorial carrier data. They make no assignment of a state-sum value at the flat centre.

Proof in Part II, page 115.

Literature input 1 (HOMFLY-PT polynomial). There is a map $D \mapsto H_D(a, z) \in \mathbb{Z}[a^{\pm 1}, z^{\pm 1}]$ on oriented link diagrams that is invariant under the three Reidemeister moves and planar isotopy, takes the value 1 on the crossing-free circle, and satisfies $aH_{D_+} - a^{-1}H_{D_-} = zH_{D_0}$ for every skein triple. Its value depends only on the oriented link presented by D .

3.4 The local source construction and its algebraic consequences

The source fact below is the literal construction used by Lickorish–Millet. Its local conclusions suffice for the algebraic proofs in this subsection. The broader source claim in Literature input 1 remains separately declared; no arbitrary ambient-isotopy theorem is inferred from the local fact. Diagrams are actual nonempty finite oriented diagrams, including all crossing-free components. Algebraic empty products are not empty links.

Literature input 2 (Exact local LM construction projection). For every such diagram D , let $F_D(l, m)$ be the function constructed in Section 1 of Lickorish–Millet [16], not an arbitrary function with similar values. The source constructs it in $\mathbb{Z}[l^{\pm 1}, m^{\pm 1}]$, independently of the auxiliary component order, nonsingular basepoints, and order of required crossing switches. It is defined modulo positive page isotopy, satisfies

$$lF_{D_+} + l^{-1}F_{D_-} + mF_{D_0} = 0, \quad \mu = -(l + l^{-1})m^{-1}, \quad (24)$$

and has the initialization $F_D = \mu^{c-1}$ for every based ordered c -component UNDER-first diagram. UNDER-first means that, traversing components in their order, from their basepoints and in their orientations, the first encounter with each crossing is the underpass. In particular every crossing-free c -component diagram has that value, regardless of nesting or orientations. Here $c \geq 1$.

The same function is invariant under each actual ordinary oriented Reidemeister move. For a supplied finite sequence, take the maximum of the finitely many crossing counts and apply source Proposition 4 at that level one move at a time. No bound by the endpoint crossing counts is required. The imported construction and its Propositions 1–6 stop before the final p120 inference from arbitrary ambient isotopy to a finite move sequence.

Literature input 3 (Uniqueness of the source polynomial over its own ring). Lickorish–Millet prove [16, Theorem, pp. 112–113, with the uniqueness argument on p. 120] that the function F_D of Literature input 2 is the only map from oriented link diagrams (modulo plane isotopy) to $T = \mathbb{Z}[l^{\pm 1}, m^{\pm 1}]$ that depends only on the isotopy class of the oriented link, takes the value 1 on the unknot, and satisfies $lF_{D_+} + l^{-1}F_{D_-} + mF_{D_0} = 0$ on every skein triple; the same statement is Theorem 15.2 of Lickorish [17]. Uniqueness is asserted among all such maps, with no restriction on the support or on the coefficients of a competitor. A map on diagrams that is invariant under planar isotopy and the three Reidemeister moves depends only on the isotopy class of the oriented link by Reidemeister’s theorem [26], the source’s own input for its descent on p. 120; the uniqueness therefore applies to every such invariant map with the source skein and the unknot value. The ring is the source’s: no statement over $\mathbb{Z}[a^{\pm 1}, z^{\pm 1}]$ is imported here (Registry 11.1).

Lemma 3.17 (Gaussian coefficient transport of skein uniqueness). *Put $R = \mathbb{Z}[a^{\pm 1}, z^{\pm 1}]$. Any two maps $D \mapsto Q_D \in R$ on oriented link diagrams that are invariant under planar isotopy and the three Reidemeister moves, take the value 1 on the crossing-free circle, and satisfy $aQ_{D_+} - a^{-1}Q_{D_-} = zQ_{D_0}$ on every skein triple coincide. No restriction on the support or on the coefficients of the maps is imposed. The lemma transports the uniqueness clause of Literature input 3 to the ring R ; it asserts no existence, no ambient-isotopy invariance and no descent theorem.*

Proof in Part II, page 116.

Theorem 3.18 (The local polynomial and its algebraic properties). *Set $R = \mathbb{Z}[a^{\pm 1}, z^{\pm 1}]$. The same source construction has an evaluation $P_D \in R$ with*

$$aP_{D_+} - a^{-1}P_{D_-} = zP_{D_0}, \quad \delta = (a - a^{-1})z^{-1}. \quad (25)$$

Its value on every UNDER-first c -component diagram is δ^{c-1} . It equals H_D of Literature input 1. It is also the unique function on this exact diagram domain satisfying this skein and all these initialization values. This uniqueness statement explicitly includes the initialization hypothesis; it is not a claim proved from only an unknot value and unspecified ambient invariance. For every c -component diagram,

$$P_D \in z^{1-c}\mathbb{Z}[a^{\pm 1}, z^2], \quad P_D \neq 0. \quad (26)$$

In particular $P_\circ = 1$, and knot evaluations are polynomials in z^2 , with no negative z exponents. All local invariances in Literature input 2 are retained.

[Proof in Part II, page 116.](#)

Lemma 3.19 (A diagrammatically split circle). *Let D' be the actual diagram formed from D by adding one simple crossing-free component having no crossings with D . It may surround some components of D ; it need not lie in the unbounded complementary face. Then*

$$P_{D'} = \delta P_D. \quad (27)$$

In particular every crossing-free c -component diagram has value δ^{c-1} .

Proof. Order and base D , and put the added component last with any basepoint. This adds no crossing encounter, so N and b are unchanged. Induct on (N, b) for D . If $b = 0$, the two source initialization values are δ^{c-1} and δ^c , which have the required relationship. Otherwise choose the same first bad crossing in D and D' . Switching or smoothing does not touch the extra component. The switched pair and the smoothed pair are therefore again related by adding exactly that crossing-free component. In either solved recurrence the coefficients are scalar elements of the commutative ring R . Substitute the inductive factor δ for both smaller values of D' , and factor it out; the remaining expression is the solved recurrence for P_D . This proves the formula. Nesting never enters this induction. This lemma does not claim that an arbitrary spatially split presentation already has this diagram form; such a use requires its own certificate or product proof. $\square$

3.5 Scalar equality from named records

We use the same source function F_D from Literature input 2 and its factor $\mu = -(l + l^{-1})m^{-1}$. Its UNDER-first initialization includes every positive component count and every crossing-free circle.

Lemma 3.20 (Polynomial equality from a named decorated record). *Let D, D' be two actual nonempty finite generic oriented link diagrams. Suppose there is a bijection of their components and crossing occurrences which preserves oriented cyclic successor, crossing pairing, over/under bits and crossing signs. The component bijection must also include all components with no crossing occurrences. Then $F_D(l, m) = F_{D'}(l, m)$. Consequently any common Laurent-ring substitution of these two source values, and every coefficient of the substituted values, agrees. This statement does not require or assert realizability of an arbitrary abstract record or an ambient isotopy between the two diagrams.*

[Proof in Part II, page 118.](#)

Lemma 3.21 (Presentations of the same decorated record). *Suppose two actual finite nonempty decorated diagrams have a specified bijection of their oriented parameter circles and crossing occurrences, preserving cyclic successor, pairing, over/under designations and signs. Their evaluations by the same local LM construction agree. This includes positive page-chart changes of an actual*

spherical diagram, clean crossing-free replacements and different compatible height choices, whenever the endpoint page diagrams belong to the finite-page domain. All crossing-free circles must be included in the component bijection. No arbitrary ambient isotopy is a hypothesis or conclusion of this scalar statement.

Proof. Apply Lemma 3.20 and the common substitution $l = ia$, $m = -iz$. For the stated presentation examples, the bijections can be checked directly. A positive regular page chart does not alter a parameter circle or its order; its positive differential multiplies a crossing determinant by a positive number. An orientation preserving finite piecewise-linear chart has the same cyclic ray order, so after its regular finite-page model the signs are the same. The specified crossing decorations retain the over/under bits. Choosing a different point in a face to omit inserts no crossing occurrence. These are comparisons of two actual diagrams, not a realization of an arbitrary word or a claim that a continuous chart is differentiable.

A clean replacement of one crossing-free interval, fixed on its endpoint collars and missing every other strand, inserts no visit. The original circle parameter therefore supplies the same successor between its two collars and identifies every old crossing germ. Each crossing-free circle is explicitly retained. Finally, compatible heights mean precisely that the prescribed over branch remains higher than the under branch in the chosen positive projection frame. Changing those heights changes none of the decorated record, even if the heights themselves are only continuous. No smoothness of such a height change, or polynomial invariance under its ambient extension, is inferred. $\square$

Lemma 3.22 (A single self crossing has scalar value one). *An actual oriented one-circle diagram with just one self crossing has local LM evaluation $P_D = 1$, for either crossing sign. A crossing-free one-circle diagram also has value one.*

Proof. In the one-crossing case choose the nonsingular basepoint in the cyclic interval immediately before its under occurrence. That occurrence is first; there is no other crossing. The diagram is therefore UNDER-first with one component and has source value $\mu^0 = 1$. The common Laurent substitution keeps this value. The zero-crossing case is the same one-component initialization. This does not assume that a selected planar monogon is the bounded face, or discard another component. $\square$

3.6 Marked products and mixed rows

We use the solved recurrences and nonempty smoothing domain of Theorem 3.18, and the named-record and presentation equalities of Lemmas 3.20 and 3.21. The marked operations below do not assert a relative tangle classification.

Marked joining operation and its actual realizations. A marked diagram is a nonempty actual diagram with a specified closed nonsingular oriented interval I on one component; I contains no crossing and is contained in a clean disc. All crossing discs used in a recursion are chosen disjoint from I . A marked join of (A, I_A) and (B, I_B) cuts a smaller open interval inside each mark and joins the oriented ends crosswise, preserving the surviving collars. There are no additional crossings in the joining arcs. Its finite data are precise: concatenate the marked component cycles at their marked gaps, retain every other component, and keep exactly all old crossing pairings, bits and signs. The component number is $c(A) + c(B) - 1$. The two old marks in the separate factors remain distinguished during recursion; a smoothing selects the unique resulting component containing the relevant interval. The operation is not defined as an unspecified ambient isotopy class.

These actual clean joins exist, including at an interior-face mark.

Details in Part II, page 119.

Theorem 3.23 (Multiplicativity for marked joins). *For two nonempty actual marked diagrams and any actual clean marked join $J(A, B)$ just specified,*

$$P_{J(A, B)} = P_A P_B. \quad (28)$$

This includes smoothing-created multi-component factors and an arbitrary crossing-free marked component. No assertion about a marked terminal tangle being ambient-trivial is needed.

Proof in Part II, page 120.

Theorem 3.24 (Stacked products and split unions). *Let D be an actual nonempty diagram whose components are partitioned into $q \geq 1$ nonempty tagged blocks. Assume that at every crossing between different blocks the smaller-index block is under the larger one. Let D_i be the actual restriction retaining all components in block i and all crossings internal to it. Then*

$$P_D = \delta^{q-1} \prod_{i=1}^q P_{D_i}. \quad (29)$$

There is no bound on the number of components in a block, no restriction on crossings inside a block, and no requirement that their page images be separated. A crossing-free disjoint union is the special case with no crossings between blocks. In particular $P_{A \sqcup B} = \delta P_A P_B$ for two nonempty diagrams presented without mixed crossings.

Proof. Order components block by block, choosing any based order within each. All between-block crossings are first encountered under. Let N be the sum of the INTERNAL crossing numbers and b the sum of their bad crossing numbers. These need not be the total crossing number of D .

If $b = 0$, the full diagram is ascending. With $c_i = c(D_i)$ its scalar is $\delta^{\sum_i c_i - 1}$. Each factor has scalar $\delta^{c_i - 1}$, so their product times δ^{q-1} has the same exponent $q - 1 + \sum_i (c_i - 1) = \sum_i c_i - 1$.

Otherwise resolve a bad internal crossing. Switching it lowers b , preserves the block assignment and the between-block hypothesis, and is the same switch in its block restriction. Smoothing lowers N . Every new component inherits the same block as the strands smoothed; no two blocks merge, and no block becomes empty. All between-block crossing visits persist on their original strands, so their under/over relation still has the required block order. The restriction of the smoothed diagram in that block is exactly the oriented smoothing of the old block restriction; the other restrictions do not change. Lexicographic induction and the appropriate solved skein now give (29), since its factors outside the resolved block and the scalar δ^{q-1} are identical in all three terms. This proof supplies the scalar without moving any component in space. $\square$

Lemma 3.25 (Mixed signed crossings in a stack). *For two distinct components of an actual generic oriented plane diagram, the sum of $\text{sgn det}(u_1, u_2)$ over their transverse intersections, in this fixed component order, is zero. Consequently the half-sum of decorated crossing signs is an integer, and it is zero if one component is always over the other.*

Proof in Part II, page 121.

Theorem 3.26 (Lowest mixed coefficients). *Let D be an actual diagram with $c \geq 1$ tagged oriented components, and let D_i be their actual knot restrictions. Define $\ell_{ij} = \frac{1}{2} \sum \sigma(x)$ using ALL mixed crossings between the original components i, j , and put $\Lambda = \sum_{i < j} \ell_{ij}$. Then*

$$[z^{1-c}]P_D = a^{-2\Lambda} (a - a^{-1})^{c-1} \prod_{i=1}^c [z^0]P_{D_i}. \quad (30)$$

For $c = 2$ this is the two-component mixed row; the two tagged restrictions are intrinsic knot diagrams while ℓ_{12} is computed in their original common diagram.

Proof. Set $h(D) = [z^{1-c}]P_D$. At a mixed crossing the smoothed diagram has $c - 1 \geq 1$ components and support at least z^{2-c} by Theorem 3.18. Its multiplication by z therefore has support at least z^{3-c} , so contributes nothing to the $1 - c$ row of the skein. Hence $ah(D_+) = a^{-1}h(D_-)$ and $h(D_-) = a^2h(D_+)$. The switch from positive to negative changes one mixed sign from $+1$ to -1 , and hence changes Λ by -1 . It follows that $a^{2\Lambda}h$ is unchanged by that switch. Reversing the equality proves the same for a negative-to-positive switch.

Switch exactly the mixed crossings necessary to put each smaller-index component UNDER every larger-index component. These are finitely many switches of an actual diagram, and no self-crossing or intrinsic component restriction changes. By Lemma 3.25 every pair of final components has linking number zero, hence final $\Lambda = 0$. Apply Theorem 3.24 with each component as one block. Its value is $\delta^{c-1} \prod_i P_{D_i}$. Knot support is nonnegative and even in z , so the $1 - c$ coefficient of this expression is exactly $(a - a^{-1})^{c-1} \prod_i [z^0]P_{D_i}$: any positive z power in one factor raises that exponent. Restoring the preserved weight $a^{2\Lambda}$ proves the formula. For $c = 1$ there are no switches, $\Lambda = 0$, and the assertion is the identity $[z^0]P_D = [z^0]P_D$. No coefficient or whole polynomial has been assumed nonzero or divided out. $\square$

Lemma 3.27 (Nested connected interlacement blocks). *Let an actual oriented one-circle decorated record have a nonempty crossing set, partitioned into the connected components of its interlacement graph. Suppose that for every component H an actual retained diagram C_H with exactly that restricted named cyclic record is supplied. Then a finite succession of clean marked joins of these actual diagrams realizes the full record, and every actual diagram with that full record has polynomial $\prod_H P_{C_H}$. The joins preserve the sign of every crossing and the writhe is the sum of the writhes of C_H .*

Proof in Part II, page 121.

Definition 3.28 (corner coefficient and the state sum). For a subpolygon Q of a decomposition of the generic polygon P , let H_Q^+ be the HOMFLY-PT polynomial of its positive lift and put

$$d_Q = 1 - m_Q - |r_Q|, \quad c(Q) = [a^{d_Q} z^0] H_Q^+(a, z) \in \mathbb{Z},$$

the coefficient of $a^{d_Q} z^0$ (zero if that monomial is absent). The *corner state sum* of P is

$$C(P) = (-1)^{\ell(P)} \sum_{\substack{S \in \text{Ind}(G_P) \\ S \text{ uniform}}} (-1)^{|S|} \prod_{Q \text{ subpolygon of } S} c(Q).$$

Figure 11 applies Definition 3.28 to the polygon in Figure 10. A smoothing contributes only when its selected crossings are independent and all its carriers are uniform.

Definition 3.29 (the state sum on the weakly generic locus). For a weakly generic polygon $P \in \mathcal{W}_n$, form its decompositions and carriers from its actual crossings as in Lemma 3.7; let H_Q^+ be the HOMFLY-PT polynomial of the positive lift (Definition 3.8) of a carrier Q , and put $d_Q = 1 - m_Q - |r_Q|$ and $c(Q) = [a^{d_Q} z^0] H_Q^+(a, z)$, with r_Q , m_Q and uniformity the data of Definition 3.6, defined by that lemma. The *corner state sum* of P is

$$C(P) = (-1)^{\ell(P)} \sum_{\substack{S \in \text{Ind}(G_P) \\ S \text{ uniform}}} (-1)^{|S|} \prod_{Q \text{ carrier of } S} c(Q),$$



Figure 11: Four smoothing choices for P ; only (a) contributes.

$\ell(P)$ the number of left turns of P . On $\mathcal{U}_n$ this is Definition 3.28 term by term, and the same symbol C is used. No value of an amplitude or of a gate at zero enters this definition.

Lemma 3.30 (local constancy of the weak state sum). *The function C of Definition 3.29 is invariant under the cyclic shift and locally constant on $\mathcal{W}_n$; hence it is constant on every visible chamber. It is the unique locally constant function on $\mathcal{W}_n$ that agrees with Definition 3.28 on $\mathcal{U}_n$.*

[Proof in Part II, page 122.](#)

Lemma 3.31 (selector form of the carrier state sum). *For a carrier L of an independent support S , let $k(L)$ be its number of corners. Put $\text{wt}(L) = 1$ if all its turns are right, $(-1)^{k(L)}$ if all are left, and 0 if its turns are mixed. Set $\text{wind}(S) = \prod_L \text{wt}(L)$. Then*

$$C(P) = \sum_{S \in \text{Ind}(G_P)} \text{wind}(S) \prod_L c(L), \quad (31)$$

where $c(L)$ remains the coefficient of the actual positive carrier lift in Definition 3.28.

Proof. Every original vertex belongs to exactly one carrier and retains its turn. Every selected crossing contributes one left and one right smoothing corner by Lemma 3.5. Thus the total number of left carrier corners is $\ell(P) + |S|$. If S is uniform, precisely its all-left carriers contribute to this count. Their selector product is therefore $(-1)^{\ell(P) + |S|}$, the sign of this support in Definition 3.28. If S is not uniform, a mixed carrier makes its selector product zero; this support is absent from that definition. The coefficient factors are unchanged in both cases. Summing these term identities proves (31). $\square$

Remark 3.32 (well-definedness). The positive lift of a subpolygon is an honest knot diagram on an actual closed plane curve (Lemma 3.5); no realization of an abstract Gauss word is involved, and

Literature input 1 is consumed only on such diagrams. The selector lemma is an identity for this definition. It asserts no equivalence with a residual-piece polynomial construction in a separate source; any use of such an equivalence would require a separate printed factorization proof.

3.7 The polynomial bound for an individual front

Definition 3.33 (Finite oriented fronts). A front F is an actual map of a nonempty finite union of parameter circles to the oriented (x, z) plane, with finitely many transverse double points and ordinary semicubical cusps, no other singularities, and no vertical tangencies on regular arcs. The limiting tangent at every cusp is also nonvertical: in a semicubical parameter u with cusp at $u = 0$, require $x''(0) \neq 0$. Cusps meet no other strand or singularity. At a crossing the branch with smaller dz/dx is over. A downward cusp is traversed from its locally upper arm to its locally lower arm. Write $D(F)$ for the number of downward cusps, $w(F)$ for the sum of the over-first tangent-determinant crossing signs, and $s(F)$ for the number of crossings plus cusps.

In disjoint clean cusp discs replace each cusp by a simple regular arc with the same oriented attachments and no crossing. The resulting ordinary diagram is denoted $S(F)$.

Lemma 3.34 (independence of the rounding). *Let F be a front on the domain of Definition 3.33. Any two ordinary diagrams $S(F)$ obtained by the cusp replacement of that definition have the same full named record: the same component circles, including crossing-free ones, and the same crossing occurrences, cyclic orders, over/under bits and signs. Consequently $P_{S(F)}$ does not depend on the choice of rounding.*

Proof. A cusp replacement inside a clean cusp disc creates no crossing and changes no successor of an old crossing visit along the oriented parameter circle: the replacing arc has the same oriented attachments as the cusp it replaces. Every crossing pairing, sign and over/under bit is determined by germs outside the cusp discs, which are unchanged, and the component correspondence is the identity on the parameter circles, including those without crossings. Thus any two such diagrams have a named decorated-record isomorphism, and Lemma 3.20 identifies their polynomial values. This does not use an ambient identification with a spatial cusp-shaped link. $\square$

The polynomial used here is the same Lickorish–Millet construction in the variables a, z . Its precise inputs, from Literature input 2, Theorem 3.18, and Lemmas 3.19 and 3.20, are:

1. For every actual nonempty finite oriented ordinary diagram E , P_E is a nonzero element of $\mathbb{Z}[a^{\pm 1}, z^{\pm 1}]$.
2. Positive page isotopy and each supplied ordinary oriented Reidemeister-I, -II or -III move preserve P_E .
3. Full oriented named records determine the value, with all component cycles retained.
4. An actual oriented skein triple satisfies $aP_+ - a^{-1}P_- = zP_0$.
5. $P_U = 1$ for a simple crossing-free circle, and $P_{E \sqcup U} = \delta P_E$, where $\delta = (a - a^{-1})/z$, if E is nonempty and U is a simple crossing-free component without mixed crossings. Nesting is allowed.

No knot-type equality, arbitrary ambient isotopy, global Legendrian completeness, mirror or reversal symmetry, braid representation theorem, or contact-neighborhood theorem is an input to this argument.

Definition 3.35 (degrees in a). For a nonzero Laurent polynomial f in a with coefficients in the integral domain $\mathbb{Z}[z^{\pm 1}]$, $\deg_a f = \max \deg_a f$ denotes the largest a -exponent with nonzero coefficient and $\min \deg_a f$ the smallest; both are integers. The same symbols with the subscript z denote the largest and smallest z -exponents of a nonzero element of $\mathbb{Z}[a^{\pm 1}, z^{\pm 1}]$.

With Definition 3.35, set

$$d(F) = \deg_a P_{S(F)}, \quad B(F) = D(F) - w(F) - d(F) - 1. \quad (32)$$

We will prove $B(F) \geq 0$ for each front, not a statement about a maximum over representatives.

3.7.1 Ordinary certificates for the local front moves

Use Rutherford's elementary front words: l_m inserts a left cusp in positions $m, m+1$, r_m joins those positions at a right cusp, and σ_m crosses them. Number cut positions from top to bottom, read words from left to right, and type each operation by its incoming and outgoing strand counts. In the standard crossing the incoming upper strand travels down to the lower output and is over; the incoming lower strand travels up to the upper output and is under. With both arrows pointing right, the over-first tangent representatives $(1, -1), (1, 1)$ have determinant 2, so this crossing is positive.

Every operation below has an empty support rectangle, fixed ordered boundary attachments, and an unchanged common exterior. All exterior crossings, including mixed crossings, and all components are retained. The permitted orientations need not produce a single component.

Lemma 3.36 (Commutations and nonsingular deformation). *Disjoint-gadget commutations and deformations through fronts without a singular event preserve D, w, d , and hence B . Every supplied finite front can be represented by a finite elementary front word.*

Proof. Disjoint gadgets commute by sliding their empty rectangles. The index changes track the same physical positions after insertion or removal of two strands. After cusp rounding this is a positive page isotopy; equivalently, the full named records are the same. The cusp directions and crossing signs are unchanged, so D and w are unchanged as well.

A deformation without a singular event preserves the records, D and w : signs, cusp directions and cyclic attachments cannot change. Lemma 3.20 gives scalar equality, without a global move-generation theorem.

If singularities have equal x coordinates, separate them by small local x translations, constant near the singularities and interpolated on the intervening graph arcs. There are finitely many clean discs and compact regular graph pieces. Their nonzero x derivatives and transverse crossings persist for sufficiently small translations, and constant translations preserve the cusp germs. Reading successive vertical cuts then gives the finite elementary word. No classification of fronts is used. $\square$

Lemma 3.37 (Front type I). *The front type-I moves preserve B .*

Proof. Each word $l_m \sigma_{m-1} r_m$ or $l_m \sigma_{m+1} r_m$ replaces one through-strand. After rounding its two cusps, the unique crossing bounds an empty ordinary monogon. In the first word, the strand crosses over heading right, passes the right cusp and then the left cusp, and returns under heading right. Both crossing arrows are therefore rightward, or both are leftward when the entire strand is reversed. The other word interchanges the over/under order along the through-strand but again has simultaneously directed crossing tangents. Its crossing is also positive. Exactly one of the two

cusps is downward and one is upward. Ordinary Reidemeister I gives, in the direction creating this front curl,

$$\Delta w = 1, \quad \Delta D = 1, \quad \Delta d = 0. \quad (33)$$

Equation (32) now gives $\Delta B = 0$. An arbitrary front zigzag is not a type-I move. $\square$

Lemma 3.38 (Front type II). *The front type-II moves preserve D, w, d , and hence B .*

Proof. The words $l_{m-1}\sigma_m\sigma_{m-1}$ and $l_{m+1}\sigma_m\sigma_{m+1}$ replace l_m . In the first, the through-strand is under at both crossings; in the second, it is over at both. The other crossing branches are the oppositely directed arms of the same cusp. Their signs are opposite for every permitted through-strand and cusp orientation. After rounding the cusp, the arcs bound an empty ordinary bigon with one common over-strand: an oriented Reidemeister-II site. The same upper and lower cusp arms remain after deletion. Thus $\Delta w = \Delta D = \Delta d = 0$. For right-cusp versions, follow these same local strands backwards; no global reversal identity for P is assumed. $\square$

Lemma 3.39 (Front type III). *The front type-III moves preserve D, w, d , and hence B .*

Proof. The words $\sigma_{m+1}\sigma_m\sigma_{m+1}$ and $\sigma_m\sigma_{m+1}\sigma_m$ form an actual ordinary Reidemeister-III configuration. Label the three physical strands by their initial top-to-bottom order. The three over/under choices give one strict height order, not a cyclic order. Each physical pair crosses on both sides with the same over/under bit and transported arrows. An ordinary oriented Reidemeister-III move therefore applies, with unchanged D, w, d and B . In a reflected local template the strict height order reverses, and an ordinary Reidemeister-III move still applies. Checking this local template does not identify an entire diagram with its reflection. $\square$

Lemma 3.40 (Zigzag and crossed-cusp deletion). *Deleting an empty zigzag lowers s by two and cannot increase B ; applying the crossed-cusp shortcut lowers s by one and cannot increase B .*

Proof. A zigzag consists of two consecutive cusps with no crossing in an empty rectangle. Their directions are both downward or both upward: the strand reverses its x direction twice while continuing its vertical passage. After rounding, it is a simple ordinary arc, positively page isotopic relative to its endpoints to the straightened arc. Deletion therefore gives

$$\Delta w = \Delta d = 0, \quad \Delta D \in \{0, -2\}. \quad (34)$$

Thus it cannot increase B ; the two cusps are removed and no crossing is touched, so s decreases by two. This covers either oriented destabilization without a knot-type assertion.

The word procedure also encounters $l_i\sigma_i$ or $\sigma_i r_i$, a cusp whose own arms cross once. Replace it by the uncrossed cusp with the same two oriented boundary attachments. After rounding, ordinary Reidemeister I deletes its empty monogon. The arms are oppositely directed, so the old crossing has sign -1 . Deletion raises w by one and flips the cusp direction because the boundary arms exchange places. If the old cusp was downward, $\Delta D = -1$; otherwise $\Delta D = 1$. Consequently

$$\Delta d = 0, \quad \Delta w = 1, \quad \Delta B \in \{-2, 0\}. \quad (35)$$

One crossing is removed and the cusp count is unchanged, so s decreases by one; this is the compressed dashed-arrow case of Ng's Figure 1 [21]. It is not asserted to be a Legendrian isotopy. The direct ordinary certificate avoids inserting a temporary higher-complexity front merely to recognize its stabilization. $\square$

Lemma 3.41 (Standard front circles). *Deleting a separated standard front circle with nonempty remainder preserves B . A single standard front circle, and any union of such circles, has $B = 0$.*

Proof. A simple crossing-free component with exactly one left and one right cusp has $D = 1$, $w = 0$ and ordinary value one, in either orientation. A standard circle separated by the word procedure has no mixed crossings. Its smoothed component may lie in a bounded face of the remainder; Lemma 3.19 permits this nesting.

For nonempty remainder, $P_{\text{before}} = \delta P_{\text{after}}$. The leading a coefficient of δ is z^{-1} , in degree one. Its product with the leading coefficient of P_{after} is nonzero because the coefficient ring is an integral domain. Therefore

$$d_{\text{before}} = d_{\text{after}} + 1, \quad D_{\text{before}} = D_{\text{after}} + 1. \quad (36)$$

The writhe is unchanged, so $B_{\text{before}} = B_{\text{after}}$. If this is the last component, stop instead: $D = 1$ and $w = d = 0$ give $B = 0$. Repeated deletion shows that a union of standard front circles also has $B = 0$. No empty-link polynomial or degree is used. $\square$

3.7.2 The two oriented cusp-skein branches

Lemma 3.42 (Cusp-skein inequality). *For either principal direction of an oriented cusp-skein interchange, B of the earlier branch is at least the minimum of B of the other principal branch and B of the unique compatible smoothing. The smoothing has one fewer singularity; the principal branches have the same singularity count.*

[Proof in Part II, page 123.](#)

3.7.3 The finite geometric reduction

Literature input 4 (Finite front-word reduction). For an actual finite front word on the domain of Definition 3.33, the procedure of Ng [21, pp. 6–8 and Figure 1] and the elementary-word proof of Rutherford [28, Lemma 3.2, pp. 8–12 of the arXiv version] supplies a finite principal chain using:

1. typed disjoint-gadget commutations and the three local front moves in Lemmas 3.37, 3.38 and 3.39;
2. the cusp-skein crossing interchange (198), or its reflected pattern;
3. zigzag deletion, the crossed-cusp shortcut of Lemma 3.40, or deletion of a separated standard front circle.

Stop at the first strict decrease of s or at the standard-circle base. Before that decrease the selected procedure does not increase s . Secondary progress is measured at the boundaries of complete finite blocks of the word procedure: a completed nonterminal block moves a singularity to the left of the selected rightmost left cusp, and hence decreases the number on its right. This is not a strict decrease at every individual elementary move. Arm-string extension within a block is bounded by the active strand count. Thus every nonbase front has a finite principal chain to a front of smaller s . At each cusp-skein interchange, the compatible smoothing branch has smaller s than the front at that stage.

This is the constructive word procedure, not a quantification over arbitrary Legendrian isotopies. The imported content is the source’s finite descent, not its printed sentence: Ng’s Lemma 1 gives one step, iterated here until s decreases; Rutherford’s two ruling-polynomial terminal branches are replaced here by the geometric shortcuts of Lemma 3.40; and the finiteness of arm-string extension, which Rutherford asserts, is given its reason here.

[Details in Part II, page 124.](#)

3.7.4 The degree induction

Theorem 3.43 (Polynomial bound for a front). *For every front F on the domain of Definition 3.33, with the same polynomial evaluated on its actual ordinary cusp rounding,*

$$w(F) - D(F) \leq -\deg_a P_{S(F)} - 1. \quad (37)$$

Proof in Part II, page 125.

Remark 3.44 (Boundary of the front bound). Theorem 3.43 is only a polynomial inequality for one actual front. Etnyre’s separate local formula identifies $w(F) - D(F)$ with the self-linking number of its chosen positive transverse pushoff. Identifying $P_{S(F_T)}$ with the value on a separately supplied transverse diagram requires a separate endpoint comparison. The contact composition below supplies it using the explicitly retained global HOMFLY source premise, after constructing the actual ordinary rounding and ambient family. That premise is not part of this individual-front argument; it is Literature input 1, consumed as the registry records.

3.8 Contact conventions and a bound for each representative

All Legendrian and transverse knots and isotopies in this subsection are smooth embedded curves and smooth families in standard contact space. The neighbourhood, generic-front, ambient-isotopy and linking/framing constructions below are internal proofs, using smooth calculus, the inverse function theorem and smooth ODE existence, uniqueness, continuation and dependence.

Published proof comparison. The neighbourhood construction uses the relative Moser method of Geiges [12], Theorem 2.32 and Example 2.33, with the stability proof of Theorem 2.27 (statement p. 21, proof pp. 22–25 of the retained arXiv version). His form $dZ + X dY$ pulls back to $dz - y dx$ under $(X, Y, Z) = (-y, x, z)$; this coordinate change has determinant $+1$. His Hamiltonian proof, Theorem 2.22 and Corollary 2.23, pp. 16–17, allows real-valued Hamiltonians; the $\mathbb{R}^+$ in Theorem 2.22’s statement is inconsistent with its proof. Here compact support supplies global flows on $\mathbb{R}^3$.

Geiges’s Lemma 2.45, pp. 38–39, gives an open-interval front perturbation and the local cusp model. Its instruction to perturb one coordinate and integrate the Legendrian equation does not by itself preserve the period of a closed curve. The closed generic-front construction here instead uses ambient contact flows and carries the chosen pushoff. The linking integrand agrees with the Euclidean formula in DeTurck–Gluck [5], p. 3; the crossing-count and uniform-framing arguments are printed here separately. These comparisons retain the distinction between the published inputs and the constructions printed here.

Lemma 3.45 (A transverse neighbourhood and its Legendrian pushoff). *Let $\alpha = dz - y dx$ and let $T : \mathbb{R}/(2\pi\mathbb{Z}) \rightarrow \mathbb{R}^3$ be a smooth embedded oriented circle with $\alpha(T') > 0$. There are $\delta > 0$ and a smooth embedding $H : S^1 \times D_\delta \rightarrow \mathbb{R}^3$ fixing the parametrized core T such that*

$$H^* \alpha = h \alpha_0, \quad \alpha_0 = d\theta + u dv - v du, \quad h > 0. \quad (38)$$

There is an oriented Legendrian knot L in this neighbourhood whose positive transverse pushoff is transversely isotopic to T . An explicit compactly supported ordinary ambient isotopy carries the parametrized L to the parametrized T , preserving their orientations.

Proof in Part II, page 126.

Lemma 3.46 (Compact parameter avoidance). *Let K be a compact subset of a smooth d -dimensional coordinate manifold and $B \subset \mathbb{R}^m$ a closed parameter ball. Suppose that F is smooth on a neighbourhood of $K \times B$, takes values in $\mathbb{R}^q$, and has derivative of rank $q > d$ at every zero in that product. The parameters a for which $F(t, a) = 0$ for some $t \in K$ form a closed set with empty interior. The same assertion holds simultaneously for a finite collection of such maps.*

Proof. The zero set is compact, so its parameter projection is closed. At each zero a nonzero q -minor and the inverse function theorem express the local zero set as a smooth graph of $e = d + m - q < m$ variables. Finitely many such charts restricted to compact cubes cover the zero set; their parameter projections are Lipschitz maps of those cubes. Subdividing an e -cube into $O(n^e)$ pieces of side $O(n^{-1})$ covers its projected image by m -boxes of total volume $O(n^{e-m})$. This tends to zero. The image cannot contain a fixed open m -box, because any finite box cover has total volume at least the covered volume: partition at the finitely many box faces to prove this inequality. A finite union of the chart images has the same estimate. This proves empty interior and the finite-collection assertion without Sard's theorem. $\square$

Lemma 3.47 (Explicit contact motions). *For a smooth compactly supported $H : \mathbb{R}^3 \rightarrow \mathbb{R}$, the vector field*

$$X_H = -H_y \partial_x + (H_x + yH_z) \partial_y + (H - yH_y) \partial_z \quad (39)$$

has a global flow ϕ_H^s of coorientation-preserving contact diffeomorphisms for $\alpha = dz - ydx$: $(\phi_H^s)^* \alpha = c_s^H \alpha$ with a smooth positive function c_s^H , its conformal factor. Finite compositions of their small-time flows depend smoothly on the times and are arbitrarily C^k -close to the identity on compact sets for each fixed finite k .

Proof. Substitution gives $\alpha(X_H) = H$ and $\iota_{X_H} d\alpha = H_z \alpha - dH$; therefore $\mathcal{L}_{X_H} \alpha = H_z \alpha$. Differentiating the pullback and solving the resulting scalar ODE gives

$$(\phi_H^s)^* \alpha = \exp \left(\int_0^s H_z \circ \phi_H^v dv \right) \alpha. \quad (40)$$

This exponential is the conformal factor c_s^H ; it is positive. The vector field is smooth, compactly supported and bounded; trajectories cannot escape to infinity in finite time. ODE continuation gives all finite times, and backwards uniqueness gives the inverse. Smooth ODE dependence proves the last assertion. $\square$

Theorem 3.48 (A generic front preserving the chosen positive pushoff). *Every smooth oriented Legendrian embedding $L : S^1 \rightarrow \mathbb{R}^3$ admits a smooth ambient coorientation-preserving contact isotopy to a Legendrian embedding L_g whose xz front has only finitely many semicubical cusps and transverse double points, with no triple point and no cusp on another branch. At each cusp there is a smooth local coordinate $u = y - y_0$ in which*

$$x = x_0 + Au^2, \quad y = y_0 + u, \quad z = z_0 + Ay_0 u^2 + \frac{2}{3}Au^3, \quad A \neq 0; \quad (41)$$

the coordinate u may increase or decrease along the prescribed orientation. The isotopy carries a chosen thin positive-pushoff annulus and preserves its positive-pushoff transverse isotopy class, as well as the oriented topological knot type.

[Proof in Part II, page 127.](#)

Lemma 3.49 (Linking calculus and uniform transverse framing). *For two disjoint smooth oriented parametrized circles C_1, C_2 in oriented $\mathbb{R}^3$, use the normalized linking pairing*

$$\ell(C_1, C_2) = \frac{1}{4\pi} \int_{S^1 \times S^1} G \cdot (G_u \times G_v) du dv, \quad G(u, v) = \frac{C_2(v) - C_1(u)}{|C_2(v) - C_1(u)|}. \quad (42)$$

It is symmetric and is constant under smooth families of disjoint oriented pairs. Call a direction $\nu \in S^2$ generic for the pair if at every parameter pair (u, v) at which $C_2(v) - C_1(u)$ is parallel to ν the tangents of the two circles projected along ν are linearly independent; the projection along such a ν displays the pair with finitely many transverse mixed crossings. Let ν point toward the observer and orient the projection plane so that its positive basis followed by ν is positive in $\mathbb{R}^3$; at a mixed crossing the strand nearer the observer is over. Then $\ell(C_1, C_2)$ equals one half the sum over the mixed crossings of the overpass-first sign $\text{sgn det}(u_o, u_u)$ of the projected over and under tangents. Thus it has exactly the linking normalization used in the source's front calculations.

More generally, let C_s , $0 \leq s \leq 1$, be a smooth family of embedded oriented circles and let v_s be a smooth vector field along them, everywhere linearly independent of $\partial_u C_s$. One common sufficiently small positive ϵ gives disjoint framed pairs $(C_s, C_s + \epsilon v_s)$ throughout the family. Their pairing is independent of that radius and of s . This includes a homotopy of normal framings on a fixed curve.

If $T_s : S^1 \rightarrow \mathbb{R}^3$, $0 \leq s \leq 1$, is a smooth family of embeddings with $(dz - y dx)(\partial_u T_s) > 0$, there is one $\epsilon_0 > 0$ such that T_s and $T_s + \epsilon \partial_y$ are disjoint positive transverse embeddings for every s and $0 < \epsilon \leq \epsilon_0$. The number

$$sl(T_s) = \ell(T_s, T_s + \epsilon \partial_y) \quad (43)$$

is independent of such ϵ and of s .

[Proof in Part II, page 130.](#)

Remark 3.50 (the self-linking convention). In (43) the global nonzero contact section ∂_y is the section in Etnyre's standard-space self-linking convention (Etnyre [8, Sections 2.6.3–2.6.4, p. 15]; independence of the chosen nowhere-zero section of the contact structure: Geiges [12, Definition 3.2 and the following paragraph, p. 45], whose ∂_X is $-\partial_y$ under the coordinate change $(X, Y, Z) = (-y, x, z)$); no surface-dependent choice or extension is being made. This identifies the number of Lemma 3.49 with the literature's self-linking number; the lemma itself proves only the invariance and the normalization stated there.

3.9 Exact regular rounding of the cusp endpoint

Lemma 3.51 (Ordinary rounding of the exact cusp germs). *Let L be a smooth oriented spatial embedding of a finite nonempty union of parameter circles in $\mathbb{R}^3$, and write $p = (x, z)$ for its projection. Its only failures of regularity are finitely many isolated cusps; all other projected coincidences are finitely many transverse double points. Assume there are no triple points or cusps on another branch, and the two y heights at every double point are distinct. At every cusp assume the exact germ, on a parameter interval with smooth coordinate $u = y - y_0$,*

$$x = x_0 + Au^2, \quad y = y_0 + u, \quad z = z_0 + Ay_0u^2 + \frac{2A}{3}u^3, \quad A \neq 0. \quad (44)$$

The coordinate u may increase or decrease along the prescribed component orientation, which is not changed. The formula is a hypothesis, not an appeal to a classification of singularities.

There is a jointly smooth family L_λ , $0 \leq \lambda \leq 1$, of oriented spatial embeddings with $L_0 = L$, fixed outside disjoint cusp parameter intervals, such that for every $\lambda > 0$ its xz projection is an ordinary

finite regular generic diagram. It cleanly smooths the cusps, creates no crossing, and retains every original crossing with its oriented decorated data. All original parameter circles, component labels and traversal orientations are retained. With no cusps take the constant family. This is an ordinary spatial deformation, not asserted Legendrian or positive transverse, and it carries no self-linking transport statement.

Proof in Part II, page 132.

Corollary 3.52 (Scalar independence of clean cusp smoothing). *For a front satisfying Lemma 3.51, define a permitted smoothed-front diagram by replacing each cusp with one regular embedded oriented arc in a clean cusp neighbourhood, agreeing with the old germs in endpoint collars and creating no crossing. Retain every double point with its original height choice. Every such actual diagram has the same named decorated record as the diagram $D_\epsilon = p(L_1)$ of the lemma, including crossing-free components. Consequently their source values $F_D(l, m)$ and polynomials $P_D(a, z)$ agree.*

Proof. A cusp replacement introduces no crossing visit and changes no successor of an old visit along the oriented parameter circle. All crossing pairings, signs and O/U bits remain unchanged because their germs lie outside the cusp neighbourhoods. The component correspondence is the identity on the original circles, including those without crossings. Thus the two actual diagrams have a named decorated-record isomorphism. Lemma 3.20, with the same source construction of Literature input 2, gives equality of F_D ; the common substitution of Theorem 3.18 preserves it. This is a scalar statement, not a classification of arbitrary relative arc embeddings or an ambient completeness assertion. $\square$

Remark 3.53. The exact cusp formula is supplied in the actual use by Theorem 3.48. The geometric lemma alone neither proves a per-front polynomial inequality nor asserts that an arbitrary smooth transverse knot has a finite generic specified projection. Polynomial transport to a supplied ordinary endpoint uses the distinct compact smooth-family endpoint comparison supplied below. Self-linking uses the separate genuine transverse-pushoff family, never the ordinary rounding above.

3.10 Contact-path polynomial transport and the representative bound

Lemma 3.54 (The supplied contact-path diagram endpoints). *Let C be a finite nonempty disjoint union of oriented parameter circles, and let $L : C \rightarrow \mathbb{R}^3$ be a smooth embedding with nonvanishing parameter derivative. Assume its xz projection F is regular except at finitely many cusps, each with the exact germ*

$$x = x_0 + Au^2, \quad y = y_0 + u, \quad z = z_0 + Ay_0u^2 + \frac{2}{3}Au^3, \quad A \neq 0. \quad (45)$$

Assume the only multiple points of F are finitely many transverse double points, none at a cusp, and the two y values differ at every double point. Suppose a supplied jointly smooth family of oriented spatial embeddings G_t starts at this parametrized L and ends at T , whose specified xz projection D_T is an ordinary finite regular generic diagram. Then every clean ordinary cusp smoothing $S(F)$, with smaller y over at the unchanged crossings, has the same polynomial as D_T :

$$P_{S(F)} = P_{D_T}. \quad (46)$$

No contact condition is imposed on G_t or on the rounding family.

Proof in Part II, page 134.

Definition 3.55 (generic positive transverse front). In $(\mathbb{R}^3, \ker(dz - y dx))$, a *generic positive transverse front* is the oriented knot diagram in the (x, z) plane obtained from a smooth oriented embedded knot T with $z' - yx' > 0$ whose xz projection is an immersion of the parameter circle with finitely many transverse double points and no triple point, the over strand at each double point being the branch of smaller y . It has no cusp. At a vertical tangent, $x' = 0$, the inequality gives $z' > 0$: every vertical tangent of a generic positive transverse front points upward, a consequence and not a hypothesis. The diagram determines its writhe and its over/under counts; the knot T is named separately where it is used.

Literature input 5 (Front and positive-pushoff inputs). Use standard contact space $(\mathbb{R}^3, \ker(dz - y dx))$, with positive transverse orientation $z' - yx' > 0$ (Etnyre [8, Sections 2.1 and 2.4] for the conventions). Etnyre’s front and pushoff statements [8, Section 2.6.2, equations (5) and (7); Section 2.6.4, equation (9); Section 2.9, Lemma 2.22, equation (17)] give, for an oriented Legendrian front with downward and upward cusp counts D, U — every cusp of an oriented front is traversed either downward or upward, so the total cusp count in his tb formula is $D + U$ —

$$r = \frac{D - U}{2}, \quad tb = w - \frac{D + U}{2}, \quad sl(T_+(L)) = tb(L) - r(L). \quad (47)$$

(For the second formula, the Legendrian tb , the printed proof used is Geiges [11, Proposition 3.5.9, p. 117]; Etnyre’s equation (7) reduces to his Remark 2.14, which is stated there without proof.) For a generic positive transverse front (Definition 3.55), self-linking equals its front writhe: Geiges [12, Section 3.1, Lemma 3.3, p. 46] (printed proof) and Geiges [11, Proposition 3.5.32, p. 127]; Etnyre’s equation (9), Section 2.6.4, states it. These are the local front and pushoff source formulas only.

Lemma 3.45 constructs an oriented Legendrian knot near each smooth positive transverse knot T , a chosen positive pushoff transversely isotopic to T , and an explicit ordinary ambient flow to T . Theorem 3.48 supplies a finite generic front with the exact germs (45), by a smooth ambient contact isotopy carrying the chosen pushoff. These are internal constructions. The linking normalization and invariance along smooth positive transverse families are proved in Lemma 3.49. The local diagram proof in Theorem 3.43, using the finite source reduction 4, gives the following bound, restated here as a lemma so that its consumer below cites it.

Lemma 3.56 (the individual-front bound in self-linking form). *Let F be a front on the domain of Definition 3.33, with $c_\downarrow(F) = D(F)$ its number of downward cusps, and let $S(F)$ be an actual clean ordinary cusp smoothing of F as in that definition. Then*

$$sl_{Ng}(F) := w(F) - c_\downarrow(F) \leq -\max \deg_a P_{S(F)}(a, z) - 1. \quad (48)$$

This is not an application of an ambient-invariant hypothesis to the local diagram construction, nor a maximum over front representatives.

Proof. Theorem 3.18 gives $P_{S(F)} \neq 0$, and $P_{S(F)}$ does not depend on the rounding by Lemma 3.34, so $\deg_a P_{S(F)}$ of display (32) is the largest a exponent with nonzero coefficient, $\max \deg_a P_{S(F)}$. Inequality (37) of Theorem 3.43 reads $w(F) - D(F) \leq -\deg_a P_{S(F)} - 1$, which is the display with $c_\downarrow(F) = D(F)$. $\square$

Lemma 3.54 supplies the separate polynomial transport to the particular ordinary transverse diagram when it is used.

Theorem 3.57 (The contact dictionary and finite-diagram representative bound). *In the xz front page the smaller- y branch is over, and its crossing sign is $\text{sgn det}_{xz}(u_O, u_U)$. For a generic positive transverse front (Definition 3.55) one has*

$$sl(T) = \sum_q \text{sgn det}_{xz}(u_O(q), u_U(q)). \quad (49)$$

Let T be an individual smooth positive transverse knot whose specified xz projection D_T is an ordinary finite regular generic diagram. With P_T denoting the original polynomial of this actual diagram, one has

$$sl(T) \leq -\max \deg_a P_T(a, z) - 1. \quad (50)$$

The finite specified-diagram hypothesis is part of this auxiliary bound; no assertion that every smooth knot has a generic specified projection is required. The carrier-floor construction supplies such a diagram.

Proof in Part II, page 135.

3.11 Tangent lifts and the full corner construction

For a local polygon $L = (q_1, \dots, q_c)$ put $\delta_i = q_{i+1} - q_i$. Its principal angle is written ϑ_i , as in Definition 1.18; τ_i elsewhere in this manuscript remains the sign of that angle. The rotation of a direction loop and of a smooth regular curve is Definition 3.58; the lemma that follows it constructs the lifts and proves independence. The notation P_D is the local polynomial, identified with the original H_D by Theorem 3.18. Its global reversal and mirror deductions below take the uniqueness they need from Lemma 3.17 and retain the existence clause of Literature input 1.

Definition 3.58 (direction loops and smooth rotation). A *direction loop* is a continuous map $T: \mathbb{R}/\mathbb{Z} \rightarrow S^1$ into the unit circle of the oriented plane. A *tangent-angle lift* of T at a seam is a continuous $\theta: [0, 1] \rightarrow \mathbb{R}$ with $T(s) = (\cos \theta(s), \sin \theta(s))$ for $0 \leq s \leq 1$; put

$$\text{tw}(T) = \frac{\theta(1) - \theta(0)}{2\pi}.$$

For a closed C^1 regular oriented curve $\gamma: \mathbb{R}/\mathbb{Z} \rightarrow \mathbb{R}^2$, that is, with $\gamma' \neq 0$ everywhere, put

$$T_\gamma = \frac{\gamma'}{|\gamma'|}, \quad \text{rot}(\gamma) = \text{tw}(T_\gamma).$$

For a polygon in the regular locus, rot remains that of Lemma 1.20. For either a polygon or a C^1 regular closed curve put $R(L) = |\text{rot}(L)|$. The lemma that follows constructs the lift and proves that these values do not depend on its choice, on the seam, or on an orientation-preserving regular reparametrisation.

The rounding, rational-curl, affine-insertion and dual-vector constructions follow the geometric proof in the separate treatment [32].

Lemma 3.59 (tangent lifts and principal turns). *Here rot is as in Lemma 1.20 for polygons and Definition 3.58 for direction loops and closed C^1 regular curves.*

- (i) *Every continuous direction loop has a tangent-angle lift. The integer $\text{tw}(T)$ is independent of the lift and seam, is unchanged by an orientation-preserving reparametrisation, and is constant under a continuous homotopy through direction loops. Consequently the rotation of a closed C^1 regular curve is invariant under regular homotopy, and reversing its orientation negates it.*

(ii) If $L \in \mathcal{R}_c$ has principal turns ϑ_i , then

$$2\pi \operatorname{rot}(L) = \sum_i \vartheta_i.$$

Consequently rotation is constant along every path in $\mathcal{R}_c$, is unchanged by a positive flat subdivision, and is negated by orientation reversal. Every regular three-corner polygon has rotation $+1$ or -1 according to the common sign of its three turns.

(iii) Suppose a closed C^1 regular curve is obtained from L by replacing every corner by a regular arc whose compatible tangent-angle lift has increment ϑ_i . Its rotation equals $\operatorname{rot}(L)$.

More generally, let b, c be regular oriented arcs having the same endpoints and the same oriented tangent rays at both endpoints, and suppose replacing b by c in a closed curve gives closed C^1 regular curves F_b, F_c . For compatible tangent lifts with equal initial values, write their increments as Δ_b, Δ_c . Then

$$\operatorname{rot}(F_c) - \operatorname{rot}(F_b) = \frac{\Delta_c - \Delta_b}{2\pi}.$$

If A_t , $0 \leq t \leq 1$, is a supplied continuous path in $\operatorname{GL}^+(2, \mathbb{R})$ with $A_0 = I$ and $A_1 = A$, applying A to both arcs leaves $\Delta_c - \Delta_b$ unchanged.

Proof in Part II, page 136.

Lemma 3.60 (rounding). Here rot is as in Lemma 1.20 for polygons and Definition 3.58 for closed C^1 regular curves. Let L be a closed polygon with corners $q_1, \dots, q_c$, and let D be an oriented diagram whose underlying plane curve is L — L together with an over/under assignment at each of its double points. Assume the principal turns of L all exist and are nonzero; that L has finitely many double points, all transversal, none of them a corner; and that no corner of L lies on an edge of L other than the two incident to it. Then there is a clearance $\varepsilon_0(L) > 0$ such that for every $\varepsilon \in (0, \varepsilon_0(L))$ there is a C^∞ regular closed plane curve L_ε , and a diagram D_ε carried by it, with these properties.

- (a) L_ε coincides with L outside the union of the discs of radius ε about the corners.
- (b) Inside the disc about q_i the unit tangent moves strictly monotonically, in the sense of $\operatorname{sgn} \vartheta_i$, from $\delta_{i-1}/|\delta_{i-1}|$ to $\delta_i/|\delta_i|$, sweeping an arc of length exactly $|\vartheta_i|$ and no more, and attaining each direction of that arc at exactly one parameter. Moreover the unit-tangent map is an immersion on the open junction arc: parametrized by arclength, its angular derivative is nonzero at every interior parameter, while at the two ends it and all its derivatives vanish, the junction meeting the straight edges flat to infinite order.

Both clauses are exported because both are read downstream, and the second does not follow from the first: a strictly monotone angle may still have a vanishing derivative at a parameter, and the direct count below needs the derivative itself to be nonzero, not merely the monotonicity. "Moves monotonically" alone is weaker still — it is satisfied by a junction that pauses on one direction for a whole sub-arc, and then a count of the parameters carrying a prescribed direction is infinite. Both hold because the profile fixed in the proof below is strictly increasing on $(0, 1)$ with $\varphi' > 0$ there, and flat to infinite order at 0 and 1: the tangent angle at arclength σ along the junction of length ℓ is $\theta_u + \vartheta_i \varphi(\sigma/\ell)$, whose derivative is $\vartheta_i \varphi'(\sigma/\ell)/\ell \neq 0$ for $0 < \sigma < \ell$.

- (c) L_ε has the same double points as L , with the same strands, the same over/under assignment and hence the same crossing signs and the same writhe.
- (d) $\text{rot}(L_\varepsilon) = \text{rot}(L)$.
- (e) The disc package is returned. There are pairwise disjoint closed discs $D_1, \dots, D_c$, one about each corner q_i , such that each D_i meets no edge of L that is not incident to q_i , contains no double point of L , meets the two incident edges exactly in the two sub-segments of length ε at q_i , and contains the whole of the modification made at q_i . This clause is a conclusion of the lemma, not a package handed over by a paragraph of its proof; its consumer asks for a disc meeting the diagram in one embedded arc and reads it here.

Proof in Part II, page 137.

Lemma 3.61 (exact negative-curl replacement). *Here rot is as in Lemma 1.20 for polygons and Definition 3.58 for closed C^1 regular curves. Let F be a connected C^∞ immersed circle in the plane — one component, with finitely many transverse double points and no triple points — given with an oriented diagram, and let p be a point of F at which the tangent points in a fixed direction u , isolated among such points, lying in an embedded arc of F that contains no double point and along which the tangent turns strictly positively. Then F may be modified inside a disc Δ meeting the rest of the diagram only in that arc, so that the resulting diagram F'*

- (i) *is again such an oriented diagram, satisfies $P_{F'}(a, z) = P_F(a, z)$, and has the same double points outside Δ , with the same signs;*
- (ii) *has no point of Δ at which the tangent equals u , and exactly one at which it equals $-u$;*
- (iii) *has exactly one double point inside Δ , and it is negative;*
- (iv) *satisfies $\text{rot}(F') = \text{rot}(F) - 1$ and $w(F') = w(F) - 1$.*

Proof in Part II, page 139.

Theorem 3.62 (Geometric lower bound for the corner coefficient).

- (R) *Here rot is as in Lemma 1.20 for polygons and Definition 3.58 for the underlying plane curve of a smooth diagram. Let D be an oriented knot diagram and $-D$ the same diagram with every arrow reversed. Then $P_{-D} = P_D$; consequently an oriented knot and its reverse have the same polynomial. Moreover $-D$ has the same crossing signs and the same writhe as D , and the rotation of its underlying plane curve is the negative of that of D .*
- (A) *Let L and D satisfy the hypotheses of Lemma 3.60. For every $\varepsilon \in (0, \varepsilon_0(L))$, the construction in the proof of that lemma returns one curve and one diagram at those data: the junction inserted at the corner q_i is determined by ε , by the two incident unit directions and by the transition profile, which is fixed once and for all in that proof and is not chosen per corner or per curve; the arc length ℓ is then determined by the endpoint condition, and the rest of the curve is L itself. Write*

$$\text{Round}(L, D, \varepsilon) = (L_\varepsilon, D_\varepsilon)$$

for that pair, the rounding record of (L, D) at ε , and call L_ε the rounded curve and D_ε the rounded diagram.

- (B) Here rot is as in Lemma 1.20, applied to the polygon L . Let L be a closed polygon with nonzero edges and nonzero principal turns, finitely many transverse double points, no triple points, no corner at a double point and no corner on a non-incident edge, carrying a diagram D . Assume that one of the two orientations has either all principal turns positive, or exactly one negative and all others positive. Choose and fix such an orientation, reversing both L and D if necessary. In the remainder of clause (B), L and D denote this oriented pair. Put $R = |\text{rot}(L)|$. Then there are a direction $u \in S^1$ and an $\varepsilon_1 > 0$ such that for every $\varepsilon \in (0, \varepsilon_1)$ the rounded curve L_ε of the record $\text{Round}(L, D, \varepsilon)$ from clause (A) has exactly R points at which its unit tangent equals u and exactly R at which it equals $-u$; at each of them the tangent crosses that direction in the positive sense.
- (C) Here rot is as in Lemma 1.20, applied to the polygon L . Let D be an oriented knot diagram all of whose crossings are positive, with writhe w , whose underlying plane curve is a closed polygon L with all principal turns existing, nonzero, and of magnitude below π , with finitely many double points, all transversal, with no triple points, none of them a corner of L , and no corner of L lying on a non-incident edge; and let R be the absolute value of its Whitney rotation number. The finiteness/transversality and the two corner conditions are what Lemma 3.60 asks of its input. The no-triple condition is preserved by rounding and is needed below when diagram records are compared. These conditions are stated here so that no generic parent polygon is assumed. Assume that, after reversing the orientation if necessary — which by clause (R) changes neither P_D , nor the crossings and their signs, nor w , nor R — either every principal turn is positive, or exactly one is negative and every other is positive. Then

$$\mindeg_a P_D(a, z) \geq 1 - w - R, \quad (51)$$

and the same bound holds for $f_D(a) = [z^0]P_D(a, z)$ whenever $f_D \neq 0$.

Proof in Part II, page 145.

Remark 3.63 (the rounding record). The record $\text{Round}(L, D, \varepsilon)$ of clause (A) exists to be quantified over. An earlier revision of this floor construction, written when the transition profile was a per-use choice, had every consumer say “every curve the lemma produces” — a quantifier over an unnamed family that read differently in each consumer. The profile is now fixed once and for all in the proof of Lemma 3.60, so the record is one curve and one diagram at each ε ; and clause (B) states its hypotheses and produces its clearance itself rather than importing them by “as in Lemma 3.60”. This is commentary on the form of the statement; the theorem’s content is clauses (R)–(C).

Theorem 3.64 (corner theorem; Theorem 2 of the letter). *Let Q be a subpolygon of a decomposition of a generic polygon, and suppose that after possibly reversing its orientation either all turns are left, or exactly one turn is right. Then*

$$\mindeg_a H_Q^+ \geq 1 - m_Q - |r_Q| = d_Q, \quad H_Q^+ \in \mathbb{Z}[a^{\pm 1}, z^2], \text{ so } \mindeg_z H_Q^+ \geq 0.$$

Proof. For a carrier, Lemma 3.5 supplies every geometric hypothesis of clause (C) of Theorem 3.62 except that no corner lies on a non-incident edge, which holds directly: a carrier’s corners are original vertices and selected crossing points; an original vertex lies on no non-incident edge segment of P by (G1), and the only subsegments containing it are its two incident ones; a selected crossing point lies in exactly two edge interiors by (G2) and is an endpoint, not an interior point, of the subsegments into which the crossing sites divide those two edges. Its actual positive lift has exactly m_Q crossings, each positive, so its writhe is m_Q . The rotation in that clause is $R = |r_Q|$. Thus its lower degree bound is $1 - m_Q - |r_Q| = d_Q$ for the polynomial of that actual diagram. Theorem 3.18

identifies this polynomial with H_Q^+ and proves its nonvanishing and its support in $\mathbb{Z}[a^{\pm 1}, z^2]$. Every supported z exponent is therefore nonnegative, which gives the second bound. A zero coefficient at the corner is permitted. $\square$

Corollary 3.65 (the corner is the south-west slot). *For Q as in Theorem 3.64, $(d_Q, 0)$ is the south-west corner of the Newton polygon of H_Q^+ in the sense that no monomial $a^i z^j$ of H_Q^+ has $i < d_Q$ or $j < 0$; $c(Q)$ is the coefficient at that slot, which may be 0.*

Proof. Theorem 3.64 gives both that every nonzero monomial has a -exponent at least d_Q and that all z -exponents are nonnegative. Consequently no exponent lies strictly below or to the left of the stated slot. By Definition 3.28, $c(Q)$ is exactly the coefficient $[a^{d_Q} z^0]H_Q^+$; neither bound requires that this coefficient be nonzero. $\square$

3.12 Actual carrier blocks and elementary corner values

Definition 3.66 (Undominated blocks and their owners). Fix a generic polygon P , its interlacement graph G_P , and an independent support S . For a vertex set T write $N_G(T)$ for the union of its graph neighbourhoods, and put

$$U(S) = V(G_P) \setminus (S \cup N_{G_P}(S)). \quad (52)$$

The connected components of $G_P[U(S)]$ are called its *blocks*. An undominated crossing has both visits on one carrier by Lemma 3.5; that carrier is its owner. Write D_A for the actual positive diagram of a carrier A , and $P_A = P_{D_A}$. These polynomials equal the corresponding H_A^+ by Theorem 3.18.

Lemma 3.67 (Actual block diagrams and the carrier product). *Every block H has one owner and admits an actual positive carrier diagram D_H with exactly its restricted named cyclic record. Its polynomial P_H is independent of the further smoothings used to produce it. For every original carrier A ,*

$$P_A = \prod_{H \text{ owned by } A} P_H, \quad m_A = \sum_{H \text{ owned by } A} |H|. \quad (53)$$

With no owned blocks the actual diagram has value 1 and $m_A = 0$.

[Proof in Part II, page 148.](#)

Lemma 3.68 (An isolated crossing forces a zero corner coefficient). *Let A be a uniform carrier of S , and suppose one of its self-crossing labels c interlaces no other self-crossing of A . Then $c(A) = 0$.*

[Proof in Part II, page 148.](#)

Lemma 3.69 (Exterior-angle count for an embedded polygon). *Every embedded regular polygon has rotation $+1$ or -1 , according to its traversal orientation around the bounded complementary region.*

Proof. Lemma 3.11 supplies a finite straight triangulation of the bounded complementary region, a disc with Euler characteristic one. Its boundary is the polygon with possible straight subdivisions. Write I for the number of interior vertices, B for boundary vertices, E for all edges and F for triangular faces. A simple boundary cycle has B boundary edges. Euler's count is

$$I + B - E + F = 1. \quad (54)$$

Every interior edge belongs to two triangles and every boundary edge to one. Counting triangle-edge incidences gives

$$3F = 2(E - B) + B = 2E - B. \quad (55)$$

Equation (54) first gives $E = I + B + F - 1$. Substituting this into (55) gives $3F = 2I + B + 2F - 2$, and hence $F = 2I + B - 2$.

The sum of all Euclidean triangle angles is πF . At each interior vertex the angles fill a full turn 2π . If α_j is the angle of the bounded region at boundary vertex j , it follows that $\sum_j \alpha_j = \pi F - 2\pi I$. Substituting the preceding face count yields $\sum_j \alpha_j = \pi(B - 2)$. Traverse the boundary with this disc on the left. Its principal turn at j is $\pi - \alpha_j$, including zero at a straight subdivision. The region is an embedded disc, so $0 < \alpha_j < 2\pi$ and these are the principal turns. Their sum is therefore $\pi B - \pi(B - 2) = 2\pi$. Subdivision contributes zero and changes no other turn. Lemma 1.20 gives rotation 1 for the original polygon. The opposite traversal negates every turn and gives -1 . This counts exterior angles directly; it makes no use of Whitney's formula or of a bound depending on the number of polygon vertices. $\square$

Lemma 3.70 (elementary corner values). *(i) If Q is uniform and embedded ($m_Q = 0$) then $|r_Q| = 1$, $d_Q = 0$ and $c(Q) = 1$.*

(ii) If Q is uniform and $\{y\}$ is a crossing of Q interlacing no other crossing of Q , then $c(Q) = 0$.

Proof. (i) Lemma 3.69 gives $|r_Q| = 1$. Thus $d_Q = 1 - 0 - 1 = 0$. The actual positive diagram is crossing-free and has polynomial 1 by Theorem 3.18; its $a^0 z^0$ coefficient is 1. (ii) The carrier Q and its isolated self-crossing satisfy Lemma 3.68, which gives the conclusion directly. The condition is interlacement isolation; no visible planar monogon is assumed or needed. $\square$

4 Laws for the corner state sum

Throughout, C is the function of Definition 3.28 on generic polygons, and the named walls are those of Definition 1.29.

Remark 4.1 (crossing records and knot presentation). Three phrases of this section and of Section 3 have the following fixed meanings. (a) The *crossing record* of a diagram, here always the positive lift of a carrier, is its named record in the sense of Definition 3.9: the crossing visits in their cyclic traversal order, the pairing of the two visits of each crossing, the over/under bit of each visit, and the crossing signs. For a positive lift every crossing sign is $+1$ (Definition 3.8), so the record is the cyclic order, the pairing and the over/under bits; the words *order and crossing records* in the proof of Lemma 1.30 name the edgewise visit lists of the Gauss word (Definition 1.12), the part of this datum without the over/under bits. (b) The over/under bit is a chirotope datum: at a crossing of E_i and E_j the over strand of the positive lift is E_i exactly when $\det(\tilde{\lambda}_i, \tilde{\lambda}_j) > 0$. On a chamber the crossing set is constant (Proposition 1.11) and this determinant is continuous and nonvanishing, the crossing being transverse, so its sign, and with it every over/under bit, is constant on each chamber; the chamber-constancy and silence proofs below print this argument where they use it. (c) Two diagrams *present the same oriented knot* in the sense fixed by the statement of Theorem 3.14: their compatible lifts in S^3 are joined by an ambient isotopy preserving the traversal orientation. The only proof of this document that uses this notion is the identification of the positive star with a torus knot in Section 7, through that theorem. In Literature input 1 the phrase *the oriented link presented by D* is the source's notion, the oriented link type of the diagram, which Reidemeister's theorem (registry L-1, Section 11) carries from the diagram.

Proposition 4.2 (chamber constancy). *The state sum C of Definition 3.28 is constant on every chamber.*

Proof in Part II, page 150.

Proposition 4.3 (silence). *At a simple exterior-extension wall (E) or a simple pure cut (C),*

$$C(P_+) = C(P_-). \quad (56)$$

Proof in Part II, page 150.

Theorem 4.4 (flat law; W2). *At a simple flat wall at j , with $n \geq 4$,*

$$C(P_{\text{right}}) - C(P_{\text{left}}) = C(P(0) \setminus j). \quad (57)$$

Proof in Part II, page 151.

Lemma 4.5 (products and the two-component row). *For oriented knots K, J , $H_{K\#J} = H_K H_J$ and $H_{K \sqcup J} = \frac{a-a^{-1}}{z} H_K H_J$. For a two-component oriented link diagram $D = D_1 \cup D_2$ with linking number lk , $[z^{-1}]H_D = (a - a^{-1})a^{-2\text{lk}}[z^0](H_{D_1}H_{D_2})$.*

Proof. Choose actual knot diagrams for K, J with clean marked intervals. The clean joining construction preceding Theorem 3.23 gives a diagram of their ordinary oriented connected sum. That theorem gives $P_{K\#J} = P_K P_J$ on these actual representatives, and Theorem 3.18 identifies each P with H . For split union, choose representatives with disjoint page images. Theorem 3.24, with two one-component blocks and no mixed crossings, gives $P_{K \sqcup J} = \delta P_K P_J$; the same identification gives the asserted H formula. Passing from these chosen representatives to the knot-class notation uses the full global source premise explicitly retained in Literature input 1; it is not a conclusion of the local construction alone.

For the actual two-component diagram D , apply Theorem 3.26 with $c = 2$. Its linking half-sum is computed from the original common presentation, with the crossing signs of Definition 3.8, exactly the lk in the statement. Its two intrinsic restrictions are the actual diagrams D_1, D_2 . Theorem 3.18 identifies the result with

$$[z^{-1}]H_D = (a - a^{-1})a^{-2\text{lk}}[z^0]H_{D_1}[z^0]H_{D_2}.$$

Both knot polynomials have only nonnegative even z exponents by that same theorem. Thus a product monomial has z exponent zero precisely when each factor does. It follows that the last two factors equal $[z^0](H_{D_1}H_{D_2})$, which is the required expression. No division by either knot polynomial or by a corner coefficient has been used. $\square$

Theorem 4.6 (vertex-edge law; W1). *At a simple vertex-edge wall at $(M; a)$, of bigon or sliding type, with halves λ_1, λ_2 (Definition 1.33) and contact sign $s = \chi_{a,a+1,M}(P_-)$,*

$$C(P_+) - C(P_-) = s C(\lambda_1) C(\lambda_2).$$

Proof in Part II, page 152.

Theorem 4.7 (empty-cusp zero). *At a simple empty cusp, $C(P_{\text{no}}) = 0$.*

Proof in Part II, page 161.

Theorem 4.8 (soft theorem for C). *Let P be generic, j a vertex, q admissible, and P_ε the soft insertion of Definition 2.22, with attachment signs $\chi_\pm$. Then for all sufficiently small $\varepsilon > 0$*

$$C(P_\varepsilon) = \frac{\chi_- + \chi_+}{2} C(P).$$

[Proof in Part II, page 162.](#)

Proposition 4.9 (reversal). $C(\overline{P}) = (-1)^n C(P)$.

Proof. Reversal preserves the crossing set, interlacement, decompositions, subpolygons and their crossings; it reverses every strand direction, so the positive lift is preserved (both u_o, u_u change sign and $\det$ is unchanged), H_Q^+ is unchanged, since the positive lift of the reversed carrier is the reversed diagram of the positive lift of Q , Theorem 3.62(R) gives $P_{-D} = P_D$ for every oriented knot diagram D , and Theorem 3.18 identifies P with H ; m_Q and $|r_Q|$ are unchanged, uniformity is unchanged, and $\ell(\overline{P}) = n - \ell(P)$ (Lemma 1.24(iii)). The prefactor gives the sign $(-1)^n$. $\square$

4.1 The triple wall

Theorem 4.10 (triple-wall invariance (W3)). *At every simple triple wall, $C(P_+) = C(P_-)$.*

Proof. The separate R proof [33] establishes equality of the polygon coefficient X_1 across every simple transversal triple event in the conventions and precise statement of [32, Sections 1 and 10, ax:R]. The bridge [34, Sections 2–3] takes a simple triple wall in the present conventions to such an event: B1–B3 verify the event domain, the exact forced guard bundle, and transversality, while B4 identifies X_1 with C on each side. Applying the R theorem and these two side identifications gives the asserted equality. This uses only the required inclusion of wall classes. $\square$

Lemma 4.11 (the two triple types). *At a simple triple wall at $\{e, f, g\}$ with edge vectors u_e, u_f, u_g , put $\delta_1 = \text{sgn det}(u_e, u_f)$, $\delta_2 = \text{sgn det}(u_e, u_g)$, $\delta_3 = \text{sgn det}(u_f, u_g)$. Exactly one of the following holds.*

- (i) *Transitive: $(\delta_1, \delta_2, \delta_3) \notin \{(+, -, +), (-, +, -)\}$. Then the relation “ E_i passes over E_j ” of the positive lift is a total order on the three strands, and for every subpolygon Q retaining all three crossings the positive lifts on the two sides differ by a Reidemeister III move.*
- (ii) *Cyclic: $(\delta_1, \delta_2, \delta_3) \in \{(+, -, +), (-, +, -)\}$. Then each strand passes over the next cyclically. For a subpolygon retaining all three crossings, the two sides have that same cyclic over/under pattern and opposite local triangle configurations.*

[Proof in Part II, page 163.](#)

Remark 4.12 (Historical terminology). The cyclic triangle change is conventionally called a Δ -move [20]. This terminology and its general knot-type behavior are not used as premises here. The preceding lemma asserts only the explicitly described local crossing pattern in its cyclic case; it supplies no cyclic knot-invariance theorem.

Proposition 4.13 (what is proved about the triple wall). *(i) For a transitive triple wall and a decomposition S that contains none of the three crossings and in which all three lie on one subpolygon, the term of S is the same on both sides.*

- (ii) *Across the wall the interlacement of every pair among the three crossings toggles and no other interlacement changes; the local interlacement graph on the two sides is either K_3 and the empty graph, or a path and its complement (an edge plus an isolated vertex).*

- (iii) Under the hypothesis that no crossing outside the three interlaces any of them (outside isolation; fix an outside decomposition Q), the sum of the terms of all decompositions $Q \cup J$, J a subset of the three crossings, is the same on both sides, for both local graph orbits.

Proof in Part II, page 165.

Remark 4.14 (The general triple cancellation). In the transitive case, decompositions smoothing one or two of the three crossings must be reindexed across the wall: their retained crossings and even their existence can change. In the cyclic case the unsmoothed term can change as well. Theorem 4.10 includes these cancellations; Proposition 4.13 isolates cases whose arguments are useful for understanding the local mechanism.

5 Fixed-count transport

5.1 The circular stars

Definition 5.1 (stars and the bow-tie). For $r \geq 1$ put $N = 2r + 1$, $u_k = (\cos \frac{2\pi k}{N}, \sin \frac{2\pi k}{N})$ for $k \in \mathbb{Z}/N$, and

$$K_r = (\mu_1, \dots, \mu_N), \quad \mu_t = u_{r(t-1)},$$

the polygon visiting every r -th point of the regular N -gon in turn. Put $K_{-r} = \overline{K_r}$. The *bow-tie* is $K_0 = ((0, 0), (2, 2), (0, 2), (2, 0))$.

Lemma 5.2 (the stars are generic). *Let $r \geq 1$ and $N = 2r + 1$. Then $\gcd(r, N) = 1$, and the following hold.*

- (i) $\sigma K_r = R(K_r)$, where R is the rotation of the plane by $2\pi r/N$ about the origin. Consequently every principal turn of K_r equals $2\pi r/N \in (0, \pi)$, all turns are left, and $\text{rot}(K_r) = r$.
- (ii) Every edge line of K_r is tangent to the circle of radius $\rho = \cos(\pi r/N)$ about the origin, at the midpoint of the edge, and the origin lies strictly to the left of every directed edge.
- (iii) K_r is generic; hence K_{-r} is generic with all turns right and $\text{rot}(K_{-r}) = -r$.
- (iv) K_0 is generic, with $(\chi_{123}, \chi_{124}, \chi_{134}, \chi_{234}) = (+1, -1, -1, +1)$, turns $(\tau_1, \tau_2, \tau_3, \tau_4) = (-1, +1, +1, -1)$, $\text{rot}(K_0) = 0$ and $X(K_0) = \{\{1, 3\}\}$.

Proof in Part II, page 174.

5.2 Connectivity of the regular fibres

Lemma 5.3 (positive closing lengths along a direction path). *Let $u_1(t), \dots, u_n(t)$ be continuous unit vectors on a compact interval, lying in no closed semicircle at any parameter. There are continuous strictly positive lengths $l_i(t)$ with $\sum_i l_i(t)u_i(t) = 0$. At either endpoint, any prescribed positive closing lengths can be joined to the selected lengths while keeping the directions fixed.*

Proof in Part II, page 174.

Lemma 5.4 (lifting a semicircle to one real interval). *If a finite real sequence $\theta_0, \dots, \theta_N$ has $|\theta_{i+1} - \theta_i| < \pi$, and all its unit directions lie in a closed semicircle, then all its real entries lie in one interval of length π . Conversely, containment of all entries in such an interval puts their unit directions in a closed semicircle.*

Proof. The inverse image of a fixed closed semicircle under $\theta \mapsto (\cos \theta, \sin \theta)$ is a union of intervals $[\alpha + 2\pi k, \alpha + \pi + 2\pi k]$. Distinct consecutive intervals have a gap of length π . A step of absolute size strictly less than π cannot move from one interval to another, so induction fixes one interval for the whole sequence. The converse follows by projection of that interval. $\square$

Theorem 5.5 (connectivity of the fibres). *Let (n, r) be admissible and let $P, P' \in \mathcal{R}_n$ be labelled tuples with $\text{rot}(P) = \text{rot}(P') = r$. If $(n, r) \neq (4, 0)$, a continuous path in $\mathcal{R}_n \cap \text{rot}^{-1}(r)$ joins P to P' . If $(n, r) = (4, 0)$, there is a $k \in \mathbb{Z}/4$ and such a path from P to $\sigma^k P'$. The four shifts of K_0 lie in the four distinct labelled components, distinguished by their turn words. Thus every fibre of cyclic polygon orbits is path connected.*

[Proof in Part II, page 175.](#)

Collision boundary of this statement. The regular locus forbids zero edges and antiparallel consecutive directions; it does not forbid equality of nonadjacent vertices. Such equality may already occur at the endpoints above, so no collision-free strengthening is asserted for arbitrary regular endpoints.

5.3 Transport through simple walls

Lemma 5.6 (transport). *Let P, Z be generic labelled tuples in the same fibre (n, r) . There are $k \in \mathbb{Z}/n$, with $k = 0$ unless $(n, r) = (4, 0)$, and a piecewise-affine path in $\mathcal{R}_n$ from P to $\sigma^k Z$ that is generic except at finitely many parameters, at each of which it is a simple wall germ of type (F), (V), (T), (E) or (C); at each (F) the deletion is generic, and at each (V) both halves are generic, all with fewer than n vertices.*

Proof. Theorem 5.5 supplies the path from P to $\sigma^k Z$ in $\mathcal{R}_n$. Its rotation is constant by Lemma 1.20(ii). Theorem 1.36 replaces it by a path with only the five listed simple wall types. Lemma 1.34 gives generic deletions and halves. $\square$

5.4 Anchors

Lemma 5.7 (rotation of a soft insertion). *Let P be generic, q admissible at j , and ε small (Lemma 2.23). Then $\text{rot}(P_\varepsilon) = \text{rot}(P)$ in the same-sign and mixed sectors, and $\text{rot}(P_\varepsilon) = \text{rot}(P) - \tau_j(P)$ in the loop sector. Moreover the same-sign sector is the open cone of directions strictly between $\tilde{\lambda}_{j-1}$ and $\tilde{\lambda}_j$ (the wedge of the turn), the loop sector is the open cone strictly between $-\tilde{\lambda}_{j-1}$ and $-\tilde{\lambda}_j$, and the mixed sector is the complement of the closures of these two cones.*

[Proof in Part II, page 177.](#)

Definition 5.8 (anchors). Let (n, r) be admissible with $n \geq 4$.

- (Z) If $(n-1, r)$ is admissible, a *zero anchor* is a generic n -gon of the form $P_\varepsilon^{(j,q)}$ (Definition 2.22) in the mixed sector, for some generic $(n-1)$ -gon P of rotation r , some vertex j of P , some admissible q and some small $\varepsilon > 0$.
- (L) If (n, r) is minimal with $|r| \geq 2$, a *loop anchor* is a generic n -gon of the form $P_\varepsilon^{(j,q)}$ in the loop sector, for some generic $(n-1)$ -gon P of rotation $r - \text{sgn}(r)$ and some vertex j of P with $\tau_j(P) = -\text{sgn}(r)$.

(L₀) If $(n, r) = (4, 0)$, a *loop anchor* is a generic 4-gon of the form $P_\varepsilon^{(j, q)}$ in the loop sector for a triangle P and any of its vertices j .

In all three cases q is admissible and $0 < \varepsilon < \varepsilon_0$, where $\varepsilon_0 = \varepsilon_0(P, j, q) > 0$ is the bound of Lemma 2.23: on $(0, \varepsilon_0)$ the insertion is generic and lies in one chamber, so that Lemma 5.7 gives its rotation throughout. This geometric bound is part of the anchor data; it is not a bound uniform over the functions to which a soft theorem is applied. In each case P is the *parent* of the anchor, j its *insertion vertex*, and the soft edge of the anchor is its edge E_j (from μ_j to $\mu_* = \mu_{j+1}$ in the labelling of the anchor).

Proposition 5.9 (anchors exist). *For every admissible (n, r) with $n \geq 4$, exactly one of the cases (Z), (L), (L₀) of Definition 5.8 applies. An anchor of that type exists, is generic, has n vertices and rotation r . For a given generic parent P and vertex j , the mixed sector and the loop sector both contain admissible vectors q , so in case (Z) the insertion vertex may be prescribed.*

Proof in Part II, page 177.

Proposition 5.10 (anchor values). *Let F be an integer-valued function on generic polygons of all arities, constant on chambers and satisfying the soft theorem in the form of Theorem 4.8 at every admissible soft insertion into a generic polygon. With P the parent and j the insertion vertex,*

$$F(Z) = 0 \quad \text{for every zero anchor,} \quad F(Y) = \tau_j(P)F(P) \quad \text{for every loop anchor.}$$

In case (L), $\tau_j(P) = -\text{sgn}(r)$; in case (L₀), it is the orientation sign of the parent triangle. The function C satisfies these hypotheses. The same identities hold for A_g at every anchor root other than the soft edge, with the corresponding parent root used on the right-hand side.

Proof. Fix an anchor with parameter ε and geometric bound ε_0 from Definition 5.8. For its fixed parent, vertex and admissible vector, the soft hypothesis for F supplies a bound $\delta_F > 0$. Choose $0 < \varepsilon' < \min\{\varepsilon, \delta_F\}$. The whole parameter interval between ε' and ε is contained in $(0, \varepsilon_0)$, so both tuples lie in the same chamber by Lemma 2.23. Chamber constancy equates their F values. At ε' the soft theorem applies. Its multiplier is zero in the mixed sector, since the attachment signs are opposite. In the loop sector both attachment signs equal $\tau_j(P)$, so the multiplier is $\tau_j(P)$. This proves both identities at the original ε , without requiring a bound uniform over functions.

Proposition 4.2 and Theorem 4.8 supply the two hypotheses for C . For A_g , fix the physical nonsoft edge throughout the labelled family and its corresponding edge of the parent. Theorem 2.24 gives the required identity on an initial interval for that root. Throughout the geometric interval the tuple remains generic, so every initially nonzero determinant keeps its sign along the connected parameter interval. Proposition 2.7 therefore makes A_g constant along that family. The same smaller-parameter argument proves the rooted assertion, without using root independence. $\square$

6 Root independence, uniqueness, and the comparison theorem

6.1 Root independence of the tree coefficient

Lemma 6.1 (triangle and bow-tie values). *(i) For every labelled triangle satisfying (G1), all turns have a common sign τ , and $A_g = -\tau$ at every physical root.*

(ii) For the bow-tie of Definition 5.1, $A_g(\sigma^k K_0) = -1$ for every physical root g and every integer k .

Proof. For a noncollinear triangle the three turn triples are cyclic permutations of one ordered triple. Lemma 1.4(i) therefore gives the common sign τ . For any physical root, the boundary word has two leaves. Its unary root contribution is $b_{[0,2]} = 0$, since its only ordinary composition has two parts and vanishes by Lemma 2.3. The binary root has $d = h = \chi(a_2, a_1, a_0) = -\tau$, so its factor $(d + h)/2$ is $-\tau$. This proves (i) directly on (G1).

For (ii), use the four coordinates of Definition 5.1. All two-leaf open sums vanish, because their near and far signs agree. On a boundary word (a_0, a_1, a_2, a_3) , put $d_1 = \chi(a_2, a_1, a_0)$, $d_2 = \chi(a_3, a_2, a_1)$, $h_1 = \chi(a_3, a_1, a_0)$ and $h_2 = \chi(a_3, a_2, a_0)$. Every two-part composition has a two-leaf child and vanishes. The unary root term is therefore minus the ordinary three-part gate product, and the three-part root term is its barred counterpart. Hence

$$A_g = \frac{(d_1 + h_1)(d_2 + h_2) - (d_1 - h_1)(d_2 - h_2)}{4}. \quad (58)$$

Expanding the two products cancels $d_1 d_2$ and $h_1 h_2$ and leaves $2d_1 h_2 + 2h_1 d_2$. Thus $A_g = (d_1 h_2 + h_1 d_2)/2$. At the four physical roots of K_0 , direct determinants of its displayed coordinates give

g	d_1	d_2	h_1	h_2	A_g
1	-1	1	-1	1	-1
2	1	1	-1	-1	-1
3	1	-1	1	-1	-1
4	-1	-1	1	1	-1

(59)

Cyclic shifts merely permute these four root words by Proposition 2.26(i). Consequently every root of every $\sigma^k K_0$ has value -1 . $\square$

Theorem 6.2 (root independence). *For every generic polygon P and any two of its edges g, h , $A_g(P) = A_h(P)$.*

Proof in Part II, page 179.

Corollary 6.3 (A is a chamber function on polygons). *$A(P) := A_g(P)$ (any g) is well defined on generic polygons and satisfies: chamber constancy and silence (Proposition 2.7, Theorem 2.15(ii)); the flat law $A(P_{\text{right}}) - A(P_{\text{left}}) = A(P(0) \setminus j)$; the cusp law $A(P_{\text{loop}}) - A(P_{\text{no}}) = -\kappa A(P(0) \setminus j)$ when the deletion satisfies (G1); the vertex-edge law $A(P_+) - A(P_-) = sA(\lambda_1)A(\lambda_2)$; the triple law $A(P_+) = A(P_-)$; the soft theorem $A(P_\epsilon) = \frac{\chi_- + \chi_+}{2} A(P)$ in every sector; reversal $A(\overline{P}) = (-1)^n A(P)$; and $A(K_1) = -1$, $A(K_{-1}) = +1$ on the counterclockwise and clockwise triangles.*

Proof. Theorem 6.2 removes the root choice, and Proposition 2.26(i) gives cyclic descent. Theorems 2.12, 2.13, 2.14, 2.15, 2.24 (evaluated at any root other than the soft edge), Proposition 2.26(ii), and the triangle values in Lemma 6.1(i) give the listed laws and base values. $\square$

Corollary 6.4 (root independence and cyclic descent of the continuation). *For $n \geq 3$ and any two roots g, h , the continuations of Theorem 2.16 satisfy $\tilde{A}_g = \tilde{A}_h$ on $\mathcal{W}_n$. For every root g and every labelled weakly generic tuple P ,*

$$\tilde{A}_g(\sigma P) = \tilde{A}_{g+1}(P), \quad (60)$$

with root indices read modulo n ; consequently $\tilde{A}_g(\sigma P) = \tilde{A}_g(P)$, and each $\tilde{A}_g$ is a function on weakly generic polygons.

Proof. By Theorem 6.2, $A_g = A_h$ on $\mathcal{U}_n$. By Theorem 2.16, $\tilde{A}_h$ is constant on visible chambers and agrees with $A_h = A_g$ on $\mathcal{U}_n$; the uniqueness clause of that theorem for the root g gives $\tilde{A}_h = \tilde{A}_g$.

The cyclic shift σ permutes the coordinates of a labelled tuple, and every condition of Definition 1.15 is a condition on the polygon: nonzero edge vectors and turns, no vertex on a nonincident closed edge, remote edges meeting at most once and transversally, and (G2) are all invariant under relabelling. Hence σ and σ^{-1} map $\mathcal{W}_n$ onto itself, are continuous, and permute its connected components, the visible chambers; likewise for $\mathcal{U}_n$. The function $P \mapsto \tilde{A}_g(\sigma P)$ is therefore constant on every visible chamber, and at a generic tuple its value is $A_g(\sigma P) = A_{g+1}(P)$ by Proposition 2.26(i). The uniqueness clause of Theorem 2.16 for the root $g+1$ gives $\tilde{A}_g \circ \sigma = \tilde{A}_{g+1}$, which is (60). With the first assertion, $\tilde{A}_g(\sigma P) = \tilde{A}_{g+1}(P) = \tilde{A}_g(P)$; iterating, $\tilde{A}_g$ is invariant under every power of σ and descends to the orbits of Definition 1.1. $\square$

Definition 6.5 (the amplitude on the weakly generic locus). For a weakly generic polygon $P \in \mathcal{W}_n$ put

$$A^{\text{cont}}(P) = \tilde{A}_g(P) \quad (\text{any root } g), \quad (61)$$

which is well defined on polygons by Corollary 6.4. On $\mathcal{U}_n$ it is the function $A(P) = A_g(P)$ of Corollary 6.3, by the agreement clause of Theorem 2.16. The letter's amplitude on its generic locus, which is $\mathcal{W}_n$ (Definition 1.15), is read as A^{cont} . On the silent locus $\mathcal{S}_n$ this is a definition, not an identity with the tree sum of Definition 2.4, which assigns no value to a gate at zero (Remark 2.17).

Corollary 6.6 (finite evaluation of the continuation). For every $P \in \mathcal{W}_n$ and every root g ,

$$A^{\text{cont}}(P) = A_g(\chi(P)), \quad (62)$$

the polynomial form of Corollary 2.21 evaluated at the chirotope of P , zero entries included. Moreover A^{cont} is the unique function on $\mathcal{W}_n$ that is constant on visible chambers and agrees with A on $\mathcal{U}_n$.

Proof. By Definition 6.5, $A^{\text{cont}}(P) = \tilde{A}_g(P)$, and Corollary 2.21 identifies this with $A_g(\chi(P))$ on all of $\mathcal{W}_n$. For the uniqueness, A^{cont} is constant on visible chambers and agrees with $A = A_g$ on $\mathcal{U}_n$ (Theorem 2.16, Corollary 6.3); any function with these two properties equals $\tilde{A}_g = A^{\text{cont}}$ by the uniqueness clause of Theorem 2.16. $\square$

6.2 Uniqueness

Theorem 6.7 (uniqueness on generic polygons). Let F be a function assigning an integer to every generic polygon of every arity $n \geq 3$ and satisfying:

- (a) F is constant on chambers and unchanged across simple silent walls (types (E) and (C));
- (b) the flat law $F(P_{\text{right}}) - F(P_{\text{left}}) = F(P(0) \setminus j)$ at every simple flat wall;
- (c) the vertex-edge law $F(P_+) - F(P_-) = s F(\lambda_1) F(\lambda_2)$ at every simple vertex-edge wall, both branches;
- (d) $F(P_+) = F(P_-)$ at every simple triple wall;
- (e) the soft theorem $F(P_\varepsilon) = \frac{\chi_- + \chi_+}{2} F(P)$ for all small ε , at every soft insertion of an admissible vector into a generic polygon;

(f) $F(K_1) = -1$ and $F(K_{-1}) = +1$.

Then $F(P) = A(P)$ for every generic polygon P .

Proof. A satisfies (a)–(f) by Corollary 6.3. Put $\Delta = F - A$, a function on generic polygons, and argue by strong induction on n .

Base $n = 3$. The generic triangles form two chambers, the counterclockwise and the clockwise triangles, containing K_1 and K_{-1} respectively; by (a) and (f), $\Delta = 0$ on both.

Step. Let $n \geq 4$ and assume $\Delta = 0$ on every generic polygon with fewer than n vertices. Let P be generic with $\text{rot}(P) = r$, so that (n, r) is admissible, and let Z be an anchor for (n, r) (Definition 5.8, Proposition 5.9). Lemma 5.6 gives a path from P to $\sigma^k Z$, which as a path of polygons ends at Z , with finitely many simple walls of types (F), (V), (T), (E), (C). Between walls Δ is constant by (a); at (E), (C) it is unchanged by (a); at (T) by (d). At a flat wall at j the jump of Δ is $\Delta(P(0) \setminus j) = 0$ by induction, the deletion being generic with $n - 1$ vertices. At a vertex–edge wall the jump is $s(F(\lambda_1)F(\lambda_2) - A(\lambda_1)A(\lambda_2)) = s(\Delta(\lambda_1)F(\lambda_2) + A(\lambda_1)\Delta(\lambda_2)) = 0$ by induction, the halves being generic with fewer than n vertices. Hence $\Delta(P) = \Delta(Z)$. At the anchor, Proposition 5.10 applies to F and to A (hypotheses (a) and (e), Corollary 6.3): for a zero anchor both values are 0; for a loop anchor with parent P' (a generic polygon with $n - 1$ vertices) and insertion vertex j , $F(Z) = \tau_j(P')F(P')$ and $A(Z) = \tau_j(P')A(P')$, and $\Delta(P') = 0$ by induction. So $\Delta(Z) = 0$ and $\Delta(P) = 0$. $\square$

Remark 6.8 (what the theorem does not assume). No invariance under reversal is assumed. Cyclic relabelling is already built into F being a function on polygons, and the two triangle values are prescribed separately. The cusp law is not assumed either; transport paths avoid cusps (Theorem 1.36). Anchor values follow from the soft theorem together with chamber constancy on the initial soft family. The prescribed values at the triangle stars K_1 and K_{-1} are the base data; no value at a higher star or at the bow-tie is assumed. Those remaining values in Section 7 are consequences.

Corollary 6.9 (uniqueness on the weakly generic locus). *For every $n \geq 3$ let F assign an integer to every weakly generic polygon in $\mathcal{W}_n$ (Definition 1.15). Suppose that F is constant on every visible chamber and that its restrictions to the generic loci $\mathcal{U}_n$, for all arities, satisfy (a)–(f) of Theorem 6.7. Then $F(P) = \tilde{A}_g(P)$ for every $P \in \mathcal{W}_n$ and every root g ; in particular the continuation $\tilde{A}_g$ of Theorem 2.16 does not depend on g .*

Proof. By Theorem 6.7, $F = A$ on every $\mathcal{U}_n$, and by Theorem 6.2, $A = A_g$ there for every root g . Fix a root g . By Theorem 2.16, $\tilde{A}_g$ is the unique function on $\mathcal{W}_n$ that is constant on visible chambers and agrees with A_g on $\mathcal{U}_n$. The function F is constant on visible chambers by hypothesis and agrees with A_g on $\mathcal{U}_n$ by the first sentence; hence $F = \tilde{A}_g$ on $\mathcal{W}_n$. The left side does not involve g , so $\tilde{A}_g = \tilde{A}_h$ for any two roots g, h . $\square$

Remark 6.10 (Relation to Theorem 1 of the letter). The letter’s Theorem 1 concludes on its generic polygons, which are the weakly generic polygons of Definition 1.15, and its chamber-constancy hypothesis is constancy on the visible chambers. Taking any root g in Corollary 6.9 and Definition 6.5 gives $F = A^{\text{cont}}$ on $\mathcal{W}_n$: the corollary is the form in which Theorem 6.7 covers the letter’s statement, its amplitude at a silent configuration being read as the continuation, which is how this document defines it (Remark 2.17); no gate is evaluated at zero. Constancy on the chambers of $\mathcal{U}_n$ alone would leave the values of F on the silent locus unspecified. The common value is also the polynomial value $A_g(\chi(P))$ by Corollary 6.6, which is not used in the proof.

6.3 The comparison theorem

Theorem 6.11 (comparison). $C(P) = A(P)$ for every generic polygon P .

Proof. C satisfies (a)–(f) of Theorem 6.7: (a) Propositions 4.2, 4.3; (b) Theorem 4.4; (c) Theorem 4.6; (d) Theorem 4.10; (e) Theorem 4.8; (f) Lemma 3.70(i) gives corner coefficient 1 for either oriented triangle. Its only decomposition is empty, so Definition 3.28 gives -1 when $\ell = 3$ and $+1$ when $\ell = 0$, as required. $\square$

Corollary 6.12 (laws inherited by C ; including W4). C satisfies every identity of Corollary 6.3 on the domains stated there, in particular the cusp law $C(P_{\text{loop}}) - C(P_{\text{no}}) = -\kappa C(P(0) \setminus j)$.

Proof in Part II, page 179.

Corollary 6.13 (comparison on the weakly generic locus). $C(P) = A^{\text{cont}}(P)$ for every weakly generic polygon $P \in \mathcal{W}_n$, with C the state sum of Definition 3.29 and A^{cont} the continuation of Definition 6.5.

Proof. On $\mathcal{U}_n$, Theorem 6.11 gives $C = A$, and $A = A_g$ for every root g by Theorem 6.2. By Lemma 3.30, C is constant on every visible chamber of $\mathcal{W}_n$. Fix a root g . The uniqueness clause of Theorem 2.16 identifies every function on $\mathcal{W}_n$ that is constant on visible chambers and agrees with A_g on $\mathcal{U}_n$ with $\tilde{A}_g = A^{\text{cont}}$; hence $C = A^{\text{cont}}$. No identification of a value at a silent configuration with the tree sum of Definition 2.4 is used. $\square$

Remark 6.14 (Dependencies and domains). The rooted tree laws, root independence, uniqueness and continuation do not use the triple-wall theorem for the corner sum. The comparison and the laws inherited by the corner sum use Theorem 4.10, whose proof is cited separately. The comparison on the weakly generic locus then follows from the continuation argument. The literature premises are stated in Section 11; the distinction between printed proofs and machine-checked statements is described there as well.

7 Values at the circular stars and on convex polygons

The stars K_r of Definition 5.1 are the minimal-arity vertices of the fibres (Definition 1.25). Their values are computed here on both sides. Their agreement follows from Theorem 6.11 and also checks the independent calculations from Definitions 2.4 and 3.28. The computations in this section are not premises of Theorem 6.7. The triangle and bow-tie values needed earlier are supplied independently by Lemma 6.1; its bow-tie clause is used in the root-independence proof at (4, 0).

Figure 3 in the introduction displays the triangle, bow-tie and first circular-star evaluations.

7.1 The tree coefficient of a star

Proposition 7.1 (tree coefficient of K_r). For $r \geq 1$, $N = 2r + 1$ and every root g ,

$$A_g(K_r) = (-1)^r \text{Cat}_r, \quad \text{Cat}_r = \frac{1}{r+1} \binom{2r}{r};$$

for $r \leq -1$, $A_g(K_r) = -(-1)^{|r|} \text{Cat}_{|r|}$; and $A_g(\sigma^k K_0) = -1$ for all g, k .

Proof in Part II, page 181.

Remark 7.2 (Relation to the Berends–Giele recursion). The Berends–Giele recursion and Catalan generating function in the letter correspond to the generating-function step in the proof of Proposition 7.1: the odd-length open sums $\bar{A}_{2a+1} = b_{2a+1}$ satisfy $G = 1 - tG^3/(1 - tG^2)$ in the variable $t = x^2$ after removing one factor x , whose solution is $\sum_a (-1)^a \text{Cat}_a t^a$; the root gate of the surviving binary root is $V^- = -1$.

7.2 The state sum of a star

Proposition 7.3 (decompositions of a star). *For $r \in \mathbb{Z} \setminus \{0\}$ put $N = 2|r| + 1$ and $m = |X(K_r)|$. Every nonempty decomposition of K_r has a mixed subpolygon, so only $S = \emptyset$ contributes to $C(K_r)$; the single subpolygon K_r is uniform with $|r_{K_r}| = |r|$ and $m = N(|r| - 1)$.*

Proof in Part II, page 182.

Lemma 7.4 (The star’s decorated record and polynomial). *For $r \geq 2$, let D_r be the ordinary diagram of the positive lift of K_r , with corners rounded away from crossings, and let B_r be the standard positive closure of $(\sigma_1 \cdots \sigma_{r-1})^{2r+1}$. There is a bijection of their component circles and crossing visits preserving oriented cyclic successor, crossing pairing, over/under bits and crossing signs. Consequently their local polynomials satisfy*

$$P_{D_r}(a, z) = P_{B_r}(a, z). \quad (63)$$

The rounding creates and removes no crossing visit, so the positive lift of K_r itself, the polygonal diagram of Definition 3.8, has the same decorated record as D_r , and the bijection and (63) hold for it as well.

Proof in Part II, page 183.

Proposition 7.5 (the star is a torus knot). *For $r \geq 2$ the positive lift of K_r presents the torus knot $T(r, 2r + 1)$, the closure of the braid $(\sigma_1 \cdots \sigma_{r-1})^{2r+1}$.*

Proof. Lemma 7.4, by its polygon-level clause, supplies the complete decorated record bijection between the actual positive star (the polygonal positive lift of K_r) and the standard closure. Both are knots. Theorem 3.14 therefore identifies their oriented knot types in S^3 . This proves the stated presentation; the scalar equality in (63) was established separately from the records. $\square$

7.3 Finite computation for torus knots

The following calculation prints the finite algebra behind Jones’s positive-torus formula [15, Theorem 9.7]. The specific Wenzl tableau input is stated explicitly below.

Lemma 7.6 (The local polynomial in Jones’s variables). *For each actual nonempty oriented decorated diagram define $P_J(D; t, x) = P_D(t^{-1}, x)$, with P_D the same construction of Theorem 3.18. We also write $P_{J,D}(t, x)$ for this value. It satisfies*

$$t^{-1}P_J(D_+) - tP_J(D_-) = xP_J(D_0), \quad P_J(\bigcirc) = 1. \quad (64)$$

On every ordered based UNDER-first diagram with $c \geq 1$ components its value is $((t^{-1} - t)/x)^{c-1}$. The following are retained under the substitution: invariance under planar isotopy and the three Reidemeister moves and the value 1 on the crossing-free circle (Theorem 3.18), the split-union rule $P_J(E \sqcup U) = ((t^{-1} - t)/x)P_J(E)$ for a simple crossing-free component U (Lemma 3.19), the displayed skein, and the full named-record equality of Lemma 3.20.

Proof. Apply the Laurent-ring substitution $a \mapsto t^{-1}$, $z \mapsto x$ to the skein and every initialization of Theorem 3.18. The units remain units, so this is a ring homomorphism. Applying it to each local equality — the invariances and initializations of Theorem 3.18 and the split-union rule of Lemma 3.19 — and to Lemma 3.20 retains those equalities as well. Thus this is a renaming of the same source construction, with all component initializations specified, not uniqueness inferred from an unknot value alone. $\square$

7.3.1 Finite braid-word certificates

Throughout this subsection, P is the same local Lickorish–Millet construction on actual finite nonempty decorated oriented diagrams, after the change of variables in Theorem 3.18. We use its supplied ordinary oriented Reidemeister-I/II/III and positive page-isotopy invariance (Literature input 2), its skein and initialization (Theorem 3.18), the full named-record equality (Lemma 3.20), and the split-circle factor $\delta = (a - a^{-1})/z$ (Lemma 3.19). These are not a general ambient-invariance premise.

Definition 7.7 (Standard closure of a supplied word). Fix $n \geq 1$. A finite word w uses letters $\sigma_i^{\pm 1}$, $1 \leq i < n$. Each letter is its standard oriented two-strand braid gadget, with every strand traversed in the increasing word coordinate. A positive letter has a determinant-positive crossing; its negative letter switches that crossing. At a letter the two adjacent positions exchange. Glue equal-numbered ports of successive gadgets, and close bottom port i to top port i by an oriented crossing-free exterior arc. Denote the resulting actual diagram by $D_n(w)$.

The closing arcs can be chosen without crossings: their endpoint pairs are nested in boundary order. Give the rightmost pair the innermost exterior U-shaped arc, the next pair a wider arc, and continue outward. Round the finitely many corners in disjoint clean discs. The braid is inside its rectangle and the closing interiors are outside, so no closure crossing is added. These choices affect no decorated record.

Equivalently, the record is the finite directed port network: each crossing gadget has two visits, each output is glued to the next input at that position, and the last outputs are glued to the first inputs. The oriented cycles are all components, including cycles without crossings. Each visit retains its crossing name, over/under bit and crossing sign. This follows the actual strands of $D_n(w)$; it makes no assertion that every abstract record is realizable. The empty one-strand word is one circle; the empty n -strand word has n circles.

Lemma 7.8 (Finite Artin substitutions). *Insertion or deletion of adjacent inverse letters, disjoint Artin commutations, and either uniform-sign three-strand Artin relation preserve $P_{D_n(w)}$, inside any fixed surrounding word. Consequently this scalar is well defined on the stated Artin generator-and-relation quotient.*

[Proof in Part II, page 184.](#)

Lemma 7.9 (Closure cut-slide and conjugation). *For two n -strand words u, v ,*

$$P_{D_n(uv)} = P_{D_n(vu)}. \quad (65)$$

Conjugate words also have equal closure values, by a finite derivation. The exponent sum is unchanged under these operations.

[Proof in Part II, page 185.](#)

Lemma 7.10 (Identity and free strand). *The empty n -strand word has value δ^{n-1} . If w is embedded in $n + 1$ strands without a letter $\sigma_n^{\pm 1}$, then*

$$P_{D_{n+1}(w)} = \delta P_{D_n(w)}. \quad (66)$$

Proof. The identity has no visits and exactly n components, so the source initialization after substitution gives δ^{n-1} . The additional unused strand and its closure give one component without crossing visits. Every old component keeps its exact record, including old empty occurrence words. The new actual closure has the same full named record as the old closure together with a crossing-free circle. Apply Lemmas 3.20 and 3.19. No assertion that a nested projected circle is spatially separated in a preferred presentation is needed. $\square$

Lemma 7.11 (Forward positive stabilization). *For every supplied n -strand word w , including $n = 1$,*

$$P_{D_{n+1}(w\sigma_n)} = P_{D_n(w)}. \quad (67)$$

Exactly one strand and one positive exponent have been added.

Proof. Track the strand entering the last gadget at position n . It visits the crossing over, exits at $n + 1$, and its exterior closure returns to top position $n + 1$. This position has no crossing in w . The strand therefore reaches the last gadget again, now under, and exits at n . Every other strand and old crossing visit is unchanged. The detour replaces one formerly crossing-free portion of an old closed component and creates no extra component. Its new consecutive visits are first over and then under, with positive crossing sign.

Insert an ordinary positive Reidemeister-I curl in the corresponding nonsingular interval of the actual old closure, choosing its first visit over. Such a template exists with both collars fixed. One central arc is $c(t) = (t^2 - 1, t - t^3)$, $-2 \leq t \leq 2$. Equality of its first coordinate at distinct parameters forces the parameters to be $t, -t$; equality of the second then gives $2t(1 - t^2) = 0$. Distinctness excludes $t = 0$, leaving the unique pair $\{-1, 1\}$. Its derivative is never zero: the first derivative coordinate vanishes only at 0, where the second is 1. Finally, $c'(-1) = (-2, -2)$ and $c'(1) = (2, -2)$ have ordered determinant $8 > 0$. The first visit is -1 ; assigning it over gives the required sign. Scale into a clean disc and attach crossing-free endpoint collars to obtain the isolated positive curl relative to the old interval. Its removal is precisely a supplied ordinary Reidemeister-I move.

The curled old diagram and $D_{n+1}(w\sigma_n)$ have identical full named records: the two specified visits replace one old gap, and all other components and visits remain. Record equality and local Reidemeister I give (67). The argument also applies when the chosen old component is crossing-free. It asserts no converse for arbitrary equal-link braid presentations. $\square$

Lemma 7.12 (The actual quadratic skein triple). *For the Jones notation of Lemma 7.6, the last crossing of $w\sigma_i^2$ gives*

$$P_{J, D_n(w\sigma_i^2)}(t, x) = txP_{J, D_n(w\sigma_i)}(t, x) + t^2P_{J, D_n(w)}(t, x). \quad (68)$$

Proof. Switching the last crossing gives $w\sigma_i\sigma_i^{-1}$. Its oriented smoothing gives $w\sigma_i$: both strands point in the same braid direction, so smoothing connects each incoming port to the outgoing port in the same position and removes precisely the last gadget. Lemma 7.8 removes the inverse pair. The actual Jones-convention skein is $t^{-1}P_{J,+} - tP_{J,-} = xP_{J,0}$. Multiplying by t and moving the negative-resolution term to the right proves (68). No diagram outside this local triple, and no assumed Hecke relation, enters the calculation. $\square$

7.3.2 Parameters and the selected finite source

Write q for the Hecke parameter, η for the Markov value of Wenzl's projector, and z_J for the Markov value of Jones's generator. The two auxiliary weight parameters are $\mathfrak{a}, \mathfrak{b}$; they are not HOMFLY variables. A partition is denoted by λ , whereas the Jones polynomial parameter will be λ_J . For $n \geq 1$, let $H_n(q)$ be the unital algebra with generators $g_1, \dots, g_{n-1}$ and relations

$$\begin{aligned} g_i^2 &= (q-1)g_i + q, & g_i g_j &= g_j g_i \quad (|i-j| > 1), \\ g_i g_{i+1} g_i &= g_{i+1} g_i g_{i+1}. \end{aligned} \quad (69)$$

Here H_1 is the scalar field. Initially q is indeterminate, so $q, 1 - q^h$ for $h > 0$, and $1 + q$ are nonzero field elements. The continuous matrix argument will instead use real $q > 0$, including the removable value $q = 1$. It will not evaluate an uncanceled weight at a zero of its denominator.

For a box (i, j) put $c(i, j) = j - i$. A standard tableau fills a Young diagram increasingly along its rows and columns. Write $c_T(k)$ for the content of the box occupied by k , and $s_i T$ for the tableau obtained by exchanging $i, i+1$, when it is standard. Put

$$\mathfrak{a}_d(q) = \frac{1 - q^{d+1}}{(1+q)(1-q^d)} \quad (d \neq 0), \quad e_i^W = \frac{q - g_i}{1 + q}. \quad (70)$$

Literature input 6 (Finite tableau matrices and diagram facts). For real $q > 0, q \neq 1$, the space V_λ with orthonormal standard tableau basis has the representation

$$\pi_\lambda(e_i^W)v_T = \mathfrak{a}_d v_T + \sqrt{\mathfrak{a}_d \mathfrak{a}_{-d}} v_{s_i T}, \quad d = c_T(i) - c_T(i+1). \quad (71)$$

The square root is nonnegative; the second term is zero for a forbidden swap. Deleting the largest entry gives the literal restriction

$$V_\lambda|_{H_{n-1}} = \bigoplus_{\mu \nearrow \lambda} V_\mu \quad (|\lambda| = n \geq 2), \quad (72)$$

with one block for each removable box. Two distinct entries differing by at most two have nonzero content difference. Distinct partitions of the same size $n \geq 3$ have different sets of immediate predecessors. These are [14, Lemma 2.1, eqs. (2.3)–(2.6), their relation and branching proofs, and Lemma 2.11(a),(b), pp. 361–364, 371–372]. No irreducibility, completeness, trace-weight or root-of-unity theorem is included in this input.

The defining formula fixes the first-row convention: on a one-row tableau $d = -1$, so $e^W = 0$ and $g = q$; on a one-column tableau $d = 1$, so $e^W = 1$ and $g = -1$. Wenzl's printed p. 365 base-case sentence interchanges those labels and is not used. His printed (3.6) is also not a normalized-weight premise: on one box it gives $(q-1)/(1-q) = -1$. The normalized product is proved below; neither that erroneous expression nor the weight assertion of Theorem 3.6(a) is imported after a silent correction.

7.3.3 Finite spanning and the continuous representation family

Lemma 7.13 (Spanning without a dimension assumption). *For every $n \geq 1$, $\dim H_n \leq n!$, and*

$$H_{n+1} = \text{span}(H_n \cup H_n e_n^W H_n). \quad (73)$$

Proof. For $n \geq 2$ set $u_j = g_{n-1} \cdots g_j$ for $j < n$, and $u_n = 1$. The space $\sum_{j=1}^n H_{n-1} u_j$ contains 1. Its stability under right multiplication follows from the exhaustive identities

$$\begin{aligned} u_j g_i &= g_i u_j & i < j-1, \\ u_j g_{j-1} &= u_{j-1} & j > 1, \\ u_j g_j &= (q-1)u_j + q u_{j+1} & j < n, \\ u_j g_i &= g_{i-1} u_j & i > j. \end{aligned} \tag{74}$$

The first identity commutes remote generators, the second extends the descending word, and the third uses the rightmost quadratic relation. For the fourth, commute the appended g_i past $g_j, \dots, g_{i-2}$, apply the adjacent braid relation, and commute the resulting left g_{i-1} past $g_{i+1}, \dots, g_{n-1}$. Every extracted left generator lies in H_{n-1} . If $j = n$, $u_n g_i \in H_{n-1}$ for $i \leq n-2$ and $u_n g_{n-1} = u_{n-1}$; these are the remaining cases. Word-length induction proves the span, and then $\dim H_n \leq n \dim H_{n-1} \leq n!$. At the next size every nontrivial descending word is g_n times a word in H_n , so H_{n+1} is spanned by H_n and $H_n g_n H_n$. Substitution $g_n = q - (1+q)e_n^W$ proves (73). This is an upper bound, not yet noncollapse. $\square$

Lemma 7.14 (Tableau count and complete matrix family). *Let f_λ be the number of standard tableaux of λ . For each n and real $q > 0$, the representations in Literature input 6 form the complete inequivalent irreducible family, their direct sum is faithful, and*

$$\sum_{|\lambda|=n} f_\lambda^2 = n!, \quad H_n(q) \simeq \prod_{|\lambda|=n} \text{End}(V_\lambda). \tag{75}$$

The matrices are continuous on $q > 0$, including $q = 1$, where this is the complete family of $\mathbb{C}[S_n]$ with one-row label trivial.

[Proof in Part II, page 185.](#)

7.3.4 The finite product and its two branching identities

For a partition λ define its hook length and row statistic by

$$h_\lambda(i, j) = \lambda_i - j + \lambda'_j - i + 1, \quad n(\lambda) = \sum_i (i-1)\lambda_i. \tag{76}$$

The weights to be proved are the rational functions

$$W_\lambda(\mathbf{a}, \mathbf{b}; q) = q^{n(\lambda)} \prod_{(i,j) \in \lambda} \frac{\mathbf{a} - \mathbf{b} q^{j-i}}{1 - q^{h_\lambda(i,j)}}, \quad W_\emptyset = 1. \tag{77}$$

All calculations through (83) are in a rational-function field; subsequent specialization never divides by a specialized weight.

Lemma 7.15 (Hook product, corners and one-box ratio). *Let $\mathcal{A}$ and $\mathcal{B}$ be the addable and removable contents of μ , and put $x_c = q^c$, $y_d = q^d$. Define*

$$D_c = x_c^{-1} \frac{\prod_{d \in \mathcal{B}} (x_c - y_d)}{\prod_{e \in \mathcal{A} \setminus \{c\}} (x_c - x_e)}. \tag{78}$$

Then

$$W_{\mu+c} = W_\mu \frac{(\mathbf{a} - \mathbf{b} x_c) D_c}{1 - q}. \tag{79}$$

Proof in Part II, page 187.

Lemma 7.16 (Ordinary and weighted branching). *For any partition μ ,*

$$\sum_{c \in \mathcal{A}} W_{\mu+c} = \frac{\mathfrak{a} - \mathfrak{b}}{1 - q} W_{\mu}. \quad (80)$$

If μ is nonempty and d_0 is any removable content, then

$$\sum_{c \in \mathcal{A}} \mathfrak{a}_{d_0-c} W_{\mu+c} = \frac{q\mathfrak{a} - \mathfrak{b}}{(1+q)(1-q)} W_{\mu}. \quad (81)$$

In particular, the substitutions

$$\mathfrak{a} = 1 - (1+q)\eta, \quad \mathfrak{b} = q - (1+q)\eta \quad (82)$$

give

$$\sum_c W_{\mu+c} = W_{\mu}, \quad \sum_c \mathfrak{a}_{d_0-c} W_{\mu+c} = \eta W_{\mu}. \quad (83)$$

The second identity requires a removable box, not the empty partition.

Proof. Sum (79) of Lemma 7.15 and apply the two different identities in (369); this proves (80). For the weighted identity put $u = q^{d_0}$. Addable and removable contents are disjoint, so $d_0 - c \neq 0$. Direct substitution gives

$$(1+q)\mathfrak{a}_{d_0-c} = \frac{x_c - qu}{x_c - u} = 1 + \frac{(1-q)u}{x_c - u}. \quad (84)$$

Consequently

$$(1+q) \sum_c \mathfrak{a}_{d_0-c} W_{\mu+c} = \frac{W_{\mu}}{1-q} \sum_c \frac{(x_c - qu)(\mathfrak{a} - \mathfrak{b}x_c)D_c}{x_c - u}. \quad (85)$$

The first part after the split in (84) is $\mathfrak{a} - \mathfrak{b}$, by (369). The second part, before multiplication by $(1-q)u$, is $-\mathfrak{a}/u$, by displays (370) and (371) of Lemma 7.15. Their sum is therefore $(\mathfrak{a} - \mathfrak{b}) - (1-q)\mathfrak{a} = q\mathfrak{a} - \mathfrak{b}$. Substitution in (85) proves (81). Finally (82) gives $\mathfrak{a} - \mathfrak{b} = 1 - q$ and, separately, $q\mathfrak{a} - \mathfrak{b} = (1+q)(1-q)\eta$; cancellation proves (83). $\square$

7.3.5 Construction and uniqueness of the full Markov trace

Theorem 7.17 (Arbitrary-element trace and Jones weights). *With the parameters (82) of Lemma 7.16, the formula*

$$\tau_n(X) = \sum_{|\lambda|=n} W_{\lambda} \operatorname{Tr}(\pi_{\lambda}(X)) \quad (86)$$

uses ordinary, not dimension-normalized, matrix traces and is the unique normalized compatible cyclic trace with

$$\begin{aligned} \tau_n(1) &= 1, & \tau_{n+1}|_{H_n} &= \tau_n, \\ \tau_{n+1}(Xe_n^W) &= \eta\tau_n(X) \quad (X \in H_n). \end{aligned} \quad (87)$$

Its Jones generator and weight parameters are

$$e_i^J = \frac{1+g_i}{1+q} = 1 - e_i^W, \quad z_J = q - (1+q)\eta, \quad \mathfrak{a} = 1 - q + z_J, \quad \mathfrak{b} = z_J. \quad (88)$$

In particular $\tau_{n+1}(Xg_n) = z_J\tau_n(X)$ for every X .

Proof in Part II, page 188.

7.3.6 Cycle characters and projector ranks

Lemma 7.18 (Hook character and rank calculation). *At $q = 1$, the character of an n -cycle in V_λ is zero unless $\lambda = (n - k, 1^k)$, $0 \leq k < n$; for that hook it is $(-1)^k$. For $n \geq 2$ and any $q > 0$, the hook dimension and the rank of the Jones projector are*

$$f_{(n-k, 1^k)} = \binom{n-1}{k}, \quad \text{rank}(e_1^J) = \binom{n-2}{k}, \quad \frac{\text{rank}(e_1^J)}{f_{(n-k, 1^k)}} = \frac{n-1-k}{n-1}. \quad (89)$$

Out-of-range binomial coefficients are zero. No generator-rank assertion is made at $n = 1$.

Proof in Part II, page 189.

7.3.7 Uniform full-twist words and the torus trace

Lemma 7.19 (Full-twist identity from the braid relations). *For $n \geq 2$ set*

$$\delta_n = g_1 \cdots g_{n-1}, \quad \Delta_n = (g_1)(g_2 g_1) \cdots (g_{n-1} \cdots g_1), \quad \delta_1 = \Delta_1 = 1. \quad (90)$$

Then $\Delta_n g_i = g_{n-i} \Delta_n$ for $1 \leq i < n$, and $\delta_n^n = \Delta_n^2$ is central.

Proof in Part II, page 190.

Lemma 7.20 (Continuous specialization and hook-only torus trace). *For positive coprime integers n, m with $n \geq 2$ and δ_n the word of Lemma 7.19, nonhook shapes have $\text{Tr}(\pi_\lambda(\delta_n^m)) = 0$, and*

$$\text{Tr}(\pi_{(n-k, 1^k)}(\delta_n^m)) = (-1)^k q^{(n-1-k)m}. \quad (91)$$

These statements first hold for $q > 0$ and then give the corresponding rational trace identities at generic parameters.

Proof in Part II, page 191.

7.3.8 Identification with the local diagram polynomial

Lemma 7.21 (Trace-to-skein normalization on actual braid closures). *For each actual nonempty finite oriented decorated diagram D , put $P_J(D; t, x) = P_D(t^{-1}, x)$, the same local LM construction in Jones's convention. Lemma 7.6 and Theorem 3.18 give*

$$P_J(D; t, x) = P_D(t^{-1}, x), \quad t^{-1} P_J(D_+) - t P_J(D_-) = x P_J(D_0), \quad P_J(\bigcirc) = 1. \quad (92)$$

Introduce independent nonzero variables s, ξ and set

$$q = \xi^2, \quad \lambda_J = s^2, \quad t = s\xi, \quad x = \xi - \xi^{-1}, \quad K = \frac{1 - \lambda_J q}{s(q - 1)}. \quad (93)$$

For $n \geq 1$, let β be an element of the Artin generator-and-relation group, supplied by a finite word w , and let $e(\beta)$ be its exponent sum. In the formulas below $\widehat{\beta}$ denotes the actual standard diagram $D_n(w)$ of Definition 7.7; its polynomial is independent of the representing word by Lemma 7.8. The notation $\tau_n(\beta)$ means evaluation on its image in $H_n(q)$. Then

$$P_J(\widehat{\beta}; t, x) = K^{n-1} s^{e(\beta)} \tau_n(\beta), \quad z_J = \frac{q - 1}{1 - \lambda_J q}. \quad (94)$$

No assertion that every link is a braid closure or that equal closures are Markov-equivalent is required.

Proof in Part II, page 192.

Theorem 7.22 (The finite Jones torus formula). *For positive coprime n, m with $n \geq 2$, put $[j]!_q = \prod_{i=1}^j (1 - q^i)$ and $[0]!_q = 1$. Then the positive Artin closure of δ_n^m (the word of Lemma 7.19) satisfies*

$$P_J(\widehat{\delta_n^m}; t, x) = \frac{(1 - q)\lambda_J^{(n-1)(m-1)/2}}{(1 - q^n)(1 - \lambda_J q)} \cdot \sum_{\substack{\beta, \gamma \geq 0 \\ \beta + \gamma = n-1}} \frac{(-1)^\beta q^{\beta m + \gamma(\gamma+1)/2} \prod_{i=-\gamma}^{\beta} (q^i - \lambda_J q)}{[\beta]!_q [\gamma]!_q}, \quad (95)$$

under (93) of Lemma 7.21, with the literal word closure of Definition 7.7. This agrees with the formula in Jones's Theorem 9.7; its weight, representation and local closure-evaluation reductions are proved here.

Proof in Part II, page 192.

Lemma 7.23 (Terminal coefficient of the positive torus closure). *Let $2 \leq n < m$ be coprime, and let D be the actual positive closure of $(\sigma_1 \cdots \sigma_{n-1})^m$. Then*

$$[a^{-(n-1)(m+1)} z^0] P_D(a, z) = \frac{(-1)^{n-1}}{n} \binom{m-1}{n-1}. \quad (96)$$

Proof in Part II, page 193.

Proposition 7.24 (corner coefficient of $T(r, 2r+1)$). *For $r \geq 2$,*

$$[a^{2-2r^2} z^0] H_{T(r, 2r+1)}(a, z) = (-1)^{r-1} \text{Cat}_r,$$

in the normalization of Literature input 1.

Proof. Use Lemma 7.23 with $n = r$ and $m = 2r+1$. These integers are coprime since any common divisor divides $(2r+1) - 2r = 1$, and $2 \leq r < 2r+1$. The terminal exponent is

$$-(r-1)((2r+1)+1) = -(r-1)(2r+2) = 2 - 2r^2. \quad (97)$$

The coefficient supplied by that lemma has magnitude

$$\begin{aligned} \frac{1}{r} \binom{2r}{r-1} &= \frac{(2r)!}{r(r-1)!(r+1)!} \\ &= \frac{(2r)!}{r!(r+1)!} \\ &= \frac{1}{r+1} \binom{2r}{r} = \text{Cat}_r. \end{aligned} \quad (98)$$

The first equality expands the binomial coefficient, the second uses $r(r-1)! = r!$, and the third uses $(r+1)! = (r+1)r!$. The sign from the lemma is $(-1)^{r-1}$. The positive standard closure is the specified $T(r, 2r+1)$; Theorem 3.18 identifies its local P with H in the stated normalization. This proves the unchanged coefficient statement. $\square$

Theorem 7.25 (state sum of the stars). *For every $r \neq 0$, $C(K_r) = \text{sgn}(r)(-1)^{|r|} \text{Cat}_{|r|} = A(K_r)$, and $C(K_0) = -1 = A(K_0)$.*

Proof in Part II, page 194.

Remark 7.26 (conventions of the torus sources). The following three conventions relate Proposition 7.24 and Theorem 7.25 to the cited sources. (1) Jones writes $t = \sqrt{\lambda q}$ and $x = \sqrt{q} - 1/\sqrt{q}$ (his Proposition 6.2) with the skein $t^{-1}P_+ - tP_- = xP_0$, so $a = t^{-1}$ and $z = x$ in the normalization of Literature input 1; his $[k]!$ in equation (9.6) and Theorem 9.7 is $\prod_{i=1}^k (1 - q^i)$, not the q -integer factorial, which would be off by $(1 - q)^{n-1}$; Section 7.3 uses the same bracket. (2) Gorsky, Brini–Eynard–Mariño and Dunfield–Gukov–Rasmussen use the mirror: their a (BEM’s c) is this document’s a^{-1} (DGR, Section 2.3), so the corner coefficient of Proposition 7.24 is their *top* coefficient ($k = n - 1$, Gorsky’s Corollary 3.3, first formula), not their bottom one. (3) Jones’s equation (9.6) prints its summation condition as “ $\alpha, \beta \geq 0$ ” where “ $\gamma, \beta \geq 0$ ” is required (Theorem 9.7 prints it correctly); the misprint is not copied here. The source comparison also records that $\text{mindeg}_a H_{T(r, 2r+1)}(a, 0)$ is exactly $2 - 2r^2$, so the corner coefficient $c(K_r)$ is nonzero; Corollary 3.65 permits a zero corner coefficient and nothing in this document uses the nonvanishing.

7.4 Convex polygons

Proposition 7.27 (convex polygons). *Let P be a strictly convex n -gon (vertices in convex position, traversed in boundary order), which is generic. Then $A(P) = C(P) = 1$ if P is clockwise and $A(P) = C(P) = (-1)^n$ if counterclockwise.*

Proof. Genericity: no three vertices of a strictly convex polygon are collinear and there are no crossings. C : no crossing, one uniform subpolygon with $c = 1$ (Lemma 3.70(i)) and prefactor $(-1)^\ell$, with $\ell = n$ counterclockwise and $\ell = 0$ clockwise. A : let τ be the common turn sign. For any root, three boundary vertices a_x, a_y, a_z with $x < y < z$ are in boundary order, so $\chi(a_x, a_y, a_z) = \tau$ and every near sign $d = \chi(a_z, a_y, a_x)$ and every far sign $h = \chi(a_j, a_r, a_i)$ equals $-\tau$ (Lemma 1.4(i)). Every ordinary gate $(d - h)/2$ vanishes, so $b_{[i, j]} = 0$ whenever $j - i \geq 2$ and only the all-singleton composition survives at the root, with $n - 2$ root gates $(d + h)/2 = -\tau$: $A(P) = (-\tau)^{n-2}$, which is $(-1)^n$ counterclockwise and 1 clockwise. $\square$

8 The decay region

Definition 8.1 (the one-incoming decay region). A generic polygon P on n vertices lies in $\mathcal{D}_1$ if, after an orientation-preserving affine change of coordinates, the edge vector $\tilde{\lambda}_1 = (-L, 0)$ with $L > 0$ and every other edge vector $\tilde{\lambda}_i$, $i \neq 1$, has strictly positive first coordinate. This is the condition $\tilde{\lambda}_i^+ > 0$ for $i \neq 1$ in the definition of the decay region in the letter. The letter denotes this decay condition by $\mathcal{R}_1$; here $\mathcal{D}_1$ distinguishes it from the regular locus $\mathcal{R}_n$, whose subscript is the number of vertices.




Figure 12: The maximal decomposition $S_{\max}$ in the decay region $\mathcal{D}_1$ (the letter’s $\mathcal{R}_1$): $k + 1$ embedded subpolygons.

Lemma 8.2 (positive-affine invariance). *Let $P = (\mu_1, \dots, \mu_n)$ be a generic labelled polygon, $n \geq 3$, and let $f(x) = Mx + b$, where M is a real 2×2 matrix, $\det M > 0$ and $b \in \mathbb{R}^2$. Write $fP = (f\mu_1, \dots, f\mu_n)$. Then fP is generic and*

$$\chi_{ijk}(fP) = \chi_{ijk}(P), \quad (99)$$

$$\ell(fP) = \ell(P), \quad (100)$$

$$\text{rot}(fP) = \text{rot}(P), \quad (101)$$

$$C(fP) = C(P). \quad (102)$$

The map f identifies the crossing pairs and their traversal orders, the decompositions, and the carriers obtained by oriented reconnection with incoming-visit ownership. Corresponding carriers Q and fQ have the same number of corners, the same corner-turn signs, the same number of self-crossings and the same rotation. In particular their uniform or mixed classification agrees. These identifications commute with cyclic relabelling.

[Proof in Part II, page 196.](#)

Lemma 8.3 (crossings in the decay region). *Let $P \in \mathcal{D}_1$. Then every crossing involves E_1 ; the return chain $\mu_2, \mu_3, \dots, \mu_n, \mu_1$ (the edges $E_2, \dots, E_n$) is embedded; the crossings $y_1, \dots, y_k$ along E_1 , ordered as met when traversing E_1 , are met by the return chain in the reverse order; the Gauss word is $y_1 y_2 \dots y_k y_k \dots y_2 y_1$, so no two crossings interlace and every subset of $X(P)$ is a decomposition; and $\text{rot}(P) \in \{-1, 0, 1\}$.*

[Proof in Part II, page 197.](#)

Proposition 8.4 (the state sum in the decay region). *Let $P \in \mathcal{D}_1$ with k crossings, and let $\Lambda_0, \dots, \Lambda_k$ be the $k + 1$ subpolygons of the decomposition $S_{\max} = X(P)$ (a decomposition by Lemma 8.3). Then*

$$C(P) = \begin{cases} (-1)^{\ell(P)+k}, & \text{every } \Lambda_a \text{ is uniform,} \\ 0, & \text{otherwise,} \end{cases}$$

in particular $C(P) \in \{-1, 0, 1\}$.

[Proof in Part II, page 197.](#)

Remark 8.5. The decay-region formula in the letter is Proposition 8.4. Its agreement with the corresponding values of A , and the letter's observation that $|A(P)| \leq 1$ on $\mathcal{D}_1$, are consequences of Theorem 6.11; the bound for C follows directly from the state-sum calculation.

9 Linear relations: $U(1)$ decoupling and Kleiss–Kuijf

Throughout this section A is the root-independent tree coefficient (Theorem 6.2) and a polygon is described by its cyclic edge word $w = (w_1, \dots, w_n)$, $w_i \in \mathbb{R}^2$, $\sum_i w_i = 0$: the polygon P_w has vertices $\mu_j = \sum_{i < j} w_i$ (so $\mu_1 = 0$ and $\tilde{\lambda}_j = w_j$) and $A(w) = A(P_w)$ whenever P_w is generic. The final results are identities among the values of A on the generic members of finite families of polygons obtained by reordering the letters of one word; every polygon in such a family has the same edge vectors, so all lie in one affine configuration space.

9.1 Common edge data and its block walls

For a finite set I of $n \geq 3$ edge-occurrence labels, put

$$E_I = \{w = (w_i)_{i \in I} \in (\mathbb{R}^2)^I : \sum_{i \in I} w_i = 0\}, \quad w_S = \sum_{i \in S} w_i. \quad (103)$$

Different labels remain different occurrences even if their vectors coincide. For an unordered partition $\pi = \{B_1, B_2, B_3\}$ of I into three nonempty blocks, choose one ordering of its blocks and define

$$f_\pi(w) = \det(w_{B_1}, w_{B_2}). \quad (104)$$

Another ordering changes only its sign, because $w_{B_1} + w_{B_2} + w_{B_3} = 0$. We use the open sets

$$M_I = E_I \setminus \bigcup_{\emptyset \neq S \subsetneq I} \{w_S = 0\}, \quad M_I^* = \{w \in M_I : f_\pi(w) \neq 0 \text{ for every } \pi\}. \quad (105)$$

These are auxiliary domains; the final relations will retain their original stated domains by a separate local sign argument.

Lemma 9.1 (The distinct common block polynomials). *Each f_π is an absolutely irreducible polynomial on E_I . Two such polynomials are associates only when their unordered partitions agree. Every point-triple determinant of a cyclic edge word on I is, up to sign, one f_π . In a fixed cyclic word, two different triples of cut vertices give different partitions. Each subset-zero locus removed in (105) is a linear subspace of real codimension two. The set M_I^* is dense in E_I and nonempty.*

[Proof in Part II, page 199.](#)

Lemma 9.2 (Finite generic paths for the common edge family). *Any two points of M_I^* can be joined in M_I by two straight segments. At their common endpoint all f_π are nonzero. Along each open segment the block-polynomial zeros are finite, simple, and no two distinct f_π vanish at the same parameter.*

[Proof in Part II, page 199.](#)

9.2 The point-sign interface for the common family

Lemma 9.3 (Using the coefficient on the point-generic domain). *The root-independent coefficient A extends to every polygon satisfying (G1), by its same finite tree formula. On this domain it is cyclically invariant and depends only on the point chirotope. For any finite family of reordered edge words, an identity among these coefficients transfers from a dense subset of its common edge space to every point where the polygons appearing in the identity satisfy (G1).*

Proof. The tree formula is already defined on (G1). Fix such a labelled polygon. There is a ball in its vertex coordinates on which all its finitely many nonzero point determinants retain their signs. By the density clause of Lemma 1.26 — the generic polygons, those satisfying (G1) and (G2), are dense in $(\mathbb{R}^2)^n$ — that ball contains a (G1),(G2) polygon. Proposition 2.7 identifies every rooted coefficient of the original polygon with the corresponding one of this generic polygon. Theorem 6.2 equates those roots, proving the asserted extension. A cyclic change of origin of the edge word translates the polygon and changes the displayed root. Translation preserves determinants, and root independence removes the root change. This proves cyclic invariance.

For the final assertion, at a point of the indicated family domain take one ball in the closed edge space preserving every point determinant of every polygon appearing in the identity. This is

possible because there are only finitely many such nonzero determinants, each continuous in the edge data. A point of the dense subset exists in that ball. Every term has the same value there by Proposition 2.7, so the identity transfers term by term. No value is assigned to a polygon failing (G1) by this argument. $\square$

Lemma 9.4 (A point-triple response without a crossing hypothesis). *The formulas (12) and (13) remain valid for a path of pairwise distinct vertices whose sole zero point triple at the central parameter changes sign. No condition on crossing concurrences is needed. The coefficients $A_g(P_\pm)$ on the left are those of the two sides, which satisfy (G1) because every point triple other than the critical one is nonzero near the centre; all closed coefficients on the right are understood on (G1), as in Lemma 9.3.*

Proof in Part II, page 200.

Lemma 9.5 (Block-wall jumps in an edge word). *Let $n \geq 4$ and let a path in M_I cross just one block polynomial f_π simply, with all others nonzero there. For a quantity g that is constant on each side of that crossing, write g^- and g^+ for its values before and after the crossing along the path and put $\Delta g := g^+ - g^-$; the jump of a coefficient across the wall is Δ of that coefficient. A cyclic word whose three blocks of π are not consecutive has zero jump. For a word whose three blocks are consecutive, its jump is as follows.*

1. *If two blocks are singleton letters u, v , write its cyclic order as (u, v, T) and put $\delta = -\frac{1}{2}\Delta \operatorname{sgn} \det(w_u, w_v)$. Then its jump is $\delta A(u + v, T)$, for either spatial order of these three collinear vertices. Swapping u, v reverses δ and preserves the merged polygon.*
2. *If exactly one block is a singleton d , write the word (d, U, V) and put $\delta = -\frac{1}{2}\Delta \operatorname{sgn} \det(w_d, w_U)$. The jump is zero unless w_U, w_V are both negative multiples of w_d . In that case, with $d_U = -w_U$ and $d_V = -w_V$, it is*

$$\delta A(U, d_U) A(V, d_V). \quad (106)$$

The vectors d_U, d_V are positive parts of w_d and add to it. Interchanging the two blocks reverses δ and leaves these same oriented half words and their product.

3. *If no block is a singleton, the jump is zero.*

Every closed lower polygon displayed here has (G1) and between three and $n - 1$ vertices.

Proof in Part II, page 201.

9.3 Wall terms in the shuffle sum

For disjoint nonempty ordered lists X, Y of labels, and a distinguished closing label c outside them, let $I = X \sqcup Y \sqcup \{c\}$ and set

$$S_c(X, Y) = \sum_{\sigma \in \operatorname{Sh}(X, Y)} A(\sigma, c). \quad (107)$$

Here Sh preserves each list's linear order. A cyclic word in this family is tested for membership by cutting immediately after c : its remaining linear word must restrict to X and to Y . All labels refer to occurrences; a merged or split edge receives a new occurrence label when it enters a lower family.

Lemma 9.6 (Cancellation on a common block wall). *Suppose that every nontrivial shuffle sum of total polygon arity less than n vanishes whenever all of its polygons satisfy (G1). For $|I| = n \geq 4$, the function $S_c(X, Y)$ has zero total jump at each wall of a path supplied by Lemma 9.2.*

Proof in Part II, page 202.

9.4 Soft normalization and the shuffle identity

Lemma 9.7 (Redistributing the soft compensation). *Fix a generic closed hard edge word of length $m \geq 3$ and an admissible vector q for one of its insertion gaps. Let a closed family insert exactly the soft edge εq in that gap, with each hard edge equal to its original vector plus $O(\varepsilon)$. For all sufficiently small $\varepsilon > 0$ its polygon satisfies (G1), and its coefficient, rooted at any corresponding nonsoft edge, equals that of the ordinary soft insertion in the same gap. If the adjacent hard edges in traversal order are l, r , then*

$$A(\text{distributed insertion}) = \frac{h_q(r) - h_q(l)}{2} A(\text{hard word}), \quad h_q(e) = \text{sgn det}(e, q). \quad (108)$$

This equality is on a sufficiently short positive tail, not merely an equality of numerical limits.

Proof. Translate to fix one hard vertex. Each other vertex differs from its hard limit by $O(\varepsilon)$, since it is a finite sum of hard edge perturbations and possibly the soft edge. The two soft endpoints converge to the same hard vertex z_j . Every triple not containing both of them has three distinct hard limits, so its determinant converges to the corresponding nonzero hard determinant. For a triple containing both endpoints and another vertex, subtract the first endpoint from the other two in its determinant. The vector to the second endpoint is exactly εq . The vector to the third is $z_k - z_j + O(\varepsilon)$. Consequently the ordered triple determinant, with its appropriate alternating sign, is

$$\pm \varepsilon \det(q, z_k - z_j) + O(\varepsilon^2). \quad (109)$$

Admissibility makes its leading coefficient nonzero. This proves (G1) on a tail and determines all its signs. The ordinary insertion has precisely the same hard limits and the same leading determinants, so its full chirotope agrees on a common tail. Proposition 2.7 equates their coefficients at the corresponding nonsoft root. The ordinary family is covered by Theorem 2.24; its attachment signs from Definition 2.22 are $\chi_- = -h_q(l)$ and $\chi_+ = h_q(r)$. Substitution gives (108). Lemma 9.3 removes the root choice. No (G2) assertion about the distributed family is required. $\square$

Theorem 9.8 (Shuffle vanishing on the point-generic domain). *For nonempty disjoint ordered lists X, Y and a closing vector $c = -\sum_{x \in X} x - \sum_{y \in Y} y$, if every polygon (σ, c) , $\sigma \in \text{Sh}(X, Y)$, satisfies (G1), then*

$$S_c(X, Y) = 0. \quad (110)$$

[Proof in Part II, page 203.](#)

9.5 The word-algebra passage to the standard relation

Lemma 9.9 (The distinguished-letter shuffle identity). *Let words consist of distinct occurrence labels, and write $U \amalg V$ for the formal sum of their shuffle words. For a fixed finite alphabet, let L be the rational linear span of the proper shuffle sums $U \amalg V$, where U, V are nonempty and together use the alphabet exactly once. In its word space modulo L , for words X, Y and a distinguished letter i whose labels together use this fixed alphabet exactly once, one has*

$$XiY \equiv (-1)^{|X|} i(X^{\text{rev}} \amalg Y). \quad (111)$$

Empty X or Y is allowed. Concatenation is extended linearly in this formula; no assertion that L is a concatenation ideal is made.

[Proof in Part II, page 205.](#)

9.6 The stated linear relations

Theorem 9.10 (*U(1) decoupling*). *Let $(b_1, \dots, b_{n-1}, q)$ be edge vectors with $\sum_i b_i + q = 0$ such that every insertion polygon $P^{(i)} = P_{(b_1, \dots, b_i, q, b_{i+1}, \dots, b_{n-1})}$, $0 \leq i \leq n-2$, is generic. Then*

$$\mathcal{Q}_n(q \mid b) := \sum_{i=0}^{n-2} A(b_1, \dots, b_i, q, b_{i+1}, \dots, b_{n-1}) = 0. \quad (112)$$

Proof. Apply Theorem 9.8 with $X = (q)$, $Y = (b_1, \dots, b_{n-2})$ and closing edge $c = b_{n-1}$. The shuffles insert q in each of the $n-1$ gaps before or among $b_1, \dots, b_{n-2}$, with b_{n-1} last. They are exactly the polygons indexed by $0 \leq i \leq n-2$ in the statement. Their genericity implies (G1), so that theorem gives the displayed sum. $\square$

Theorem 9.11 (Kleiss–Kuijf). *Let X, Y be nonempty disjoint ordered words of edge vectors and $c = -\sum_{x \in X} x - \sum_{y \in Y} y$, and suppose every shuffle polygon $P_{(\sigma, c)}$, $\sigma \in \text{Sh}(X, Y)$, is generic. Then*

$$\mathcal{S}_c(X, Y) := \sum_{\sigma \in \text{Sh}(X, Y)} A(\sigma, c) = 0. \quad (113)$$

Consequently, for disjoint ordered words α, β and letters $1, n$ with all polygons below generic,

$$A(1, \alpha, n, \beta) = (-1)^{|\beta|} \sum_{\gamma \in \text{Sh}(\alpha, \beta^{\text{rev}})} A(1, \gamma, n). \quad (114)$$

Proof. The shuffle-sum assertion is Theorem 9.8 restricted to the stated generic family.

For the consequence, first take edge data in M_I^* for the labels appearing in the formula. Every ordering then satisfies (G1). Fix the closing occurrence n and extend $W \mapsto A(W, n)$ linearly to the rational word space on the other labels. Theorem 9.8 says exactly that this functional vanishes on the space L of Lemma 9.9. Apply that lemma with $X = \beta$, $i = 1$, $Y = \alpha$. Its left word is $\beta 1 \alpha$; cyclicity from Lemma 9.3 identifies $A(\beta, 1, \alpha, n)$ with $A(1, \alpha, n, \beta)$. Its right side is the signed sum of $A(1, \gamma, n)$ for $\gamma \in \text{Sh}(\beta^{\text{rev}}, \alpha)$. Commutativity of shuffle interleaving makes this precisely (114), including either empty list.

For arbitrary edge data under the final statement's genericity assumption, use the density of M_I^* from Lemma 9.1. A small ball preserves every point sign of every polygon actually displayed in (114). Choose a point of M_I^* in this ball and transfer its proved identity term by term using Lemma 9.3. This last step imposes no genericity assumption on additional permutations at the original data. $\square$

10 Computations

The chamber census is reported in Table I of the letter. The finite checks below are consistency checks, not proofs of universal identities.

Definition 10.1 (visible chambers and the census). *A visible chamber is a connected component of $\mathcal{W}_n$ (Definition 1.15). It is open and path connected by Lemma 1.16, and lies in the regular locus $\mathcal{R}_n$, every turn being nonzero on $\mathcal{W}_n$; hence rot is constant on it by Lemma 1.20(ii). A visible chamber of (n, r) is a visible chamber on which $\text{rot} = r$. By Definition 6.5 with Corollary 6.4, and by Lemma 3.30, the functions A^{cont} and C (Definition 3.29) are constant on visible chambers and invariant under the cyclic shift. Write $\nu(n, r)$ for the number of visible chambers of (n, r)*

modulo the cyclic shift, and $\nu^{\text{refl}}(n, 0)$ for the number of those classes at $r = 0$ that are mapped to themselves by an orientation-reversing affine map of the plane (which fixes $r = 0$ and exchanges r with $-r$ otherwise).

Computation 10.2 (Chamber census). The following values were obtained by enumerating the realizable rank-three oriented matroids on the n edge lines and the line at infinity and gluing across the silent walls (Table I of the letter):

$n \setminus r $	0	1	2	3	total
3	—	2	—	—	2
4	1 (1)	4	—	—	5
5	8 (0)	10	6	—	24
6	57 (9)	84	68	—	209
7	772 (0)	1510	840	302	3424

where the entry at $|r| > 0$ is $\nu(n, r) + \nu(n, -r)$ and the parenthesis is $\nu^{\text{refl}}(n, 0)$. The entries $\nu(3, \pm 1) = 1$ each and $\nu(4, 0) = 1$ are proved (Theorem 6.7, base case; Theorem 5.5); the remaining entries are computational.

10.1 Enumeration by carrier arrangements

The counted equivalence is cyclic change of the starting vertex only; traversal reversal and ambient reflection are retained. To obtain the visible chambers of Definition 10.1, the carrier refinement glues these transitions: a vertex meeting a remote carrier outside its segment; concurrence of three pairwise nonadjacent carriers when at least one corresponding segment crossing is absent; and parallelism of nonadjacent carriers. These events change refinement data without changing the visible immersion. The quotient does not glue flat or cusp turn barriers, on-segment vertex-edge events, active triple events, or zero-edge boundaries. A zero edge is a boundary of the regular polygon space, not an additional manuscript wall.

Adjoin the line at infinity to the n carrier lines. Signs of all 3×3 carrier determinants encode a uniform rank-three oriented matroid on $n + 1$ elements. The finite encoding normalizes individual signs of the carrier equations, retains the infinity coorientation framing ambient orientation, and canonicalizes cyclic relabelling of finite carriers. For $n \leq 7$ there are at most eight elements. Richter–Sturmfels prove that the realization space of every uniform rank-three oriented matroid on at most eight elements is contractible, and hence has the isotopy property [36, Theorem 4.1, p. 400]. The realization-space convention here is that of the cited theorem; this is not a contractibility claim for the unnormalized space of all labelled polygon coordinates.

The recorded enumeration generates the fine states, forms their graph of single-basis mutations, and classifies mutations by exact chirotope signs and betweenness tests. A disjoint-set computation joins only the invisible edges listed above. Signed rotation is obtained by exact signed ray crossings. The pseudoline-generation recursion follows NumPSLA [38, Section 2.2 and Appendix F]. This is a finite computational implementation of the refinement and gluing, not a new connectivity theorem for arbitrary polygon size.

Each fine state is also reduced to its signed rotation, cyclic turn word and ordered signed crossings along every oriented edge. The recorded computation finds that this signature is constant on every invisible component and every signature has exactly one component through $n = 7$. At $n = 7$ both counts are 3424, with zero ambiguous signature fibres. This is the finite check of the visible-signature identification used in the census; the definition of a visible chamber does not assume it.

n	$r = -3$	-2	-1	0	1	2	3	Total
3	–	–	1	–	1	–	–	2
4	–	–	2	1	2	–	–	5
5	–	3	5	8	5	3	–	24
6	–	34	42	57	42	34	–	209
7	151	420	755	772	755	420	151	3424

Table 1: The recorded visible-chamber census by signed rotation, modulo cyclic start. Reflection and traversal reversal are not divided out. A dash marks an inadmissible pair.

10.2 Stored certificates and their scope

The fine-state counts modulo cyclic start are 2, 6, 152, 9100 and 1,270,188 for $n = 3, 4, 5, 6, 7$. In the largest case, two generation routes yield the same sorted list of 64-bit states: the NumPSLA recursion and the 3315 eight-point order types of Aichholzer–Aurenhammer–Krasser [37], expanded over 5760 choices of infinity element and cyclic labelling. The eight-point input count is also recorded in Rote’s Table 1 [38]; it counts order types, not polygon chambers. Each stored state file has 10,161,504 bytes. Their byte identity and recorded SHA-256 digest were checked when preparing this companion:

ee66db33c615c24634e6dc2c567b2e1cc4203e971194b7b5bdccb5888a0b8c93

This checks agreement of the stored fine-state inventories. It does not independently recheck the mutation classification, component computation or geometric interpretation of their quotient. No fresh enumeration is asserted here.

The recorded smaller checks are a separate integer vertex-order-type calculation through $n = 5$ and a separate Python dual-mutation computation through $n = 6$. The first is not complete for $n \geq 6$: an active triple transition can occur inside one vertex-order-type realization space, so the carrier refinement is needed. The census and finite signature test are computational results; they enter no proof of comparison or uniqueness.

10.3 The pentagon atlas

The signed pentagon row contains 24 chambers. Figure 13 shows their 12 ambient-reflection pairs, retaining traversal orientation and identifying only cyclic start before reflection is taken.

Computation 10.3 (finite check of Definition 2.4). The exact-arithmetic check implements Definitions 2.2 and 2.4 in the SM’s tail-labelled convention, with exact integer determinants and rational gates. It checks all expected reference values and exits unsuccessfully on a mismatch. Its literal integer star and convex fixtures have their circular-order chirotopes checked exactly before use. The recorded run verifies 94 reference root values: the two triangle values; every root of all four shifts of K_0 ; the values $A(K_{\pm r}) = \pm(-1)^r \text{Cat}_r$ for $r = 2, 3, 4$; and the convex values for $n = 4, 5, 6$ in both orientations. A physical four-vertex gate fixture also distinguishes the near and far input order.

With seed 1 it checks root independence and reversal on 229 accepted integer polygons satisfying (G1), with $4 \leq n \leq 7$. It checks the soft formula in all three sectors on 107 accepted insertion events, giving 531 comparisons across nonsoft roots: 39 events at arity 4, 33 at arity 5, and 35 at arity 6. Mutating the root gate to the ordinary gate, reversing the ordinary near/far order, or corrupting a nonzero reference value makes the executable fail in the recorded mutation checks.




Figure 13: One representative of each of the 12 ambient-reflection classes of visible oriented pentagons, modulo cyclic change of the starting vertex; traversal reversal is not additionally identified. Each numbered panel is an exact generic realization, decoded by the key at the top. The three $|r| = 2$ panels and five $|r| = 1$ panels represent pairs with opposite signed rotation; the four $r = 0$ panels represent the four remaining reflection pairs.

Reversing the argument order of χ inside the descent gate, or shifting the root index by one, leaves every printed value unchanged (χ is alternating, so the reversal exchanges the two gates up to sign; A is root-independent), so those two fidelity questions are code-reading obligations against Definition 2.2, not numerical ones. These are finite consistency checks of the definition; the random fixtures are not an exhaustive chamber enumeration or a proof of any universal identity.

Remark 10.4. The identification of visible chambers with the classes counted by the matroid enumeration requires that the realization set of each visible signature be connected within its fibre; for $n \leq 7$, this identification is a reported computational result. Definition 10.1 does not assume it.

Part II. Longer proofs

The proofs below are grouped by the sections in which their statements appear. Each statement links to its proof; proofs of a few paragraphs remain alongside their statements in Part I.

Proofs for Section 1

Proof of Lemma 1.16

Proof of 1.16. Fix a tuple P satisfying Definition 1.15. Nonzero edge vectors and nonzero turn determinants persist under small perturbations by continuity. Thus distinct adjacent segments still meet only at their common endpoint: their supporting lines are independent and have that point in common.

Suppose every vertex moves by at most ε . A point of an edge at affine parameter $t \in [0, 1]$ moves by at most $(1 - t)\varepsilon + t\varepsilon = \varepsilon$, by the triangle inequality. Consequently a vertex–segment distance, or a distance between two segments, decreases by at most 2ε . Every vertex of P has positive distance from every nonincident closed segment, since the segment is compact and does not contain the vertex. There are finitely many such pairs, so these separations all persist nearby. In particular, remote segments cannot acquire an endpoint contact.

Each disjoint pair of remote closed segments has positive compact distance; preserve all these distances as well. This step allows parallel or collinear disjoint segments. For an intersecting remote pair, the endpoint exclusion puts the intersection in both relative interiors, and weak genericity makes the two directions independent. Write their initial points as p, q and directions as u, v , and put $a = q - p$ and $D = \det(u, v) \neq 0$. The intersection equation is $tu - sv = a$. Taking its determinant with v gives $tD = \det(a, v)$. Taking its determinant with u gives $sD = \det(a, u)$, since $-\det(v, u) = D$. Thus the two parameters are

$$t = \frac{\det(a, v)}{D}, \quad s = \frac{\det(a, u)}{D}. \quad (115)$$

The denominator stays nonzero nearby. Both parameters, and hence the intersection point $p + tu$, vary continuously. Shrink the neighbourhood so the parameters remain strictly between zero and one. The two supporting lines then have exactly this one transverse intersection.

Intersection points belonging to distinct unordered pairs of remote edges at P are distinct: if two were equal, their union would contain at least three distinct edges whose relative interiors contain that point, contrary to (G2). Preserve the finitely many positive distances between these points. Every nearby intersecting remote pair was already intersecting at P , because the disjoint pairs remain disjoint. A nearby point in three edge interiors would involve pairwise remote edges, since adjacent interiors are disjoint. It would therefore identify the points of two distinct preserved crossing pairs, a contradiction. If there were initially fewer than two crossing pairs, the same argument gives a contradiction before any separation needs to be chosen. Thus (G2) persists.

Intersecting the finitely many neighbourhoods just constructed preserves every condition of weak genericity. This proves openness. A sufficiently small ball about any point is contained in the open set and is convex; it therefore lies in the connected component of its centre. Hence every component is open. Within one component, fix a point and consider the points reachable from it by a finite polygonal path. A small ball about a reachable point is reachable by adding a final segment. A small ball about an unreachable point cannot contain a reachable point, since a final segment would then reach its centre. The reachable set and its complement are both relatively open. Connectedness makes the reachable set the entire component, proving polygonal connectedness. $\square$

Proof of Lemma 1.26

Proof of 1.26. We use only Lemma 1.20: rotation is integer-valued and constant on paths in the regular locus; subdivision preserves it; regular triangles have rotation ± 1 ; and a regular n -gon satisfies $2|\text{rot}| < n$. The regular locus consists of polygons with nonzero edges and without antiparallel

consecutive directions. Thus the three stated conditions are necessary. We print the existence argument, including its polynomial-avoidance step.

Regular seeds. First let $r > 0$ and put $m = 2r + 1$, $\alpha = 2\pi r/m$, and $q = \exp(i\alpha)$. Identify the plane with $\mathbb{C}$ and take the ordered vertices $z_j = q^{j-1}$ for $1 \leq j \leq m$. The integers r, m are coprime, since a common divisor would divide $m - 2r = 1$. These vertices are therefore distinct. With cyclic indices, the edge directions satisfy

$$d_j = z_{j+1} - z_j = q^{j-1}(q - 1). \quad (116)$$

Here $q^m = 1$, so this includes the closing edge, and $q \neq 1$, so every edge is nonzero. Successive directions satisfy $d_j = qd_{j-1}$, also across the closing corner. Since $0 < \alpha < \pi$, every principal turn is exactly α . The seed is regular, and its total turn is

$$\sum_{j=1}^m \vartheta_j = m\alpha = 2\pi r. \quad (117)$$

Consequently its rotation is r . For negative r , reverse the seed constructed for $|r|$; reversal negates rotation and preserves regularity. This construction uses no genericity or concurrency claim about a regular star and does not refer to a later star definition.

For $r = 0$, use the four ordered vertices $(0, 0), (2, 2), (0, 2), (2, 0)$. Its four directions are $(2, 2), (-2, 0), (2, -2), (-2, 0)$, and its principal turns, starting at the first vertex, are respectively

$$-3\pi/4, \quad 3\pi/4, \quad 3\pi/4, \quad -3\pi/4. \quad (118)$$

Their sum is zero and none is antiparallel. This is a regular seed of rotation zero. In each case let m be the number of seed vertices. The stated admissibility conditions imply $n \geq m$.

Subdivision and a regular neighbourhood. Insert $n - m$ distinct interior subdivision points in traversal order on any seed edge. All refined edge vectors are positive multiples of the former vector; each new principal turn is zero and all old turns are unchanged. The resulting n -gon P_0 is regular with rotation r . The regular locus is open: nonzero vectors form an open set, and for two nonzero vectors the condition that their unit directions sum to a nonzero vector is open and excludes exactly antiparallel directions. Choose an open ball B about P_0 contained in this locus. Every point of B is connected to P_0 by a straight path in B , so it has rotation r by the preceding rotation-number lemma.

Explicit nonzero polynomials. All coordinates of all n vertices may vary in B . For every three distinct indices i, j, k , let

$$D_{ijk} = \det(\mu_j - \mu_i, \mu_k - \mu_i). \quad (119)$$

This polynomial is not identically zero: assign these three vertices the coordinates $(0, 0), (1, 0), (0, 1)$, respectively, obtaining value 1. The remaining coordinates can be arbitrary for this test.

For an edge from (x_i, y_i) to (x_{i+1}, y_{i+1}) , define its homogeneous line-coefficient row by

$$\ell_i = (y_i - y_{i+1}, x_{i+1} - x_i, x_i y_{i+1} - x_{i+1} y_i). \quad (120)$$

Both endpoints satisfy $\ell_i \cdot (x, y, 1) = 0$ by direct substitution. For each unordered triple of pairwise remote edges i, j, k , put

$$H_{ijk} = \det \begin{pmatrix} \ell_i \\ \ell_j \\ \ell_k \end{pmatrix}. \quad (121)$$

These three edges have six distinct endpoint indices. Hence we can assign their ordered endpoint pairs independently as $(0, 0), (1, 0)$; $(0, 0), (0, 1)$; and $(0, 1), (1, 1)$. The rows in this assignment are $(0, 1, 0)$, $(-1, 0, 0)$, and $(0, 1, -1)$; their determinant is -1 . Thus each H_{ijk} is a nonzero polynomial. Coinciding coordinates in this witness are harmless: it proves a polynomial identity is nonzero on the full coordinate space, rather than purporting to be a generic polygon or a point of B .

If the three supporting lines have a common finite point (x, y) , their coefficient matrix annihilates the nonzero vector $(x, y, 1)$. Its determinant then vanishes. Thus $H_{ijk} \neq 0$ excludes triple interior concurrence for these three edges. It is a sufficient condition; no converse is used.

Simultaneous avoidance in the small ball. Let F be the finite product of all the polynomials D_{ijk} , taking one order for each index triple, and all the polynomials H_{ijk} . An empty family of edge triples contributes the factor 1. The polynomial ring over $\mathbb{R}$ has no zero divisors (the highest nonzero monomials in lexicographic order multiply to a nonzero highest monomial), so F is nonzero. For completeness, a nonzero real polynomial cannot vanish on a nonempty open box. For one variable this follows from the finite root bound. Induct on the number of variables, writing the polynomial as a polynomial in the last variable with coefficient polynomials in the others. If it vanishes on a box, fixing the other variables makes it vanish on an interval; all its coefficients vanish at that fixed choice. They therefore vanish on the lower-dimensional box, and the inductive hypothesis makes every coefficient identically zero, a contradiction. Every open ball contains an open box. Hence some $P \in B$ has $F(P) \neq 0$. The same avoidance applies to any nonempty open subset of $(\mathbb{R}^2)^n$ in place of B , since it too contains an open box; with the genericity of such a P shown next, this proves the density clause.

At this P , the nonvanishing of every D_{ijk} gives (G1), in particular nonzero edges and non-collinear consecutive vertices. Adjacent edges can meet only at their common endpoint: their lines are distinct and nonparallel because the consecutive-vertex determinant is nonzero. A point in three edge interiors therefore would require three pairwise remote edges. Its concurrence would force the corresponding H_{ijk} to vanish, contrary to $F(P) \neq 0$. This proves (G2). The polygon is generic and has rotation r because it lies in B . This proves sufficiency and the lemma. $\square$

Proof of Lemma 1.30

Proof of 1.30. Write $A = \mu_{j-1}(0)$, $M = \mu_j(0)$ and $B = \mu_{j+1}(0)$. All central vertices are distinct: if two coincided, adjoining each of two further indices would produce two distinct collinear triples; there are two further indices because $n \geq 4$. This contradicts the unique-triple assumption. Thus all edges are nonzero. Every turn other than the one at j uses a different triple, so its determinant is nonzero. Strict betweenness gives $M = A + \lambda(B - A)$ for exactly one $\lambda \in (0, 1)$, proving (2). In particular the turn at j is positive-flat, not antiparallel. Continuity and the finitely many nonzero central determinants give (i).

No vertex contact at the centre. A vertex on the line of an edge not incident to it would give a collinear triple. The sole permitted triple can occur in this way only as B on the line of $[A, M]$ or A on the line of $[M, B]$. Strict betweenness puts each such point outside the corresponding closed segment. Thus no vertex lies on a nonincident closed segment. Two adjacent edges other than the flat pair have nonparallel directions, and meet only at their common endpoint. The flat pair consists of the successive subsegments $[A, M]$ and $[M, B]$, also meeting only at their common endpoint. Hence no adjacent edge interiors meet.

Crossing pairs. For a remote pair whose four endpoint-against-line triples avoid $\{j-1, j, j+1\}$, those determinants have nonzero limits and constant signs. Two segments with all endpoints off the other segment's line cross if and only if each line separates the other segment's two endpoints. This is the pair of strict opposite-sign tests. To see sufficiency, intersect the two lines: opposite

endpoint signs place their unique intersection inside each segment. Necessity follows by restricting the affine signed line equation to the other segment. The tests therefore give the same crossing answer at zero and on both sides. An actual intersection cannot be parallel: parallel intersecting lines would coincide, forcing an endpoint determinant to vanish.

The only remote pairs whose tests can use the critical triple are $\{j-2, j\}$ and $\{j-1, j+1\}$. The line of $E_{j-2}(0)$ meets the flat line only at A ; its other endpoint is off that line by the unique-triple assumption. Since $A \notin [M, B]$, the compact segments $E_{j-2}(0)$ and $E_j(0)$ are disjoint. Similarly the line of $E_{j+1}(0)$ meets the flat line only at $B \notin [A, M]$, so $E_{j-1}(0)$ and $E_{j+1}(0)$ are disjoint. Positive distances between these compact segment pairs persist under small changes of their endpoints. These pairs remain noncrossing on both sides. This covers all remote pairs, including when $n = 4$.

Order and crossing records. Each crossing at zero has two nonparallel directions. Solving its two-by-two line-intersection system therefore gives continuous crossing parameters, strictly between zero and one at zero. There are finitely many pairs, so all these strict inequalities persist on one interval. If two distinct central crossings had the same point, at least three edge interiors would contain that point. None of those edges could be adjacent, by the preceding paragraph. This would be a forbidden concurrence of three pairwise remote edges. Thus crossing points are distinct. In particular, on a fixed edge all crossing parameters are distinct. The finitely many nonzero parameter differences retain their signs by continuity; crossing visits neither collide nor reorder. The Gauss word is just these edgewise lists concatenated in traversal order. Pairing is indexed by the fixed edge pairs, and interlacement is read from that word. For each actual crossing, the nonzero direction determinant keeps its sign. The same physical strand therefore remains over in the positive over/under data of the statement. This proves (ii) without a triple-wall lemma.

Deletion and fusion. The deletion's point triples are central parent triples avoiding j , so they are all noncollinear. At zero the oriented segment $[A, B]$ is precisely the traversal of $[A, M]$ followed by $[M, B]$. No crossing of a remote edge with this segment can occur at M : that would put the parent vertex M on a nonincident edge, already excluded. Thus such a crossing lies in exactly one of the two open subsegments and inherits its transverse direction and its pairing. An edge remote from $[A, B]$ is remote from both old subsegments. The reverse implication needs care only for the neighbours at A, B : their possible opposite-subsegment pairs are exactly the two compact disjoint pairs treated above. They create no discarded crossing. One remote straight edge cannot meet both old subsegments, since it meets their common line at most once. Hence the crossing-label map is bijective.

If a crossing parameter on $[A, M]$ is a , its parameter on $[A, B]$ is λa ; if it is b on $[M, B]$, its fused parameter is $\lambda + (1 - \lambda)b$. Thus the old first-edge visits precede the old second-edge visits, in exactly their original orders. Equation (2) shows that every direction determinant involved changes only by a positive factor, preserving signs and over/under data. No triple of deletion edge interiors can concur: away from M it would give three parent edge interiors at the same point; at M it would put a remote parent edge through that parent vertex. Both are excluded. This proves genericity and all stated crossing data for Q . Finally, the finitely many strict noncollinearity, transverse crossing, disjointness and crossing-order conditions used above also prove openness of genericity near Q . By continuity, the connected small path $P(t) \setminus j$ stays in that open generic neighbourhood and hence in the chamber of Q . This proves (iii). $\square$

Proof of Lemma 1.31

Proof of 1.31. Case (F), including its central regularity and crossing data, is Lemma 1.30. We prove the remaining cases directly in the one-based tail convention $d_i = \mu_{i+1} - \mu_i$.

Central distinctness and regularity. The remoteness requirements give $n \geq 5$ in (V),(E) and

$n \geq 6$ in (T),(C). If two central vertices coincided, adjoining each of two further indices would give two different zero point triples. This contradicts the permitted zero set, which has at most one member. Thus all central vertices are distinct and every edge is nonzero. A zero turn would give a consecutive zero triple. There is no such triple in any of (V),(T),(E),(C): the point zero set is empty in (T), has no adjacent pair in (C), and has just the adjacent pair $a, a+1$ in (V),(E), because $M \notin \{a-1, a, a+1, a+2\}$. Hence every central turn in these cases is nonzero. In particular the centre has no kink and belongs to $\mathcal{R}_n$. Each other point determinant is nonzero at zero and is continuous; finitely many such determinants retain their signs nearby.

Which segment pairs can change. No remote pair can have collinear overlapping segments at the centre: its four distinct endpoint indices on a line would give more than one zero point triple. Apart from an endpoint incidence, a remote pair is therefore either disjoint or a transverse interior pair, and Lemma 1.9 applies. An endpoint incidence requires a zero point triple containing the two endpoints of its base edge. In (C) the only zero triple contains no adjacent pair, and in (T) there is no zero triple, so these cases have no endpoint incidence. In (V),(E), the only possible incidence is the specified vertex M on the line of E_a . The only possibly affected remote pairs are consequently $\{a, M-1\}$ and $\{a, M\}$. All other pairs satisfy the stability lemma.

The two vertex-edge branches, with the finite-segment test. Write $A = \mu_a, B = \mu_{a+1}, D = B-A, M = \mu_M$, and take in turn $U = \mu_{M-1}$ and $U = \mu_{M+1}$. Put

$$h = \det(D, M - A), \quad h_U = \det(D, U - A). \quad (122)$$

At the centre $h = 0$ and $h_U \neq 0$: another zero h_U would give a second point triple, since both neighbour indices differ from $a, a+1, M$. Shrink so h_U and $h_U - h$ are nonzero. The intersection of the supporting line of E_a with the line MU has the exact form

$$q_U = M + r_U(U - M), \quad r_U = -\frac{h}{h_U - h}. \quad (123)$$

Indeed its height is $h + r_U(h_U - h) = 0$. When h and h_U have opposite signs, $r_U = |h|/(|h| + |h_U|)$ lies strictly between zero and one. When their signs agree, the height of every point in the relative interior of MU has that common sign, so there is no such intersection. In the first case $r_U \rightarrow 0$ and $q_U \rightarrow M(0)$. In case (V), $M(0)$ is strictly inside $E_a(0)$. The parameter of q_U on the line AB therefore tends to a number in $(0, 1)$ and remains in $(0, 1)$ nearby. This supplies the second, finite-segment condition, not merely the straddling of the supporting line. The crossing is transverse since the two line directions have determinant $h_U - h \neq 0$.

Consequently each contact pair crosses precisely when the sign of h is opposite to its neighbour sign. The two neighbour signs are equal in the bigon branch, giving both crossings or neither. They are opposite in the sliding branch, giving precisely one crossing on each side. By the defining sign change of h , the asserted exchanges follow. The predecessor leg is E_{M-1} and the successor leg is E_M , so these are exactly the displayed tail-indexed pairs.

Exterior extension. At the centre in (E), each contact leg meets the supporting line of E_a only at $M(0)$, because $h_U(0) \neq 0$. That point is outside the closed segment $E_a(0)$. Thus each contact leg is compact-disjoint from $E_a(0)$, and both pairs remain disjoint by Lemma 1.9. All remote crossing pairs are therefore constant, including at the centre, in (E), as they already are in (C) and (T).

Orders and localization. Consider two persistent crossings on a common edge whose central parameters coincide. The other two edges then meet that base edge at the same point. If those other edges are remote, their common interior point gives a member of Z_c . If they are adjacent, their nonzero central turn implies that their lines meet only at their shared vertex. The coincidence would put that vertex on the base edge. In (E),(C) no such endpoint incidence exists; in (T) none

exists either, and its only concurrence is the named triple. It follows that all parameters in (E),(C) are centrally distinct and their orders are constant. Their fixed edge order then gives a fixed cyclic Gauss word, including at zero. In (T) the only possible ties are the two triangle visits on a bundle edge. Lemma 1.28 makes each such pair adjacent. Their differences change sign by the definition of (T), so they exchange order, and every other order is fixed by the stability lemma.

For (V), no persistent crossing has central point $M(0)$: any additional edge through that point would give a second endpoint-incidence triple. All persistent crossing parameters on E_a are therefore separated from its contact parameter. Those on either incident leg are separated from its endpoint at M . Formula (123) puts any newborn contact crossings arbitrarily close to those contact parameters. Finiteness gives neighbourhoods free of other crossings. Every other pair of persistent visits has distinct central parameters, since a tie would give either a forbidden triple concurrence or an additional vertex incidence. Their orders remain fixed. This proves the claimed localization and completes all cases. $\square$

Proof of Lemma 1.32

Proof of 1.32. First, all vertex coordinates at the centre are distinct. If two coincided, their triples with either of two further vertex indices would both be collinear. Such further indices exist because $n \geq 4$, contradicting the sole-triple hypothesis. All edges at the centre are therefore nonzero. Put $A = \mu_{j-1}$, $M = \mu_j$, $B = \mu_{j+1}$, and, as functions of t , put

$$u = M - A, \quad v = B - M, \quad T = \det(u, v). \quad (124)$$

Thus $\tau_j = \text{sgn } T$ on the punctured sides. The distinct collinear points A, M, B have a unique middle point; it is not M . This proves the two-case assertion. On each punctured side every point-triple determinant is nonzero, so its sign is constant by continuity.

The newborn crossing in case A. Let $w = \tilde{\lambda}_{j+1}$, so that E_g starts at B and has direction w . Here $\Delta = \det(u, w)$. If $\Delta(0) = 0$, the nonzero vector $w(0)$ would be parallel to the cusp line through A, M, B . Its other endpoint $\mu_{j+2}(0)$ would lie on that line too. The indices $j-1, j, j+2$ are distinct and give a second zero triple, a contradiction. Therefore $\Delta(0) \neq 0$; it has one nonzero sign on a smaller interval.

Write the intersection of the two supporting lines as $X = B + sw$. Membership in the line through E_f means $\det(u, X - A) = 0$. Since $B - A = u + v$, bilinearity gives, in separate steps,

$$\det(u, B - A) = \det(u, u) + \det(u, v) = T, \quad (125)$$

$$0 = \det(u, X - A) = T + s\Delta, \quad s = -\frac{T}{\Delta}. \quad (126)$$

At $t = 0$, $X = B$ lies strictly inside $E_f = [A, M]$ and $s = 0$. The line intersection varies continuously because Δ remains nonzero. Consequently its parameter on E_f remains strictly between 0 and 1, and $|s| < 1$ on a smaller interval. It lies in the interior of E_g exactly when $s > 0$. Thus E_f, E_g cross exactly when T and Δ have opposite signs. This also proves that the new crossing is interior and transverse on the stated side.

The newborn crossing in case B. Let $h = \tilde{\lambda}_{j-2}$, so that E_f ends at A and has direction h . Here $\Delta = \det(h, v)$. If $\Delta(0) = 0$, the other endpoint $\mu_{j-2}(0)$ of E_f would also lie on the cusp line, again giving a second zero triple. Hence Δ is nonzero with constant sign on a smaller interval. Write the intersection as $X = A - sh$, so $s > 0$ means moving into E_f from its endpoint A . The relevant determinant is

$$\det(v, A - M) = \det(v, -u) = -\det(v, u) = T. \quad (127)$$

Because $\det(v, h) = -\Delta$, its line equation becomes

$$0 = \det(v, X - M) = T - s \det(v, h) = T + s\Delta, \quad s = -\frac{T}{\Delta}. \quad (128)$$

At $t = 0$, $X = A$ is strictly inside $E_g = [M, B]$. Its parameter there remains interior, and $|s| < 1$ on a smaller interval. The two segments therefore cross exactly when $s > 0$, equivalently when T and Δ have opposite signs. Since T changes sign, the loop and no-loop sides in both cases are distinct and have exactly the signs in (i).

No other crossing pair changes. For any remote pair, strict crossing is determined by the four signs of the endpoint-against-line determinants: each segment must separate the endpoints of the other. This criterion applies on both generic punctured sides. If none of those four triples is $\{j-1, j, j+1\}$, all four signs have the same nonzero limits and hence agree across the wall. Such a crossing pair cannot change.

The only remote pairs whose tests can contain that triple are $\{j-1, j+1\}$ and $\{j-2, j\}$. Indeed the only edges joining two critical vertices are $[A, M]$ and $[M, B]$; adjoining the third critical vertex as an endpoint of a remote edge gives precisely these two pairs. In case A, the unused pair is E_{j-2}, E_j . At the centre the supporting line of E_{j-2} meets the cusp line only at A , since its other endpoint is off that line. But A lies outside $[M, B]$, so the two compact segments are disjoint and remain disjoint nearby. In case B the unused pair is E_{j-1}, E_{j+1} ; the supporting line of E_{j+1} meets the cusp line only at B , which lies outside $[A, M]$. The same compact-disjointness argument applies. This proves the last assertion of (i). It does not assert that the order of thread crossings is unchanged.

The two needle turns. In case A there is $a > 0$ with $v(0) = -au(0)$, because B lies strictly between A and M . The turn determinant at $B = \mu_{j+1}$ satisfies

$$\det(v(0), w(0)) = -a \det(u(0), w(0)) = -a\Delta(0). \quad (129)$$

It is nonzero and retains sign $-\text{sgn } \Delta$ nearby. The other needle corner is M , whose sign is τ_j . By (i), these signs are opposite on the no-loop side and equal on the loop side. In case B there is $a > 0$ with $u(0) = -av(0)$. The turn determinant at $A = \mu_{j-1}$ is therefore

$$\det(h(0), u(0)) = -a \det(h(0), v(0)) = -a\Delta(0). \quad (130)$$

It too retains sign $-\text{sgn } \Delta$, while the other corner is M . Applying (i) gives (ii) also in case B.

Threads and the empty condition. In case A, start at the newborn visit on E_{j-1} and follow the orientation. The traversal runs to M , along E_j to B , and then along E_{j+1} to the other newborn visit. This is the arc in (iii). In case B it runs from the newborn visit on E_{j-2} to A , along E_{j-1} to M , and then along E_j to the other visit. Each description contains exactly the two specified polygon vertices.

No crossing other than the newborn can have both visits on this short arc. The arc is contained in three consecutive edges. The two adjacent pairs have no interior intersection because their incident directions are linearly independent on a generic side. The remaining pair is the newborn pair; its two nonparallel supporting lines meet only once, at the newborn crossing itself. Thus any other crossing visit in the open short arc has its partner in the complementary open arc.

If the newborn visits are cyclically adjacent in the Gauss word, at least one of those two open arcs contains no crossing visit. A crossing visit in the short arc would force a partner in the complementary arc, so would make both nonempty. Hence the short arc is crossing-free. In particular the complete edge between c_1, c_2 has no crossing on the loop side. This middle edge belongs to neither member of the newborn pair. By (i), every crossing pair involving it is unchanged,

so it has no crossing on the no-loop side either. This proves (iii) without assuming that a general cusp is empty.

The rotation sign. Let $\vartheta_i(t) \in (-\pi, \pi)$ be the principal turn at vertex i on a punctured side. The quotient of consecutive unit edge directions is $\exp(i\vartheta_i)$; multiplying around the polygon gives $\exp(i\sum_i \vartheta_i) = 1$. Therefore $\sum_i \vartheta_i$ is an integer multiple of 2π , and $\text{rot}(P) = (2\pi)^{-1} \sum_i \vartheta_i$ is an integer. Every principal turn is continuous away from an antiparallel pair of incident edges. It follows that rotation is constant on each of the two connected punctured sides.

At the centre, every turn determinant except that at j is nonzero: its three vertex indices would otherwise give a second collinear triple. The corresponding principal turns therefore have the same limits from the two sides. At j the nonzero incident vectors are antiparallel. The principal turn tends to $+\pi$ on the side with $T > 0$ and to $-\pi$ on the side with $T < 0$. Set $s = \tau_j(P_{\text{loop}})$. Then the loop-side limit of ϑ_j is $s\pi$ and the no-loop-side limit is $-s\pi$. Cancelling the common limits of all other turns gives

$$2\pi(\text{rot}(P_{\text{loop}}) - \text{rot}(P_{\text{no}})) = s\pi - (-s\pi) = 2s\pi. \quad (131)$$

Dividing by 2π proves (iv). No rotation is assigned to the singular centre, and no differentiability or unthreadedness was used. $\square$

Proof of Lemma 1.34

Proof of 1.34. Clause (i) is the deletion assertion in Lemma 1.30(iii), which includes the fused-edge genericity proof. Its chamber conclusion also follows directly from openness of the generic locus and continuity of deletion: after shrinking, the image of the whole parameter interval, including zero, is connected and generic, so it lies in the connected component containing the central deletion.

For (ii), put $b = a + 1$. By Lemma 1.31 the central parent has nonzero edges and nonzero turns. Relabel cyclically for the count only, so $M = 1$ and $a = k$. The imposed restriction $M \notin \{a - 1, a, a + 1, a + 2\}$ excludes $k = 1, 2, n - 1, n$. Thus $3 \leq k \leq n - 2$. The first half has k vertices and the second has $n - k + 1$, which satisfies the same bounds. This count uses the stated remoteness condition; it does not add a separate regular-path hypothesis.

Each child vertex is a distinct inherited parent vertex. The only zero parent point triple is $\{a, b, M\}$. The first half contains a, M but not b , and the second contains M, b but not a , so neither contains this triple. Every child point determinant is therefore a nonzero parent determinant, proving (G1) for both halves.

Every inherited child edge is a parent edge. The one cut edge in either half is a positive subsegment of E_a . More explicitly, for $M = A + \lambda(B - A)$ with $0 < \lambda < 1$, its two possible directed vectors are λd_a and $(1 - \lambda)d_a$, respectively. A point in the relative interior of either cut edge is in the relative interior of E_a . If three distinct child edge interiors concurred, their parent edge images would be three distinct parent edges with an interior concurrence: each half has only one cut edge, so this edge map is injective. The parent edges must be pairwise remote, because adjacent central parent edges with nonzero turn meet only at their shared endpoint. Such a concurrence contradicts $Z_c = \emptyset$. Hence both halves satisfy (G2), and together with their already proved (G1) this is exactly genericity of Definition 1.5. $\square$

Proof of Lemma 1.35

Proof of 1.35.

The quadratic $q(u, v) = u_x v_y - u_y v_x$ has symmetric matrix, in the order (u_x, u_y, v_x, v_y) ,

$$\frac{1}{2} \begin{pmatrix} 0 & 0 & 0 & 1 \\ 0 & 0 & -1 & 0 \\ 0 & -1 & 0 & 0 \\ 1 & 0 & 0 & 0 \end{pmatrix}. \quad (132)$$

Its rank is four. A nontrivial factorization of a homogeneous quadratic is a product of two homogeneous linear forms. If their coefficient vectors are a, b , the symmetric matrix of their product is $(ab^\top + ba^\top)/2$, of rank at most two. Thus q is irreducible. The vector pairs defining Δ and the auxiliary remote direction determinant $H_{ef} = \det(d_e, d_f)$ are independent linear coordinates: complete either surjective linear map to an invertible coordinate change. The polynomial is then q with extra unused variables. A factorization cannot depend on an unused variable because polynomial degrees in that variable add. This proves irreducibility and nonzeroness of Δ and H .

For T , the six endpoint occurrences are independent. Use the invertible tail/direction coordinates $(\mu_e, d_e), (\mu_f, d_f), (\mu_g, d_g)$. Writing the cross product of the first two line rows as (A, B, C) gives $C = H_{ef}$. Expansion in the third row yields

$$T = d_{g,y}(\mu_{g,x}C - A) + d_{g,x}(B - \mu_{g,y}C). \quad (133)$$

These two coefficients have no common nonconstant divisor. Indeed any irreducible common divisor is independent of $\mu_{g,x}$ because the second coefficient is, and independent of $\mu_{g,y}$ because the first is. Comparing coefficients would make it divide A, B, C . But $C = H_{ef}$ is irreducible and does not divide both A, B : take $d_e = d_f = (1, 0)$ and tails on different horizontal lines. Then $C = 0$ while $A \neq 0$. Therefore the two coefficients in (133) are coprime. In a factorization of a polynomial homogeneous of degree one in $(d_{g,x}, d_{g,y})$, one factor is independent of both variables and divides both coefficients. It must be a unit. This proves irreducibility of T ; the same expansion proves it is nonzero.

Each Δ depends on every vertex in its three-element support. Each T depends on every endpoint in its six-element support: with one endpoint fixed, the line through it can be varied by moving its other endpoint; choose the other two lines to intersect away from the fixed endpoint to make the determinant vary. Thus different supports cannot give associates, and the three- and six-vertex families cannot give associates. For a fixed six-vertex support, the three pairwise remote cycle edges form a perfect matching. A proper subset of a cycle induces disjoint paths, each of which has at most one perfect matching (remove its first edge recursively). The only possible two matchings occur when $n = 6$ and the support is the whole cycle; they are the two alternating matchings. At the specialization

$$(\mu_1, \dots, \mu_6) = ((1, 0), (2, 0), (0, 1), (0, 2), (1, 1), (2, 2)) \quad (134)$$

the tail-indexed values are $T_{135} = 0$ and $T_{246} = 2$. Hence they too are nonassociate. One representative per unordered point triple already removes the sign associates among the Δ 's.

Every scalar endpoint coordinate occurs affinely in its edge line row, and the three rows of T have disjoint endpoint supports. Hence T is affine in each scalar coordinate. The same assertion for Δ follows by expanding its area determinant. Genuine dependence therefore has a nonzero slope a . Irreducibility of $ax_j + b$ implies that a, b are coprime in the polynomial ring of the other coordinates: a nonconstant common divisor would factor F . Likewise c, d are coprime. If $ad - bc = 0$, coprimality of a, b in this unique factorization domain implies $a \mid c$; put $c = au$. The identity then gives $d = bu$. Coprimality of c, d forces u to be a unit, so $G = uF$, contradicting nonassociation. Thus $ad - bc$ is a nonzero polynomial. After specialization, a nonzero slope makes a linear root simple, and $cF - aG = cb - ad$ excludes a common root when $ad - bc \neq 0$. $\square$

Proof of Theorem 1.36

Proof of 1.36. Open neighbourhoods and endpoint collars. The regular locus is open: nonzero edge lengths and avoidance of antiparallel direction pairs are open conditions. The generic locus is open as well. To see this directly under SM (G1),(G2), finitely many nonzero point determinants retain their signs near a generic tuple. Every central remote pair is either compact-disjoint or an interior transverse pair (Lemma 1.6(v)); disjointness persists, and a transverse interior crossing persists with continuous parameters, by Lemma 1.9. Distinct crossing points have positive pairwise separation, since a coincidence would place at least three edge interiors at one point. Thus no triple interior concurrence appears in a sufficiently small neighbourhood. This proves openness even when a concurrence polynomial vanishes at an inactive configuration, such as three line extensions concurrent outside the segments. No all-polynomials-nonzero endpoint assumption is needed.

Cover the compact image of γ by axis-parallel open cubes with closures in $\mathcal{R}_n$ and diameters less than δ . The endpoint cubes may be chosen inside the generic locus. A sufficiently fine subdivision $0 = t_0 < \dots < t_N = 1$ has each image $\gamma([t_{i-1}, t_i])$ inside one such cube C_i , with the first and last cubes generic. Consecutive cubes overlap in a nonempty open set because they both contain $\gamma(t_i)$. Refine if necessary so $N \geq 3$.

Choose independent internal waypoints $z_i \in C_i \cap C_{i+1}$, while $z_0 = \gamma(0)$ and $z_N = \gamma(1)$ remain fixed. Connect z_{i-1} to z_i by changing their $2n$ scalar coordinates one at a time in a fixed order. Every intermediate hybrid belongs to C_i , since an axis-parallel cube is closed under coordinate hybridization. Parameterize these legs over $[t_{i-1}, t_i]$. The old and new points at each time belong to the same cube, proving the uniform δ bound. All legs are regular. The first and last coordinate polylines are entirely generic; they are endpoint collars. Only legs joining two independently variable internal waypoints need further conditions.

One joint algebraic choice for the central legs. The product of the open waypoint sets is an open subset of a finite-dimensional real coordinate space. For each central leg moving scalar coordinate x_j , impose the following nonvanishing requirements:

1. Every $F \in \mathcal{F}$ is nonzero at both endpoint hybrids.
2. For every x_j -dependent $F = ax_j + b$, the fixed-coordinate slope a is nonzero; for every distinct such pair F, G , its fixed-coordinate linear resultant $ad - bc$ is nonzero.
3. The scalar step is nonzero, and both coordinate differences of every pair of vertex occurrences are nonzero at the endpoint hybrids.

Every requirement is the nonvanishing of a nonzero polynomial in the joint waypoint variables. For the first two assertions, a central hybrid selects each coordinate from one of two independent internal waypoints; its projection contains an open coordinate box. A nonzero polynomial in the hybrid coordinates, or in the fixed coordinates of the leg, therefore cannot pull back to the zero polynomial. The nonzero slopes and resultants are supplied by Lemma 1.35. Coordinate-difference polynomials use distinct occurrence variables; the scalar step uses two independent values of the moving coordinate, so these polynomials also are nonzero.

There are finitely many conditions. Their product is a nonzero polynomial, which cannot vanish on a nonempty real open set. For completeness, this last fact follows by induction on the number of variables: a univariate nonzero polynomial has finitely many roots; on a product of open intervals, fixing all but one variable would force every coefficient polynomial to vanish, and the induction hypothesis then forces the original polynomial to be zero. Thus one joint waypoint choice satisfies all requirements. Fixed original endpoints are excluded from this choice, which is why collars were separated.

Finite simple roots and collision avoidance. On a central leg a factor independent of x_j stays at its nonzero endpoint value. A dependent factor is a genuinely linear polynomial in the moving coordinate with nonzero slope, hence has at most one root on the leg. Different factors have no common root by their nonzero linear resultant. No root is at a leg endpoint, and each root changes the sign of exactly one named factor because the scalar step is nonzero. Consequently all central roots are finite, isolated, smooth hypersurface crossings. Tangencies and intersections of distinct control hypersurfaces have been avoided by the explicit slope and resultant conditions, not by an unproved generic-position assertion.

On a leg moving the x -coordinate of one vertex, its y -difference from every other vertex is fixed and nonzero. It cannot collide with any of them. All other vertex pairs are fixed and distinct. The same argument exchanges x, y on a y -coordinate leg. The collars are generic and hence have distinct vertices. Thus the whole resulting path is collision-free, including its fixed endpoints. Every physical edge occurrence stays nonzero because all legs lie in their regular cubes. A chosen labelled edge can therefore be followed along the path without relabelling or deletion; no root-independence statement is used.

If all members of $\mathcal{F}$ are nonzero, (G1) holds. Any triple interior concurrence would, under (G1), use pairwise remote edges and make one T zero, so (G2) holds too. Hence every nongeneric parameter is among the finite central roots. Some roots of an inactive T may still be generic; they are not declared wall events. Collars contain no nongeneric parameters even when an inactive polynomial there has nonisolated zeros.

A point-determinant root. At a root of Δ_K , exactly the point triple K is collinear, all vertices remain distinct, and all T 's are nonzero. Thus $Z_{\text{pt}} = \{K\}$ and $Z_c = \emptyset$. The determinant changes sign simply on the leg. Count the cyclically adjacent pairs in K . If there are at least two and $n \geq 4$, $K = \{j-1, j, j+1\}$ is a consecutive triple. Its middle vertex must lie strictly between its neighbours, because otherwise the two consecutive directions there are antiparallel, contrary to regularity. This is exactly (F). When $n = 3$, a distinct collinear triple forms a closed collinear triangle and necessarily has an antiparallel consecutive pair; such a root cannot lie in the regular cube.

If there is exactly one adjacent pair, write it as $a, b = a+1$ and call the third index M . The lack of a second adjacent pair is precisely $M \notin \{a-1, a, a+1, a+2\}$. The contact point is not either endpoint, by collision freedom. If M is strictly inside $[\mu_a, \mu_b]$, this is (V). Both neighbours of M have nonzero heights over that line, since another zero height would be another point triple. Equal height signs give the bigon branch and opposite signs the sliding branch. If M lies outside the closed segment, this is (E). Finally, no adjacent pair gives exactly the pairwise nonadjacent point triple of (C). These alternatives are exhaustive and disjoint.

An active concurrence root. At a root of a single T_{efg} all point determinants are nonzero. If it is nongeneric, some three edge interiors concur; they must be pairwise remote, and their concurrence polynomial must be this unique zero member. Thus $Z_{\text{pt}} = \emptyset$ and $Z_c = \{\{e, f, g\}\}$. At their common interior point every pair is transverse, since coincident remote supporting lines would force a point-determinant zero. All three crossing pairs persist on nearby punctured sides by their continuous interior intersection parameters.

To verify the sign-change part of (T), put $H_{ef} = \det(d_e, d_f)$ and write $q_{ef} = \mu_e + t_{ef}d_e$. The cross product of the first two line rows in (6) equals $H_{ef}(q_{ef,x}, q_{ef,y}, 1)$, as follows from its third coordinate H_{ef} and the two line equations. Therefore

$$T_{efg} = H_{ef} \det(d_g, q_{ef} - \mu_g) = -(t_{ef} - t_{eg})H_{ef}H_{eg}. \quad (135)$$

The second equality subtracts $q_{eg} = \mu_e + t_{eg}d_e$, which lies on the g line, and uses $\det(d_g, d_e) = -H_{eg}$. The two H factors are nonzero with constant sign near the root. The simple sign change of T

therefore changes the sign of $t_{ef} - t_{eg}$. Apply the same calculation after cyclic permutation of (e, f, g) to obtain the other two required parameter-difference sign changes. This is exactly (T).

Every nongeneric parameter has now been classified as one named simple wall. There are no cusp walls, because their central antiparallel pair is excluded throughout the regular cubes. Both the endpoints and every occurrence label have been kept exactly as required. $\square$

Proofs for Section 2

Proof of Lemma 2.9

Proof of 2.9. For (i), the one-part composition in coordinate $[i, j]$ contributes X_{ij} with coefficient one. Every other composition has at least two parts, all of shorter interval length. Given any target Y , first put $X_{i,i+1} = Y_{i,i+1}$ and then, in increasing interval length, subtract the already determined nonunary sum from Y_{ij} . This uniquely solves $\Phi_{D,H}(X) = Y$, by polynomial operations in the finite arrays. Writing the solution map as Ψ , construction gives $\Phi_{D,H} \circ \Psi = \text{id}$. Applying uniqueness to the target $\Phi_{D,H}(X)$ gives $\Psi \circ \Phi_{D,H}(X) = X$, the other inverse.

Expand $F_H(G_D(X))_{ij}$. Its terms are indexed by an outer composition of $[i, j]$ and an inner composition of each outer part. Their ordered union is a refined composition $\rho = (r_0 < \dots < r_s)$, with the subset of its interior cuts marked as outer cuts. Conversely, cut ρ at the marked cuts to recover the outer parts and restrict ρ to each part to recover its inner composition. This is a bijection, including empty and full marked subsets and unary inner or outer compositions.

At a marked cut r_k the coefficient is $-H(i, r_k, j)/2$. At an unmarked cut its immediate neighbors in ρ are precisely its neighbors inside the same inner part, since that part's endpoints are also in ρ . Its coefficient is $D(r_{k-1}, r_k, r_{k+1})/2$. Each interior cut supplies one factor, and the product of child coordinates is exactly the refined product. Summing over every subset of marked cuts independently chooses the near or negative far summand at each cut. Distributivity yields the product of their sums, precisely the coefficient of ρ in Definition 2.8. This proves (7), with all neighboring near gates accounted for.

For (ii), the identities $d\Theta(-hd) = (d - h)/2$ and $d\Theta(hd) = (d + h)/2$, for $d, h \in \{\pm 1\}$, turn the open-sum recursion into $\Phi_{D,H}(b) = E$ and the full root sum into $\Phi_{D,-H}(b)_{0N}$. In particular its unary root term is retained. By (7), $F_H(G_D(b)) = E$; uniqueness of the inverse gives $G_D(b) = c$. Applying the factorization with $-H$ gives

$$\Phi_{D,-H}(b) = F_{-H}(G_D(b)) = F_{-H}(c). \quad (136)$$

This proves the assertion about the complete output. It asserts no near-array independence for individual trees or open sums.

For (iii), restriction to J commutes with all coordinate equations: each equation on a subinterval of J reads only its own subintervals and triples. Triangular uniqueness therefore identifies the restriction of c with the inverse constructed directly on that restricted word. The equation $F_H(c) = E$ gives the claimed $\mathcal{E}$ values. On a one-leaf interval the only composition is unary, its inverse coordinate is one, and both transforms equal one. This is a formal open-word convention, not a two-gon amplitude. On an interval with at least two leaves, the closed endpoint polygon has at least three vertices. Cutting its closing edge from a_j to a_i gives exactly the word $a_i, \dots, a_j$. If it satisfies (G1), part (ii) on that word identifies $\mathcal{B}_J$ with its amplitude at that specified physical root. The auxiliary coordinates c may be rational and their binary coordinates need not vanish. $\square$

Proof of Theorem 2.10

Proof of 2.10. The wall hypothesis makes every distinct-triple determinant except the critical one nonzero at zero; after a common shrink its sign is the same on both punctured sides. Neither gap contains the entire critical triple, so both gap arrays and every nontrivial closed gap polygon are nondegenerate and unchanged. Singleton gap values have only their formal open-word meaning. At least one epsilon is positive: otherwise both summands in $\xi_z - \xi_x = (\xi_y - \xi_x) + (\xi_z - \xi_y)$ would have sign opposite to their nonzero sum.

Use the far-only representation of Lemma 2.9 on each side. The changing triple appears as far data in exactly one coordinate formula of F_H , namely $[x, z]$. Induction in interval length therefore shows that c is unchanged on every interval not containing I_* .

In the $[x, z]$ equation every proper child is unchanged. A composition has a changing coefficient exactly when it cuts at y . Splitting its cut list at y is a bijection with a pair of arbitrary compositions of L and R ; concatenating the two lists is its inverse. Its one changing far gate $-H(x, y, z)/2$ has difference $-\delta$. For every other cut in these compositions, the wall geometry gives

$$H(x, r, z) = \epsilon_L H(x, r, y) \quad (x < r < y), \quad (137)$$

$$H(x, r, z) = \epsilon_R H(y, r, z) \quad (y < r < z). \quad (138)$$

For the first identity, the vectors from a_x to a_z and to a_y differ by the scalar $(\xi_z - \xi_x)/(\xi_y - \xi_x)$. Taking their determinants with $a_r - a_x$ gives the indicated sign ratio. For the second use the vectors from a_z to a_x and to a_y ; their scalar ratio has sign ϵ_R . All these determinants are nonzero at the wall. Their signs on both sufficiently close sides are therefore the wall signs, proving both displayed identities on those sides without evaluating a zero gate.

Consequently the composition sum on a gap is $F_H(c)_J = \mathcal{E}_J$ if $\epsilon_J = 1$, and $F_{-H}(c)_J = \mathcal{B}_J$ if $\epsilon_J = -1$. The changed nonunary sum in $F_H(c)_{xz} = E_{xz} = 0$ is thus $-\delta \mathcal{U}_L \mathcal{U}_R$. Since the unary coefficient is 1, moving this sum to the other side gives

$$c_{xz}^+ - c_{xz}^- = s_*, \quad s_* = \delta \mathcal{U}_L \mathcal{U}_R. \quad (139)$$

The nonunary part of $F_{-H}(c)_{xz}$ must be calculated separately. Its changing gate is $+H(x, y, z)/2$, with difference $+\delta$; all proper children are still unchanged. Reversing the remaining far signs exchanges $\mathcal{E}_J$ with $\mathcal{B}_J$ in each of the two gap sums. Its direct coefficient change is therefore $+\delta \mathcal{V}_L \mathcal{V}_R$.

It remains to propagate (139) through larger intervals. Contract I_* in the boundary word to the formal leaf e_* , and let c^Q be the inverse $F_{H^Q}^{-1}(E)$ on that contracted word. This inverse is defined even when the contracted word has only one leaf; that is an open-word convention, not a two-gon amplitude. For an interval containing I_* write I/e_* for its contraction. We prove by interval length that

$$c_I^+ - c_I^- = s_* c_{I/e_*}^Q \quad (I \supseteq I_*), \quad c_I^+ - c_I^- = 0 \quad (I \not\supseteq I_*). \quad (140)$$

The smallest containing interval is I_* itself. Its contracted coordinate is the leaf value 1, so this is (139). For a strictly larger I , its far coefficient formula is unchanged. In a top composition at most one child can contain the positive-length interval I_* , because distinct children have disjoint interiors. If no child contains it, every child coordinate is unchanged. Otherwise the product difference is exactly the difference of that unique child times all other, unchanged child coordinates. There is no product of two changing factors or choice of side for the others.

The compositions having such a child are in bijection with compositions of I/e_* . Forward, erase the interior vertices of I_* inside that one child; no top cut lies there. Backward, the distinguished leaf e_* belongs to one unique contracted child, and expanding it recovers the original child and top cut list. The operations are inverse. All top endpoints and cut vertices survive, so every far gate has the same indices and signs. Every other child is disjoint from the interior of I_* , and its c coordinate agrees with c^Q by restriction and uniqueness of the unchanged triangular equations. Insert the induction hypothesis on the changing child. The difference of the equation $F_H(c)_I = 0$ becomes s_* times the equation $F_{H^Q}(c^Q)_{I/e_*} = 0$. The contracted interval has at least two leaves when I strictly contains I_* , so this latter right-hand side is indeed zero. Equivalently, its unary term is the negative of its full nonunary sum, giving precisely $c_I^+ - c_I^- = s_* c_{I/e_*}^Q$. This proves the induction with its sign.

When $I_* \neq F$, the same composition bijection, now including the unary term, applies to the barred transform F_{-H} on F . Its coefficients are unchanged because its endpoints are not both x, z . Substitution of (140) gives $s_* F_{-HQ}(c^Q)_{F/e_*}$. At least one further leaf survives outside I_* , so Q has at least three vertices. Its vertices have no zero triple, since the only critical triple lost its middle occurrence. Its endpoints and hence physical root survive. Lemma 2.9 identifies the last output with $s_* A_g(Q)$, proving (12).

When $I_* = F$, all proper children are unchanged and the unary contribution is (139). Add the separately computed direct barred coefficient change $+\delta \mathcal{V}_L \mathcal{V}_R$ to obtain (13). In this case there is no closed contracted amplitude of arity at least three. $\square$

Proof of Theorem 2.12

Proof of 2.12. Put $Q = P(0) \setminus j$. Its vertices are a subset of the wall vertices missing the middle occurrence of the sole zero triple. Hence all its distinct-triple determinants are nonzero, so it satisfies (G1), and $n - 1 \geq 3$. The flat hypothesis puts the deleted vertex strictly between its neighbors. Let $P^\pm$ denote the sides on which $\chi(\mu_{j-1}, \mu_j, \mu_{j+1}) = \pm 1$, and write $u = (\mu_{j-1}, \mu_j)$, $v = (\mu_j, \mu_{j+1})$, $w = (\mu_{j-1}, \mu_{j+1})$. The determinant defining this chirotope is $\det(\tilde{\lambda}_{j-1}, \tilde{\lambda}_j)$, since $\mu_{j+1} - \mu_{j-1} = \tilde{\lambda}_{j-1} + \tilde{\lambda}_j$. Thus it is τ_j , and $P^- = P_{\text{right}}$, $P^+ = P_{\text{left}}$.

Write $A = \mu_{j-1}$, $M = \mu_j$, $B = \mu_{j+1}$. In any root cut the three critical occurrences have a cyclic permutation of the order A, M, B . A cyclic permutation is even, so their reverse-ordered far sign is always $-\chi(A, M, B)$. Thus the δ of (9), for P^+ minus P^- , is -1 .

If g is neither u nor v , the cut order is A, M, B , with no root seam between these consecutive edges. Both gaps are singleton leaves, and strict betweenness gives $(\epsilon_L, \epsilon_R) = (1, 1)$. Hence $\mathcal{U}_L \mathcal{U}_R = 1$. Their span is proper: if both extreme vertices were the root endpoints, the remaining root edge would close a three-vertex polygon, contrary to $n \geq 4$. Contracting the two-edge span deletes exactly M and retains g . Formula (12) gives $A_g(P^+) - A_g(P^-) = -A_g(Q)$.

If $g = u$, the cut order is $M, B, \dots, A$, so the critical span is full. The gap $L = (M, B)$ is singleton and $(\epsilon_L, \epsilon_R) = (-1, 1)$: with $A = 0$, $B = 1$, $M = t \in (0, 1)$, the first ratio has numerator $1 - t > 0$ and denominator $-t < 0$, and the second has numerator $-1 < 0$. The other gap R has $n - 2 \geq 2$ leaves; its closing root is $w = (A, B)$ and its closed polygon is Q . Therefore $\mathcal{U}_L \mathcal{U}_R = 0$ and $\mathcal{V}_L \mathcal{V}_R = A_w(Q)$. Formula (13) gives the desired negative of $A_w(Q)$.

If $g = v$, the cut order is $B, \dots, A, M$. Now L is the nontrivial deletion word with closing root w , $R = (A, M)$ is singleton, and $(\epsilon_L, \epsilon_R) = (1, -1)$. Indeed their numerators are -1 and t with common denominator $t - 1 < 0$. Again $\mathcal{U}_L \mathcal{U}_R = 0$ and $\mathcal{V}_L \mathcal{V}_R = A_w(Q)$. These three cases exhaust the physical roots and prove (14) with precisely Definition 2.11. $\square$

Proof of Theorem 2.13

Proof of 2.13. Let $u = (\mu_{j-1}, \mu_j)$, $v = (\mu_j, \mu_{j+1})$ and $w = (\mu_{j-1}, \mu_{j+1})$. These are physical edges; they are not renumbered root slots. The cusp hypothesis places the deleted vertex outside its neighbors' closed segment. As in the preceding theorem, $\chi(\mu_{j-1}, \mu_j, \mu_{j+1}) = \tau_j$ by bilinearity.

First prove the signed consecutive identity directly, without assuming the betweenness calculation from S3. Write $A = \mu_{j-1}$, $M = \mu_j$, $B = \mu_{j+1}$, and take line coordinates $A = 0$, $B = 1$, $M = t$, where $t < 0$ or $t > 1$. Name sides by $\chi(A, M, B) = -1, +1$ as in S3. The reverse-ordered far sign again gives $\delta = -1$ for P^+ minus P^- . For a nonincident root the two gaps are singletons, regardless of their epsilon signs; hence $\mathcal{U}_L \mathcal{U}_R = 1$. The proper-span formula gives the signed deletion identity.

For the incident roots the following full-span table computes all four cases from (10):

root	critical cut order	t	ϵ_L	ϵ_R	possibly nonzero summand
u	M, B, A	$t > 1$	1	1	$\mathcal{V}_L \mathcal{V}_R = \mathcal{B}_R$
u	M, B, A	$t < 0$	1	-1	$\mathcal{U}_L \mathcal{U}_R = \mathcal{B}_R$
v	B, A, M	$t > 1$	-1	1	$\mathcal{U}_L \mathcal{U}_R = \mathcal{B}_L$
v	B, A, M	$t < 0$	1	1	$\mathcal{V}_L \mathcal{V}_R = \mathcal{B}_L$

(141)

For u , L is singleton and R is the deletion word closed by $w = (A, B)$; the epsilon ratios are $(1 - t)/(-t)$ and $(-1)/(-t)$. For v , R is singleton and L is that deletion word; the ratios are $(-1)/(t - 1)$ and $t/(t - 1)$. These four ratio tests give the table. The other summand contains the nontrivial deletion word's $E = 0$. Formula (13) therefore gives $A_g(P^+) - A_g(P^-) = -A_{D_j(g)}(Q)$ in each incident case also.

For the rotation sign, the principal local turn at M tends to $+\pi$ on the positive-chirotope side and to $-\pi$ on the negative side, since its determinant is $\det(M - A, B - M)$ and has sign $\chi(A, M, B)$. Every other local turn is continuous at this sole consecutive degeneration: its consecutive triple is distinct from the critical triple because $n \geq 4$, and is nonzero at the wall. The edges are nonzero and the two incident directions at M tend to opposites. Thus the sum of all principal turns on the positive side minus that on the negative side tends to 2π . By Lemma 1.20, each punctured-side rotation is an integer constant on its sufficiently small connected regular side. Their difference is therefore its limit, namely 1. If the loop side is P^+ , then $\kappa = 1$ and the just-proved signed formula is (15). If it is P^- , then $\kappa = -1$ and reversing the same formula gives (15). No root-independence or transport theorem is used. $\square$

Proof of Theorem 2.14

Proof of 2.14. Write $A = \mu_a(0)$, $B = \mu_{a+1}(0)$, $X = \mu_M(0)$, $d = (A, B)$, $d_1 = (A, X)$ and $d_2 = (X, B)$. The original cyclic order of the three critical vertices is A, B, X . Both original arc words, from X to A and from B to X , have at least two leaves: their closing halves are generic and have at least three vertices by Lemma 1.34(ii). Each nontrivial closed word below has only nonzero point-triple determinants. Its amplitude at the indicated closing edge is therefore the interval value of Lemma 2.9(iii).

In any root cut, the critical order is a cyclic permutation of A, B, X . Such a permutation is even, whereas the reverse ordering in the far sign is odd. The critical far sign is consequently $-\chi(A, B, X)$. It equals $-s$ on P_- and s on P_+ . The quantity δ in (9) is therefore s . Choose affine line coordinates $A = 0$, $B = 1$, $X = t$ with $0 < t < 1$.

For $g = d$, the boundary word begins at B , passes through X , and ends at A . The critical span is full. The two ratios in (10) are $(t - 1)/(-1) > 0$ and $(-t)/(-1) > 0$. Both gaps are nontrivial, so both $\mathcal{U}$ factors are zero $\mathcal{E}$ factors. The $\mathcal{V}$ factors are respectively $A_{d_2}(\lambda_2)$ and $A_{d_1}(\lambda_1)$. The full-span formula (13) gives the desired product. Writing it in the prescribed half order uses commutativity of scalar multiplication; it does not interchange the named halves.

If g lies on the original arc from X to A , the critical cut order is A, B, X . The first gap is the singleton edge d , and the second is the nontrivial arc from B to X . Its epsilon ratios are $1/t > 0$ and $(t - 1)/t < 0$, so $\mathcal{U}_L \mathcal{U}_R = A_{d_2}(\lambda_2)$. The critical span is proper. Otherwise the physical root would be an edge from X to A , making that original arc a singleton and its closing half a two-gon, contrary to the arity just established. Contracting the span replaces the arc $A, B, \dots, X$ by d_1 . The surviving polygon is exactly λ_1 with the same physical root g . Formula (12) gives $sA_g(\lambda_1)A_{d_2}(\lambda_2)$.

If g lies on the original arc from B to X , the critical cut order is X, A, B . The first gap is the nontrivial arc from X to A , and the second is the singleton d . The ratios are $-t/(1 - t) < 0$ and

$1/(1-t) > 0$. Thus $\mathcal{U}_L \mathcal{U}_R = A_{d_1}(\lambda_1)$. The span is proper, since a full span would make the original B -to- X arc a singleton. Contraction replaces $X, \dots, A, B$ by d_2 , leaving exactly λ_2 at the same physical root. The proper-span formula gives $sA_{d_1}(\lambda_1)A_g(\lambda_2)$.

Every original root belongs to exactly one of these three cases. The inherited-arc cases include the edges incident to the named endpoints, so no additional root-seam case remains. No step restricts the side of the contact line occupied by either neighbour of X . Both bigon and sliding branches are therefore covered. The three products are precisely the three rows of the stated half-root map. $\square$

Proof of Theorem 2.15

Proof of 2.15. For (i), every point-triple determinant has a nonzero central value, so finitely many continuity conditions give a common interval on which all signs read by Definition 2.2 agree across the wall. Every gate and composition weight agrees. Starting with the one-leaf value and inducting on interval length in the open recursion then gives equality of all b 's. Substituting them in the finite root sum gives equality of A_g . Crossing-visit order is not an input to any of these formulas.

For an exterior-extension wall, use A, B, X, d as in the preceding proof, but now $t < 0$ or $t > 1$. By the remoteness condition in Definition 1.29(E), both full original arcs from X to A and from B to X have at least two leaves. A singleton in either would give a second adjacent pair in the critical triple and would be a consecutive-triple wall instead. Closed nontrivial gap and contracted words have (G1), since they omit at least one member of the sole zero triple.

If $g = d$, the critical order is B, X, A , the span is full and both gaps are nontrivial. Their epsilon ratios are $1 - t$ and t . Exactly one is negative. Hence $\mathcal{U}_L \mathcal{U}_R$ and $\mathcal{V}_L \mathcal{V}_R$ each contain a nontrivial $\mathcal{E}$ factor, which is zero by Lemma 2.9(iii). The full-span formula gives zero jump. If g lies on the original X -to- A arc, the cut order is A, B, X ; L is singleton, R is nontrivial, and $\epsilon_R = \text{sgn}((t-1)/t) = 1$ for both allowed ranges of t . Thus $\mathcal{U}_R = \mathcal{E}_R = 0$. The span is proper, since a full span would make the original X -to- A arc a singleton. Formula (12) gives zero jump. For a root on the other arc, the cut order is X, A, B ; L is nontrivial and $\epsilon_L = \text{sgn}(-t/(1-t)) = 1$. The span is proper for the analogous endpoint reason, and $\mathcal{U}_L = \mathcal{E}_L = 0$. This exhausts all exterior roots. No assertion here applies to a consecutive cusp or flat wall.

Finally consider a pure cut. In the fixed root word write its critical positions as $x < y < z$, with $N = n - 1$. Pairwise nonadjacency gives

$$y - x \geq 2, \quad z - y \geq 2, \quad N + 1 + x - z \geq 2. \quad (142)$$

The first two inequalities make both gaps nontrivial. The third shows that the span cannot be full: full-span extremes would be the two adjacent endpoints of the physical root. At least one epsilon is positive, since the two nonzero line displacements from x to y and from y to z cannot both have sign opposite to their sum from x to z . Therefore the source product $\mathcal{U}_L \mathcal{U}_R$ contains a zero $\mathcal{E}_L$ or $\mathcal{E}_R$. Formula (12) again gives zero. The argument allows $x = 0$ or $z = N$ separately, so it includes roots incident to a critical vertex and all six possible line orders. The conclusion concerns the complete root output; no constancy of the original open sums at a silent point-triple wall is asserted. $\square$

Proof of Theorem 2.16

Proof of 2.16. First, the weakly generic locus $\mathcal{W}_n$ is open in the labelled coordinate space, by Lemma 1.16; in particular no new crossing or triple concurrence appears in a small neighbourhood of a weakly generic tuple, including at exterior collinearities.

The generic locus $\mathcal{U}_n$ is dense in the open set $\mathcal{W}_n$ (Lemma 1.17), so generic points occur in every open ball contained in $\mathcal{W}_n$.

Every visible chamber, a connected component of the open set $\mathcal{W}_n$, is polygonally connected, by Lemma 1.16.

Take two generic points in one visible chamber V . Join them by a path γ in V . The compact image has a positive distance from the closed complement of $\mathcal{W}_n$; choose a smaller positive error tolerance. Since nonzero turns imply regularity, Theorem 1.36 applies and gives an endpoint-fixed path within that tolerance, hence still in $\mathcal{W}_n$. Its possible actual wall events are F,V,T,E,C. A flat event would have a zero turn, a vertex-edge event a vertex on a nonincident closed segment, and a triple event a forbidden interior concurrence. All three are excluded by membership in $\mathcal{W}_n$. Only E and C remain. Along each generic segment A_g is constant by Proposition 2.7; at each of the finitely many remaining events its jump is zero by Theorem 2.15(ii). Thus the two generic endpoint values agree.

Density ensures that every visible chamber meets the generic locus. Define $\tilde{A}_g$ on that chamber to be the common generic value just proved. This is a well-defined function, constant on visible chambers and equal to A_g on $\mathcal{U}_n$. Any other such function must have the same value on each chamber because that chamber contains a generic point. This proves uniqueness. The original finite tree formula on $\mathcal{U}_n$ is the defining one; no evaluation of a gate at zero or polynomial-form assertion at simultaneous silent zeros is used. $\square$

Proof of Lemma 2.18

Proof of 2.18. The boundary word contains every vertex label exactly once. Consequently unordered triples of boundary positions correspond bijectively to unordered triples of vertex labels; alternating the order changes only the sign of the canonical variable.

Consider a vertex with composition $r_0 < \dots < r_s$. There are no cut factors when $s = 1$. When $s = 2$, the ordered triples defining d_1 and h_1 are both $(a_{r_2}, a_{r_1}, a_{r_0})$. Thus $d_1 = h_1$ as formal variables, and the ordinary and root factors are respectively zero and d_1 .

For $s \geq 3$, the unordered near triple at cut k is $\{r_{k-1}, r_k, r_{k+1}\}$ and the far triple is $\{r_0, r_k, r_s\}$. Two near triples are equal only if their smallest positions agree, which forces their cut indices to agree. Two far triples are equal only if their interior positions agree, again forcing the same cut. Equality of a near triple and a far triple would require $r_{k-1} = r_0$ and $r_{k+1} = r_s$ by comparison of their smallest and largest positions. These equalities force $k = 1$ and $s = 2$, contrary to $s \geq 3$. Hence distinct cut factors at one vertex use disjoint sets of canonical variables.

Now consider distinct vertices that both carry cut factors. If neither is an ancestor of the other, their intervals lie in different child intervals of their least common ancestor. Those child intervals have disjoint interiors and at most one common boundary position. Thus the two vertices cannot share a triple of boundary positions.

If v is an ancestor of w , the interval of w lies in one child interval of v . Each child interval of v contains exactly two cut positions of v , namely its endpoints. Every near or far triple at v contains three distinct cut positions of v , so it cannot lie in that child interval and cannot occur at w .

For completeness, a vertex with at least two children has every child interval strictly smaller than its own. All ordinary internal vertices have at least two children by Lemma 2.5. Therefore equality of the intervals of distinct internal vertices is possible only for the unary root and its sole child. The root then has no cut factor, so this equality creates no repeated variable.

A tree with an ordinary binary vertex has zero formal weight. In every other tree, the cut factors are linear and distinct factors use disjoint sets of canonical variables. Every monomial in their product therefore has degree at most one in each variable. The sign $(-1)^{N_+(T)}$ does not change

this property, and neither does summing over the finite set of trees. This proves multi-affinity in the polynomial ring, without imposing any relations on its variables.

Finally, on (G1)-polygons Lemma 2.3 identifies every linear cut factor with the original gate. Applying this identity term by term in Lemma 2.5 proves $A_g(\chi(P)) = A_g(P)$. $\square$

Proof of Lemma 2.19

Proof of 2.19. The vertices of a weakly generic polygon are distinct. Coincident adjacent vertices would give a zero edge; coincident nonadjacent vertices would put a vertex on a nonincident closed edge. An original edge whose endpoints are among the selected line points cannot contain any other selected point in its interior. Also two successive gaps of one leaf cannot both join selected line points: they would be consecutive collinear original edges, with zero turn. These are consequences of Definition 1.15.

Choose the line coordinate so $t_0 = 0$ and $t_k = 1$. Suppose every increasing gap had one leaf. Let t_ℓ be the largest selected coordinate; it is not t_0 , so $\ell \geq 1$. Its incoming gap is increasing, hence is an original edge. Since that edge contains no other selected point, $t_{\ell-1}$ is the second largest selected coordinate. We cannot have $\ell = 1$: then this second largest coordinate would be $t_0 = 0$, whereas the distinct final point has coordinate 1 below the maximum. The maximum cannot itself be that final point when $\ell = 1$, since $k \geq 2$. Thus $\ell \geq 2$. The preceding coordinate $t_{\ell-2}$ is smaller than $t_{\ell-1}$, because only t_ℓ is larger and the selected points are distinct. That preceding gap is therefore also an increasing original edge. The two successive one-leaf gaps give the forbidden zero turn. This proves the first assertion.

For the second assertion close the selected list by the physical root edge from a_N to a_0 . Its line direction is negative. If every negative gap in the open list had one leaf, every negative step in this cyclic list would be an original edge. Start at its largest coordinate. The following point is lower, and the edge to it cannot contain another selected point, so it is the second largest coordinate. There are at least three selected points. The next point in cyclic order is therefore still lower: it is neither of those two points. The next step is another negative original edge, again giving two successive collinear original edges and a zero turn. This contradiction proves that one negative open gap is nontrivial. $\square$

Proof of Lemma 2.20

Proof of 2.20. The polynomial inverses exist by Lemma 2.9(i). We first establish a coefficient identity with the child coordinates fixed to c^0 . Fix $I = [i, j]$ and let $Z_I = \{r : i < r < j, H^0(i, r, j) = 0\}$. Only the entries indexed by $r \in Z_I$ vary in this top-coordinate formula. Write those entries as u_r . For a nonempty subset $S = \{r_1 < \dots < r_{k-1}\} \subseteq Z_I$, set $r_0 = i$, $r_k = j$ and $J_\ell = [r_{\ell-1}, r_\ell]$. All these endpoints lie on the line through a_i, a_j ; define ϵ_ℓ as in (17). For $\eta \in \{1, -1\}$ the coefficient of $\prod_{r \in S} u_r$ in the top formula is

$$\left[\prod_{r \in S} u_r \right] F_{\eta H}(c^0)_{ij} = \left(-\frac{\eta}{2} \right)^{|S|} \prod_{\ell=1}^k F_{\eta \epsilon_\ell H^0}(c^0)_{J_\ell}. \quad (143)$$

Here the bracket denotes coefficient extraction only in the top variables u_r ; every child coordinate is the fixed number c^0 .

To prove the identity, each top cut occurs at most once in a composition. A contribution to this coefficient must include every cut in S and exclude all other cuts in Z_I . Split its ordered cut list at the members of S . This is a bijection with independent compositions of the gaps J_ℓ , excluding

their other collinear cuts; concatenation reverses it. Each selected cut supplies $-\eta/2$. At a further cut r in gap J_ℓ , collinearity of the gap endpoints gives

$$H^0(i, r, j) = \epsilon_\ell H^0(r_{\ell-1}, r, r_\ell). \quad (144)$$

Indeed write the endpoints on the common line as $p + t\omega$. In the determinant defining H^0 , changing the endpoint difference from $(t_k - t_0)\omega$ to $(t_\ell - t_{\ell-1})\omega$ multiplies its sign by ϵ_ℓ ; translating the line endpoint changes no determinant with ω . This calculation also holds when a_r lies on that line: both signs are then zero. Such a cut has zero weight in the gap transform, so including it there changes no sum. The remaining cuts and child factors are exactly those of the gap transforms in (143). Finite distributivity now proves that identity, including gaps with only one leaf.

For $\eta = 1$, Lemma 2.19 supplies a nontrivial gap with $\epsilon_\ell = 1$. Its factor is

$$F_{H^0}(c^0)_{J_\ell} = E_{J_\ell} = 0. \quad (145)$$

The first equality is the defining inverse equation, and the second uses the gap's length of at least two. Thus every nonempty top coefficient vanishes. The constant coefficient is the value with all $u_r = 0$, namely $F_{H^0}(c^0)_{ij} = E_{ij}$.

Induct on interval length to prove $c = c^0$ for the entire formal array H . Leaf coordinates are one in both arrays. If all shorter coordinates agree, the equation for $[i, j]$ has its unknown unary term c_{ij} with coefficient one; every nonunary child is already c^0 . The preceding coefficient calculation says that substituting also c_{ij}^0 makes this equation equal E_{ij} for every value of the top variables. Uniqueness of its unary solution therefore gives $c_{ij} = c_{ij}^0$. This proves the induction. In particular it does not require the formal sign assignments to be realizable.

Finally apply (143) to $I = [0, N]$ and $\eta = -1$. For every nonempty S , the second assertion of Lemma 2.19 supplies a nontrivial gap with $\epsilon_\ell = -1$. Its factor is again $F_{H^0}(c^0)_{J_\ell} = 0$. All nonconstant top coefficients vanish, and the constant term is $F_{-H^0}(c^0)_{0N}$. Substituting $c = c^0$ proves (18). The physical root edge is used here; an arbitrary open interval would not justify this last negative-gap step. $\square$

Proof of Corollary 2.21

Proof of 2.21. Replace each cut factor of every gate in the finite tree recursion by the linear form $(d_k - h_k)/2$ or $(d_k + h_k)/2$ of Lemma 2.18, and regard the entries as formal variables. The same recursion is $\Phi_{D,H}(b) = E$, with root output $\Phi_{D,-H}(b)_{0N}$. By the formal clause of Lemma 2.5, the recursion and the tree sum agree as polynomials in the gate variables; sending each gate variable $V^\pm(\pi)$, for a composition $\pi = (r_0 < \dots < r_s)$, to the product over the interior cuts r_k , $1 \leq k \leq s-1$, of the linear forms of Lemma 2.18 — $(d_k - h_k)/2$ for V^+ at an ordinary vertex and $(d_k + h_k)/2$ for V^- at the root, the empty product being 1 — defines a ring homomorphism from $\mathbb{Z}[V^\pm(\pi)]$ to $\mathbb{Z}[\frac{1}{2}][X_{ijk}] \subset \mathbb{Q}[X_{ijk}]$ (Definition 2.2 makes each gate that product of its cut factors), so this root output is the polynomial A_g of that lemma, as an identity in $\mathbb{Q}[X_{ijk}]$ and not only at realizable chirotopes. The ring identity of Lemma 2.9(i), followed by its two-sided inverse, makes this formal output $F_{-H}(F_H^{-1}(E))_{0N}$. This is a polynomial identity, so it also applies when the geometric signs have zeros; it does not evaluate the original step-function gates at zero. The geometric near and far entries are the reverse-ordered chirotope entries of the root word. This changes only fixed signs and names of the canonical variables. Lemma 2.20 therefore makes A_g constant in every entry that vanishes at P , with the other entries fixed.

Fix the labelled representative P and root g . By Lemma 1.16, choose a ball about P contained in its visible chamber, small enough that every initially nonzero area determinant keeps its sign.

By Lemma 1.17, this ball contains a generic tuple Q . On Q the polynomial equals $A_g(Q)$ by Lemma 2.18, and this equals $\tilde{A}_g(P)$ by Theorem 2.16. Formal constancy gives the same value after all silent entries are set to zero, which is $A_g(\chi(P))$. This proves the identity and the cancellation assertion throughout the weak locus. $\square$

Proof of Lemma 2.23

Proof of 2.23. Retain the original names of vertices and edges, replacing only E_j by the soft and return edges. The positive parameter interval may be shortened finitely many times in the argument.

First establish (G1). Original vertex triples are unchanged. For distinct original indices $k, l \neq j$, the area determinant for $(M_\varepsilon, \mu_k, \mu_l)$ tends to the nonzero determinant for (M, μ_k, μ_l) . For $k \neq j$, the determinant for $(M_\varepsilon, M, \mu_k)$ is $-\varepsilon \det(q, \mu_k - M)$, nonzero by admissibility. These exhaust the new triples. Finiteness gives one interval on which (G1) holds and all signs are constant.

The soft edge lies within $\varepsilon|q|$ of M . Every parent segment not incident to M has positive distance from M , by compactness and (G1) of P . Hence the soft edge misses all such segments for small ε . It meets the incoming edge only at M , since $\det(u, q) \neq 0$, and the return edge only at M_ε , since $\det(q, D_\varepsilon) = \det(q, v) \neq 0$. The return edge and the unchanged edge following B meet only at B : their directions tend to the independent directions at the parent corner B . All pairs of unchanged edges retain their parent intersections.

For each parent edge remote to E_j , compare it with the return edge. A disjoint parent pair has positive compact distance; the return segment converges uniformly at each affine parameter to E_j , so that distance remains positive. No independence of the directions is required for a disjoint pair. For a crossing parent pair, the directions have nonzero determinant and both intersection parameters lie in $(0, 1)$. The intersection equations form an invertible two-by-two linear system. Its coefficients converge to the parent coefficients and its determinant remains nonzero, so its inverse varies continuously. The two parameters converge to the parent values and stay in $(0, 1)$; the point converges as well, and the ordered direction determinant keeps its sign.

These are all return-edge pairs except the incoming edge E_{j-1} . Put $T = \det(u, v)$, of sign τ . The return edge meets E_{j-1} precisely when each supporting line strictly separates the endpoints of the other segment; the directions are independent since $\det(u, D_\varepsilon) \rightarrow T \neq 0$. For the incoming line the signed heights are

$$\det(u, M_\varepsilon - A) = \varepsilon \det(u, q), \quad \det(u, B - A) = T. \quad (146)$$

They are opposite exactly when $\chi_- = \tau$. For the return line based at M_ε , the height of M is

$$\det(D_\varepsilon, M - M_\varepsilon) = \varepsilon \det(q, v). \quad (147)$$

The height of A tends to $\det(v, -u) = T \neq 0$, so has sign τ throughout a sufficiently small interval. This pair of heights is opposite exactly when $\chi_+ = \tau$. Both tests are required: their conjunction is exactly the loop sector. There is no intersection of this pair in the other two sectors. In the loop sector its supporting-line equations remain invertible as $\varepsilon \rightarrow 0$. At zero the unique intersection is M , with incoming parameter one and return parameter zero. Thus $y \rightarrow M$, and these two parameters tend to one and zero.

All parent crossing points are away from M , since (G1) excludes a vertex on a nonincident edge, and they are pairwise distinct by (G2). The finitely many inherited points therefore remain pairwise distinct and away from M . The newborn, when present, stays distinct from them after another shrink. Every possible crossing has been enumerated. A concurrence of three distinct edge

interiors would involve pairwise remote edges, since (G1) makes distinct adjacent interiors disjoint. It would identify the points of two distinct unordered crossing pairs, which was excluded. Hence (G2) holds. The connected parameter interval now maps continuously into the labelled generic locus, so its image lies in one labelled chamber and hence one polygon chamber. This proves (i).

The turn determinants at M and M_ε are $\varepsilon \det(u, q)$ and $\varepsilon \det(q, D_\varepsilon) = \varepsilon \det(q, v)$; Definition 2.22 gives their signs $-\chi_-$ and $-\chi_+$. At B , the incoming direction tends to v and the outgoing one is unchanged, so the old nonzero turn sign persists. Every other old vertex except M has unchanged incident directions. Finally $\det(u, D_\varepsilon) \rightarrow T$ gives sign τ , and antisymmetry gives $-\tau$ for the opposite ordered pair. This proves (ii).

Distinct parent crossing parameters on each edge have positive pairwise separations and positive distances from zero and one. Their convergence preserves their order on the corresponding edges, including every edge whose crossing points move. This proves (iii). Without a newborn, concatenating these edgewise visit lists recovers the parent's cyclic Gauss word. With a newborn, its incoming visit is after all inherited visits on the incoming edge, and its return visit is before all inherited visits on the return edge. The intervening soft edge has no crossing. Thus the arc from a through M, M_ε to b has no other crossing visit. Deleting a, b leaves exactly the inherited lists, proving (iv). $\square$

Proof of Theorem 2.24

Proof of 2.24. We prove the finite duplication identity used below, including its two boundary cases. This proof uses the definitions of the boundary word, gates, open sums and soft insertion. It does not use a wall law or root independence. All formal arrays in the next two paragraphs are over $\mathbb{Q}$; their near entries need not be geometric signs.

Finite transforms and freedom to choose a near array. For a linearly ordered list with boundary positions $0, \dots, L$, define

$$\Phi_{D,H}(X)_{ij} = \sum_{\pi=(i=r_0 < \dots < r_s=j)} \prod_{k=1}^{s-1} \frac{D(r_{k-1}, r_k, r_{k+1}) - H(i, r_k, j)}{2} \prod_{k=1}^s X_{r_{k-1}, r_k}. \quad (148)$$

Put $F_H = \Phi_{0,H}$, $G_D = \Phi_{D,0}$, and let E_{ij} be 1 when $j = i + 1$ and 0 otherwise. The one-part composition contributes X_{ij} ; every other term uses only shorter intervals. Thus one solves $\Phi_{D,H}(X) = Y$ uniquely in increasing interval length, by subtracting the already determined nonunary terms from Y_{ij} . This gives a polynomial inverse on the finite array space. The constructed solution map is a right inverse; uniqueness applied to the target $\Phi_{D,H}(X)$ makes it a left inverse too.

The following factorization holds for arbitrary near and far arrays:

$$\Phi_{D,H} = F_H \circ G_D. \quad (149)$$

Indeed, expand the right side using an outer composition of $[i, j]$ and an inner composition of each outer part. Their ordered union is a refined composition ρ , together with a subset of its interior cuts marked as outer cuts. Conversely, that subset recovers the outer composition, and restricting ρ to each part recovers the inner compositions. These operations are inverse, including empty or full outer-cut subsets and one-part inner or outer compositions. At an outer cut r_k the factor is $-H(i, r_k, j)/2$. At a nonouter cut its immediate neighbors in ρ are exactly its neighbors inside its inner part, so the factor is $D(r_{k-1}, r_k, r_{k+1})/2$. Each interior cut supplies one factor, and the child-interval product is the refined product. Summing independently over the outer-cut subset gives the product of their sums, which is (148).

Write $b = \Phi_{D,H}^{-1}(E)$ and $c = F_H^{-1}(E)$. Applying (149) to the equation for b gives $F_H(G_D(b)) = E$; uniqueness gives $G_D(b) = c$. Applying the same factorization with $-H$ then gives

$$\Phi_{D,-H}(b) = F_{-H}(c). \quad (150)$$

Thus the complete root transform is independent of D with H fixed. For geometric arrays $D(x, y, z) = H(x, y, z) = \chi(a_z, a_y, a_x)$, the elementary gate identities $d\Theta(-hd) = (d - h)/2$ and $d\Theta(hd) = (d + h)/2$ show that $\Phi_{D,H}(b) = E$ is exactly the open-sum recursion, and its full $\Phi_{D,-H}(b)$ coordinate is A_g , including the one-part root term. The freedom in (150) concerns the complete transform; it asserts no equality of individual auxiliary and geometric open sums.

The formal duplication problem. Take a core ordered list $z_0, \dots, z_M$, $M \geq 2$, and distinguish $A = z_s$, where $0 \leq s \leq M$. Insert a new occurrence B immediately after A . The collapse r sends B to A and fixes the other occurrences. Let H^0 be any core far array of signs, and let $t(X) \in \{-1, 1\}$ be arbitrary for each core occurrence $X \neq A$. Assume the parent far array H^1 has the following entries, where order relations refer to positions in the linear lists:

$$\begin{aligned} H^1 \text{ on a triple not containing both } A, B & \text{ is the pullback of } H^0, \\ H^1(X, A, B) &= t(X) & (X < A), \\ H^1(A, B, Y) &= t(Y) & (B < Y). \end{aligned} \quad (151)$$

Only the neighbors in the next display are read cyclically in the core:

$$\eta_- = t(z_{s-1 \bmod (M+1)}), \quad \eta_+ = t(z_{s+1 \bmod (M+1)}), \quad k = \frac{\eta_- + \eta_+}{2}. \quad (152)$$

We will prove that the full parent root output is k times the core one.

Choose auxiliary near arrays as follows. On core triples with middle occurrence A , set $D^0(X, A, Y) = t(X)$; on other core triples set $D^0 = 0$. On parent triples not containing both A, B , pull back D^0 , and set

$$D^1(X, A, B) = t(X) \ (X < A), \quad D^1(A, B, Y) = t(Y) \ (B < Y). \quad (153)$$

These are formal arrays, not a claimed chirotope. Keep the far arrays fixed. Let $b^0 = \Phi_{D^0, H^0}^{-1}(E^0)$, and propose the following parent ordinary array:

parent interval I	b_I^1
I not containing both A, B	$b_{r(I)}^0$
$[A, B]$	1
$[A, Y], \ B < Y$	$\frac{\eta_+ - t(Y)}{2} b_{[A,Y]}^0$
$[X, B], \ X < A$	$\frac{\eta_- - t(X)}{2} b_{[X,A]}^0$
$[X, Y], \ X < A < B < Y$	$k b_{[X,Y]}^0$

(154)

Consecutiveness of A, B makes these cases exhaustive. Apart from the explicit singleton $[A, B]$, the collapsed endpoints are distinct. We verify $\Phi_{D^1, H^1}(b^1) = E^1$ on every interval before identifying b^1 with the actual auxiliary ordinary solution.

Intervals without the duplicate and the singleton. On an interval not containing both occurrences, collapse is an order-preserving bijection on vertices, compositions, gates and child intervals. Thus its transform is the corresponding core E^0 entry, which equals its parent E^1 entry. On $[A, B]$ the sole composition has one part and value 1, as required.

Intervals starting at A. Fix $[A, Y]$ with $B < Y$. For a core composition of $[A, Y]$, write X for its first boundary after A . There are exactly two parent presentations. Without a cut at B , its first child contains A, B and has multiplier $(\eta_+ - t(X))/2$. With a cut at B , the first child is the singleton $[A, B]$, and the rest is the core composition with its first endpoint renamed B . Its new gate at B has near entry $t(X)$ and far entry $t(Y)$. Removing the cut at B , when present, recovers the core composition; its presence distinguishes the two inverse expansions. All other children agree, the neighboring near gate pulls B back to A , and all other far entries are the stipulated pullbacks. The sum of the two local multipliers is, for an ordinary top vertex,

$$\frac{\eta_+ - t(X)}{2} + \frac{t(X) - t(Y)}{2} = \frac{\eta_+ - t(Y)}{2}, \quad (155)$$

and for a root top vertex it is

$$\frac{\eta_+ - t(X)}{2} + \frac{t(X) + t(Y)}{2} = \frac{\eta_+ + t(Y)}{2}. \quad (156)$$

In each case $t(X)$ cancels, leaving a composition-independent multiplier of the corresponding core transform. The ordinary transform vanishes when the core has at least two leaves, since its E^0 entry is zero. When the core has one leaf, Y is the actual successor of A , so $t(Y) = \eta_+$ and the ordinary multiplier itself vanishes. Thus every such parent ordinary coordinate has its required zero value, including the two-leaf parent interval.

Intervals ending at B. Fix $[X, B]$ with $X < A$. In a core composition of $[X, A]$, let Y be the last boundary before A . Without a cut at A the last child contains the soft edge and has multiplier $(\eta_- - t(Y))/2$. With that cut it ends in the singleton $[A, B]$; the new gate has near entry $t(Y)$ and far entry $t(X)$. Remove the cut A when present and rename the endpoint B as A to recover the core list. The cut's presence distinguishes the two expansions. The neighboring near gate pulls B back to A and all other factors agree. Consequently the ordinary and root multipliers are, respectively,

$$\frac{\eta_- - t(Y)}{2} + \frac{t(Y) - t(X)}{2} = \frac{\eta_- - t(X)}{2}, \quad (157)$$

$$\frac{\eta_- - t(Y)}{2} + \frac{t(Y) + t(X)}{2} = \frac{\eta_- + t(X)}{2}. \quad (158)$$

For a core interval with at least two leaves its ordinary transform is zero. For one leaf, X is the actual predecessor of A , so $t(X) = \eta_-$ and the multiplier is zero. This checks the other two-leaf parent boundary as well.

Intervals strictly spanning A, B. Fix $[X, Y]$ with $X < A < B < Y$. Collapse every parent cut list, retaining only one copy when both A, B are cuts. If the resulting core composition does not cut at A , it has exactly one child spanning A . It has exactly one parent presentation, with neither A nor B a cut; that child strictly spans the duplicate and has multiplier k . All top gates pull back unchanged. This includes the one-part composition: the proposed leading b^1 coordinate is being checked, not assumed to satisfy a recurrence.

If the core composition cuts at A , let L, R be its neighboring cuts and put $a = t(L)$, $b = t(R)$ and $h = H^0(X, A, Y)$. Its core near entry at A is a . There are precisely three parent presentations:

cuts among A, B	exceptional child multiplier
B only	$(\eta_- - a)/2$
A only	$(\eta_+ - b)/2$
A and B	1 ($[A, B]$ is a singleton)

(159)

Collapse recovers the core composition, while its subset of cuts among A, B recovers the row; expanding according to each row is the inverse. The fourth subset, neither cut, gives the preceding no-core-cut case and cannot occur here. The neighboring gate at L has right neighbor A or B , both pulling back to A ; the neighboring gate at R has the analogous left neighbor. Their near entries are unchanged. The far entries at A and B both pull back to h , and all exterior gates and remaining child factors agree with the core.

For an ordinary top, either one-cut row has the core gate $(a - h)/2$; the two-cut row replaces it by $(a - h)(b - h)/4$. The two one-cut child multipliers sum to $k - (a + b)/2$. After subtracting k times the core gate, the residual is therefore

$$R_o = -\frac{a+b}{2} \frac{a-h}{2} + \frac{(a-h)(b-h)}{4}. \quad (160)$$

Factoring $(a - h)/4$ leaves $-(a + b) + (b - h) = -a - h$, so

$$R_o = \frac{(a-h)(-a-h)}{4} = \frac{h^2 - a^2}{4} = 0. \quad (161)$$

The last equality uses exactly $a, h \in \{-1, 1\}$. For a root top the corresponding residual is

$$R_r = -\frac{a+b}{2} \frac{a+h}{2} + \frac{(a+h)(b+h)}{4}. \quad (162)$$

Factoring $(a + h)/4$ leaves $-(a + b) + (b + h) = -a + h$, so

$$R_r = \frac{(a+h)(-a+h)}{4} = \frac{h^2 - a^2}{4} = 0. \quad (163)$$

Thus both transforms equal k times their respective core transforms on every strictly spanning interval. Such a core interval has at least two leaves, so its ordinary E^0 entry is zero. No relation among different remote values of t was used.

All interval cases have now verified $\Phi_{D^1, H^1}(b^1) = E^1$. Finite triangular uniqueness proved above identifies (154) with the actual auxiliary ordinary solution. The root transforms already checked therefore compute the actual auxiliary root outputs. We did not assume that any auxiliary binary open sum vanishes.

For the full interval, if $0 < s < M$ the strictly spanning multiplier is k . If $s = 0$, the final endpoint $Y = z_M$ is the cyclic predecessor of A , so $t(Y) = \eta_-$ and (156) gives k . If $s = M$, the initial endpoint $X = z_0$ is the cyclic successor of A , so $t(X) = \eta_+$ and (158) gives k . These are all positions. Finally (150), applied separately to the parent and core, restores their geometric near arrays without changing their complete root outputs. This proves the formal duplication identity.

Specialization to the soft polygon. Write $A = \mu_j$ and $B = A + \varepsilon q$, and let the parent have $n + 1$ vertices, so the fixed core P has $n \geq 3$ vertices. Every parent determinant not containing both A, B either is a core determinant or replaces A by $A + \varepsilon q$ in one. For example, bilinearity gives

$$\det(X - A - \varepsilon q, Y - A - \varepsilon q) = \det(X - A, Y - A) - \varepsilon \det(q, Y - X). \quad (164)$$

The potential quadratic term vanishes because $\det(q, q) = 0$; alternation gives the same bound when A is in another position. The core has finitely many nonzero distinct-triple determinants, so their absolute values have a positive minimum δ . The finitely many coefficients $\det(q, Y - X)$ have a finite maximum absolute value K . Choose a common positive tail with $\varepsilon K < \delta$ when $K > 0$; if $K = 0$, no restriction is needed. On this tail every such parent sign equals its pulled-back core

sign. All core edges have a positive minimum length; shrinking further if necessary also keeps the return edge $\tilde{\lambda}_j - \varepsilon q$ nonzero. Take this tail inside $(0, \varepsilon_0)$.

For every other core vertex X , set

$$t(X) = \chi(B, A, X) = -\operatorname{sgn} \det(q, X - A). \quad (165)$$

Indeed $A - B = -\varepsilon q$ and $X - B = X - A - \varepsilon q$, so their determinant is $-\varepsilon \det(q, X - A)$. Admissibility makes it nonzero. For a boundary triple (X, A, B) the far entry is $t(X)$; for (A, B, Y) it is $\chi(Y, B, A) = \chi(B, A, Y) = t(Y)$ by a cyclic permutation. Together with the common-tail calculation, these are exactly (151).

The two cyclic neighboring values are the stated attachments. Writing $u = \mu_j - \mu_{j-1}$ and $v = \mu_{j+1} - \mu_j$, the predecessor gives

$$t(\mu_{j-1}) = -\operatorname{sgn} \det(q, -u) = -\operatorname{sgn} \det(u, q) = \chi_-; \quad (166)$$

the equality follows from bilinearity and antisymmetry. The successor gives directly

$$t(\mu_{j+1}) = -\operatorname{sgn} \det(q, v) = \chi_+. \quad (167)$$

Cut the boundary word immediately after the specified nonsoft root. The occurrences A, B are consecutive in this linear word. If the root is the incoming edge ending at A , the word is soft-first ($s = 0$). If it is the return edge starting at B , the word is soft-last ($s = M$). Every other nonsoft root gives the strictly spanning case. Here the core word has $M = n - 1 \geq 2$, so the formal result applies even to a triangle core; no two-gon amplitude is invoked. Collapse of B to A retains each other root's physical endpoints, except that the return edge becomes the core edge $E_j = [\mu_j, \mu_{j+1}]$. Thus the collapsed word is exactly the boundary word of (P, g') , including cyclic wraparound. The soft root itself would put A, B at opposite ends of the cut word and has been excluded throughout.

The duplication identity now gives (20) with $k = (\chi_- + \chi_+)/2$. Evaluating this expression on the three attachment-sign patterns gives the sector multipliers in the statement. Neither the crossing classification of the soft family nor a rotation formula for it was needed in this calculation. $\square$

Proof of Proposition 2.26

Proof of 2.26. The relabellings preserve (G1), because they permute distinct vertex triples. Put $N = n - 1$. For the shift, the boundary vertex at position k of $(\sigma P, g)$ is μ_{g+2+k} , which is precisely position k of $(P, g + 1)$. Every interval composition, near sign and far sign is therefore identical. Induction on interval length in the open-sum recursion gives equality of all open sums, and the complete root sums agree, proving (i).

For reversal, edge $g^\vee$ runs from $\mu_{2-g^\vee} = \mu_{g+1}$ to $\mu_{1-g^\vee} = \mu_g$, the reverse of the original physical root. If $a_k = \mu_{g+1+k}$ is the original boundary word and a'_k the word at this reversed root, then

$$a'_k = \bar{\mu}_{g^\vee+1+k} = \mu_{1-g^\vee-k} = \mu_{g-k} = a_{N-k} \quad (0 \leq k \leq N). \quad (168)$$

In the last equality the indices differ by n .

Use the tree expression of Lemma 2.5. Reflect the order of the children at every internal vertex of a rooted plane tree, and relabel boundary positions by $k \mapsto N - k$. This is an involution on the tree set: the root is still allowed one or more children, ordinary vertices still have at least two, and leaf order is exactly the reversed boundary order. A composition $(r_0 < \dots < r_s)$ becomes $(N - r_s < \dots < N - r_0)$ on the reflected interval. At corresponding cuts its ordered near triple reverses its first and last entries; the same is true of its ordered far triple. Thus $d' = -d$ and

$h' = -h$. The product $h'd' = hd$, so either gate retains its step-function argument and acquires precisely one minus sign:

$$d'\Theta(\mp h'd') = -d\Theta(\mp hd). \quad (169)$$

This holds for both ordinary and root gates, including zero factors. The separate factor (-1) for each ordinary internal vertex is unchanged by the involution.

If the tree has I internal vertices and N leaves, its number of child edges is $I + N - 1$. The sum of the arities over internal vertices counts each child edge once. A vertex of arity s contributes $s - 1$ gates, so the total number of gates is

$$\sum_{v \text{ internal}} (s_v - 1) = (I + N - 1) - I = N - 1 = n - 2. \quad (170)$$

A unary root contributes zero to this count, so its allowed presence creates no exception. Every corresponding tree term therefore differs by $(-1)^{n-2} = (-1)^n$. Summing over the involution gives (21). For $n = 3$ the same count is one; no exceptional base convention is required. Neither part proves or assumes root independence on a fixed polygon. $\square$

Proofs for Section 3

Proof of Lemma 3.5

Proof of 3.5. Finite successor model. Mark the traversal circle at every original vertex and every crossing visit. This is a nonempty finite cyclic list, even if there are no crossings. Each consecutive pair bounds an oriented straight subsegment. Write ρ for the successor permutation of the list. At the two distinct marked visits a, b of a selected crossing, replace the two arrows by

$$\rho_S(a) = \rho(b), \quad \rho_S(b) = \rho(a). \quad (171)$$

At an unselected marked point keep its outgoing arrow. Selected crossings have disjoint pairs of marked visits. Thus these swaps operate on disjoint outgoing slots, commute, and give a permutation ρ_S independent of their order. Its cycles are precisely the oriented reconnections: the node a retains its incoming arc and leaves along the old outgoing arc at b . This also implements the incoming-visit convention explicitly.

Splitting and induced order. We prove by induction as selected pairs are processed that the cycles traverse their marked nodes in the cyclic order inherited from the original list, and that the two endpoints of each unprocessed selected pair lie on one current cycle. Initially there is the one original cycle. Suppose a current selected pair has endpoints a, b on a cycle L . In its inherited cyclic order write that cycle as

$$(a, A_1, \dots, A_p, b, B_1, \dots, B_q). \quad (172)$$

Either of the lists between the endpoints may be empty. Swapping the outgoing arrows gives exactly the two cycles

$$(a, B_1, \dots, B_q), \quad (b, A_1, \dots, A_p). \quad (173)$$

Each retains the induced cyclic order. Every remaining selected pair has both endpoints in one of the two open arcs between a, b on the original circle, since it does not interlace this selected pair. If it was on L , its two endpoints therefore enter the same new cycle; if it was on another cycle, it is unaffected. This proves the inductive claims and shows that every selected reconnection splits one cycle into two, never merging cycles. Hence the final count is $|S| + 1$. The endpoints a, b enter different cycles in (173), and all later operations only split. They remain on different final carriers. This proves (i).

Noncrossing away from selected endpoints. The preceding induction also gives noncrossing ownership of the unmarked open traversal arcs. Here are the details, including arbitrary points on those arcs. For no selected pairs there is one carrier, so there is nothing to check. In a splitting step the one old carrier L is divided between the two open arcs cut out by a, b , with the endpoints omitted for this paragraph. All other carrier sets are unchanged. If two new carrier blocks alternated with four points, and at most one of them came from splitting L , the same four points would violate the previous partition's noncrossing property, using L for that new block. If both came from splitting L , their points would alternate between the two open arcs of a, b . The four disjoint cyclic transitions between these four points would then each contain one of a, b . Two endpoints cannot lie in four disjoint transitions. This is impossible. Induction proves noncrossing for all traversal points away from preimages of S .

Including every selected visit. Assume four distinct crossing visits violated (iv). Move any of these that are selected a sufficiently small distance backwards along the original traversal circle; leave the unselected ones fixed. There are finitely many selected visits and four distinct chosen points, so the distances can be chosen positive and small enough that no selected endpoint is

passed and the cyclic order of the four points is unchanged. By the incoming-visit convention, each moved point has exactly the carrier assignment of its original selected visit. Now all four points avoid preimages of S and give the forbidden alternation of the previous paragraph. This contradiction proves the all-visits assertion.

Ownership of unselected crossings. Let $x \notin S$. If it interlaces no selected pair, its two endpoints start on the same original cycle and remain together after every split: at a split whose carrier contains them, noninterlacement puts them in one open arc of that selected pair. Thus they end on one carrier. If x interlaces a selected pair s , process s first, which is permitted by order independence. Its split puts the two visits of x on different cycles. Later reconnections only split, so they can never be reunited. This proves both ownership assertions in (iii).

Nonzero segments and corner signs. Genericity makes all original vertices distinct: coincident vertices would give a zero triple with any third index. It also makes every edge nonzero and every original turn nonzero. No vertex lies on a nonincident closed edge: such an incidence would give a zero triple. Two remote edges meeting at a point have neither endpoint there, and their directions cannot be parallel, since coincident lines would give collinear original vertices. Thus their intersection is an interior transverse crossing. Distinct crossings have distinct points, as a coincident pair of crossing points would lie on at least three edge interiors, contrary to genericity.

Consequently consecutive marked points on an original edge are distinct in the plane, with strictly increasing parameters; their connecting subsegment has positive length. At a selected node a , the point of the plane agrees with its twin b , but its new successor is the old successor of b , which is at positive distance along b 's outgoing subsegment. Equation (171) does not insert a zero segment from a to b . Omitting unselected crossing marks joins successive positive subsegments of the same original edge in the same direction. Hence every edge between consecutive carrier corners is nonzero and inherits an original oriented direction up to a positive scalar.

At an original vertex i no reconnection occurs. The incoming and outgoing carrier directions are positive multiples of d_{i-1} and d_i , so their determinant has sign τ_i . At a selected crossing of E_i, E_j , the node assigned by incoming E_i leaves along E_j ; its determinant is a positive multiple of $\det(d_i, d_j)$. The twin node has a positive multiple of the opposite determinant $\det(d_j, d_i) = -\det(d_i, d_j)$. Transversality makes these nonzero, with opposite signs. These are all corner types. In particular none is antiparallel or flat.

Each cycle contains a corner: a cycle having only unselected visits would follow the unchanged original successor around the full parent traversal, which contains original vertices. Thus its nonzero segments form a closed polygonal cycle with the indicated corners. One nonzero segment cannot close. A closed polygon of exactly two nonzero segments has opposite segment vectors and hence antiparallel corners, already excluded. Every carrier therefore has at least three corners. This proves (ii), including membership in the regular locus.

Exactly the retained self-intersections. Every traced carrier segment is an original straight subsegment with its inherited parameter direction. Distinct traversal subsegments of one original edge have disjoint interiors, so no overlap is introduced. Original adjacent edges meet only at their common original vertex, which has one traversal occurrence and lies on one carrier. All other possible intersections are original crossings. At a selected crossing the two incoming visits belong to different carriers by (i); a single carrier therefore passes through that smoothing point only once. An unselected crossing is visited twice by a carrier precisely when both of its original visits are assigned to that carrier. Otherwise it is a meeting of two distinct carriers, not a self-intersection of either. Every retained self-crossing is still transverse with its unchanged straight germs. It is not a corner because original vertices and selected crossing sites are all distinct from it. No carrier can have a triple point, since no point of the parent has three traversal preimages and each selected site is used only once by a given carrier. These statements prove the remaining assertions of (iii). $\square$

Proof of Lemma 3.7

Proof of 3.7. Fix a labelled representative of P . By Lemma 1.17 the visible signature is locally constant on $\mathcal{W}_n$ and generic tuples are dense there; choose a generic tuple P' in a neighbourhood of P on which the signature is constant. Its crossings have the same pairs of parent edges and the same order along each edge as those of P , so inserting the original vertex labels gives the same finite marked cyclic list, the same successor swaps, the same independent sets, the same cycle assignments and the same incoming-visit ownership for P and for P' . Lemma 3.5 for P' therefore gives (i), (iv) and the two interlacement and ownership assertions of (iii) for P , these being statements about the finite marked list only. The geometric assertions are verified directly at P .

All vertices of P are distinct: coincident adjacent vertices would give a zero edge, and if two nonadjacent vertices coincided, either edge incident to one of them would be a nonincident closed edge containing the other, contrary to Definition 1.15. All edges and all turn determinants are nonzero by that definition. Two remote edges meet at most once and transversally, and a meeting point is interior to both edges, since a vertex on a nonincident closed edge is excluded. Distinct crossings have distinct points, since two crossings at one point would put at least three edge interiors there, contrary to (G2).

Consequently the marked points on each original edge have strictly increasing parameters with positive separations. At a selected visit the new successor is the old outgoing subsegment of its twin visit, of positive length, so no zero segment is inserted between the coincident twin points; omitting unselected marks concatenates positive subsegments of one directed edge. Every carrier edge is therefore nonzero and is an original directed edge vector times a positive scalar. At an original corner the two carrier vectors are positive multiples of the parent's consecutive vectors, so the turn keeps the parent's nonzero sign; at a selected crossing of E_i, E_j the two smoothing corners have determinants with the signs of $\det(d_i, d_j)$ and $\det(d_j, d_i)$, nonzero and opposite by transversality. No corner is flat or antiparallel. Every cycle contains a corner, since a cycle with only unselected visits would follow the unchanged parent successor and reach an original vertex; one nonzero edge cannot close, and two closing nonzero edges would be opposite and give an antiparallel corner. Hence every carrier has at least three corners. This proves (ii), including membership in the regular locus.

For the geometric part of (iii): distinct traced subsegments of one parent edge have disjoint interiors; adjacent parent edges are not collinear, their turn determinant being nonzero, and meet only at their common vertex; every other meeting of two carrier segments is an original crossing. At a selected crossing the two visits belong to different carriers by (i), so a carrier passes through that site once. At an unselected crossing a carrier passes twice exactly when it owns both visits; the two germs are the parent's transverse germs, and the site is not a corner because it differs from every original vertex and every selected site. No point has three preimages on a carrier: that would give three parent traversal preimages, whereas (G2) and the vertex exclusions allow at most two. These are the remaining assertions of (iii).

Cyclic relabelling renames the marked list, the independent sets and the cycles without changing the traced oriented curves or the ownership of visits, so every construction and conclusion descends to the polygon. $\square$

Proof of Lemma 3.10

Proof of 3.10. Choose pairwise disjoint small crossing discs, each missing every remote arc. At a smooth transverse double point, in local parameters the map $(s, t) \mapsto \gamma_1(s) + \gamma_2(t) - p$ has invertible derivative. After its derivative is used as the linear coordinate frame, its inverse straightens the two

branches and has derivative arbitrarily close to the identity on a sufficiently small disc. Multiply its displacement by a cutoff which is one on a smaller disc and zero near the outer boundary. Choose the inner and outer radii with a fixed ratio; shrinking them makes the displacement derivative, including the cutoff term, have norm less than one. The interpolated maps $x \mapsto x + tv(x)$ are injective, since $|v(x) - v(y)| < |x - y|$, and surjective, since $x = y - tv(x)$ is solved uniquely by contraction. They are homeomorphisms fixed outside the disc and give the required local isotopy.

After this operation the remaining edge pieces are compact embedded regular arcs with disjoint thin strips. A strip can be constructed by $(s, u) \mapsto \gamma(s) + un(s)$: for a smooth arclength arc its Jacobian is $1 - u\kappa(s)$, nonzero for uniformly small u . Local injectivity follows in a tangent chart; pairs of parameters outside those local intervals have positive distance by compactness, giving global injectivity after one further shrink. The strips can be chosen disjoint outside the fixed crossing discs. Polygonal interpolation at a sufficiently fine finite partition has tangent directions uniformly close to the arc's tangent. Its strip projection is therefore strictly increasing, so it is one continuous normal graph $u = f(s)$, with $f = 0$ at fixed ends and arbitrarily small norm. In strip coordinates use

$$(s, u) \mapsto (s, u + t\chi(u)f(s)), \quad 0 \leq t \leq 1. \quad (174)$$

Take $\chi = 1$ near zero, $\chi = 0$ near the strip boundary, and $\|\chi'\| \|f\| < 1$. Each fibre map is strictly increasing, onto and fixed at the strip ends; its inverse is continuous. This is an ambient page isotopy. Finite polygonal corner pieces already need no straightening; a smooth replacement in a clean corner disc is treated by the same graph construction. Subdivide the fixed straight crossing germs and all resulting polygon edges. This gives a finite PL graph model. All maps preserve page orientation and the named finite data. Their products with the height coordinate give ambient isotopies of the diagram lifts. There is no application of a smooth theorem to the image of an arbitrary circle homeomorphism. $\square$

Proof of Lemma 3.11

Proof of 3.11. Here is a finite triangulation and disc construction, so no Jordan or Schoenflies conclusion is being assumed. Remove one point off the curve and put the polygon inside a large rectangle. Extend its finitely many edge lines, and also lines through its vertices in a generic auxiliary direction, across that rectangle. The arrangement has convex cells; subdivide boundary segments at all intersections and triangulate each convex cell by a fan from an interior point. Collinear redundant vertices are simply retained as subdivisions, never used to claim a strict ear. The circle is now a simple edge cycle C in a finite triangulation of the rectangle. Cap the rectangular boundary by a second triangulated disc. This cap represents the actual one-point-compactified exterior: from the rectangle's center, write its boundary as $r = R(u)$ and map an exterior point ru to $(R(u)/r)u$, with infinity sent to zero. This has a continuous radial inverse and identifies the exterior with a disc. The two-disc model is a triangulated sphere by radial identification with the upper and lower hemispheres. The original exterior is part of this same triangulation, not just a collar of C .

Work with chains over $\mathbb{F}_2$. If the sphere triangulation has V, E, F cells, then $V - E + F = 2$. This follows first for the two-disc rectangle model and is unchanged by subdividing an edge, adding a vertex inside a face, or drawing a diagonal; these are precisely the finite arrangement and fan subdivisions just used. The vertex-edge graph and the face-adjacency graph are connected. The incidence map on edges has rank $V - 1$, since a spanning tree gives $V - 1$ independent boundary columns and every edge boundary has even coordinate sum. Hence its cycle space has dimension $E - V + 1 = F - 1$. The face-boundary map has kernel exactly the two constant face assignments: across every edge a boundary-zero face assignment has equal coefficients on its two adjacent faces,

and the face-adjacency graph is connected. Its rank is therefore $F - 1$. Face boundaries are cycles, so every cycle, in particular C , is the boundary of a face assignment, unique up to complement.

Color faces by that assignment. Across an edge of C the colors differ; across any other edge they agree. At a vertex of C its two cycle edges divide the cyclic triangle fan into exactly two consecutive fans, one of each color. At every other vertex all incident faces have one color. Thus each color closure is a compact triangulated surface whose boundary is C , with no pinch point. Following C through those local fans shows that all boundary-adjacent faces of a given color lie in one connected face component. Any further face component would have no boundary edge and hence no adjacency to its complement, contradicting connectedness of the full dual graph. Thus there are exactly two connected color regions, R_0, R_1 , each with its single boundary C .

We next prove, rather than assume, that each is a disc. For any connected triangulated surface R with nonempty boundary, take a spanning tree of its triangle-adjacency graph, rooted at a triangle having a boundary edge. Remove that triangle across that boundary edge, and then remove each triangle after its parent, across their common edge. At its turn the chosen edge is free: the parent triangle is gone and no third triangle is incident to that edge. A triangle and its free edge collapse onto its other two edges. This preserves connectedness and Euler characteristic and leaves, after finitely many steps, a connected graph K . Consequently $\chi(R) = \chi(K) = 1 - \beta_1(K) \leq 1$. Since the two regions meet in C and $\chi(C) = 0$, cell counting gives $\chi(R_0) + \chi(R_1) = 2$. Both inequalities are therefore equalities, so both residual graphs are trees.

For completeness these collapses also construct disc parametrizations, not only a homotopy equivalence. Maintain a small closed regular neighbourhood of the current collapsed complex, with vertex polygons and edge rectangles. A free-triangle collapse changes its neighbourhood by removing a boundary notch, which gives a boundary-compatible PL homeomorphism. In the standard triangle $x, y \geq 0, x + y \leq 1$, whose free edge is $x + y = 1$, use the neighbourhood of the other two edges

$$\{x, y \geq 0, x + y \leq 1, \min(x, y) \leq \varepsilon\}, \quad 0 < \varepsilon < 1/2. \quad (175)$$

Its new boundary path runs from $(1, 0)$ through $(1 - \varepsilon, \varepsilon), (\varepsilon, \varepsilon), (\varepsilon, 1 - \varepsilon)$ to $(0, 1)$. Subdivide the old free edge at the intersections of the rays from $(0, 0)$ through those three corners. Map the resulting triangle fan affinely to the corresponding fan of the notched region, fixing the two retained edges pointwise. Every triangle map has positive determinant, the fans meet on their common radial edges, and the images cover the notched region. Thus the map is a PL homeomorphism, fixed where that triangle joins the rest. In thin edge rectangles use the same product map, with identity on the outer attaching boundary. This proves by induction the stated regular-neighbourhood invariant for every free collapse. Choose widths successively smaller than the positive distances between disjoint cells; finiteness makes all choices compatible.

A regular neighbourhood of a tree is a disc by an explicit leaf induction. One vertex gives a polygonal disc. Adding a leaf attaches one edge rectangle and one leaf disc along a single boundary interval. Triangulate that rectangle and disc into a fan along the interval; the same positive fan maps identify their union with a longer boundary collar. Thus each step is a disc and records its boundary parametrization. Reverse the finite free-collapse maps to parametrize R_0 and R_1 . This completes both the interior and exterior constructions.

Finally identify two resulting discs with convex polygonal unit discs. Subdivide their boundaries until a specified positive PL boundary map is affine on each segment, choose one interior vertex in each, and map the corresponding boundary fans affinely, interior vertex to interior vertex. Their cyclic order makes all determinants positive. For a continuous positive boundary map instead use radial coordinates and $(r, u) \mapsto (r, b(u))$; the inverse uses b^{-1} and continuity at the center follows from the radial bound. This proves both extension claims with the prescribed boundary, not just

an unlabelled disc type. □

Proof of Lemma 3.12

Proof of 3.12. We first give the required local motion. In a planar chart let $v(x)$ be supported in a ball and Lipschitz with constant less than one. Then $x \mapsto x + tv(x)$ is a homeomorphism for $0 \leq t \leq 1$: injectivity follows from the strict Lipschitz estimate and surjectivity from contraction applied to $x = y - tv(x)$. Its inverse depends continuously on (y, t) by the same estimate. For a sufficiently small increment $x \mapsto Mx + b$ of a positive affine motion, choose a cutoff equal to one near the moving small triangle and zero near the chart boundary, and set $v(x) = \psi(x)((M - I)x + b)$. A sufficiently fine time subdivision makes its Lipschitz constant less than one. Thus any compact path of small positive affine triangles, with a positive clearance from the outer boundary, has an ambient extension fixed on that boundary.

Here a path between two interior triangles can be given explicitly: contract each about its centroid to an arbitrarily small copy, change the first shape to the second by a positive linear path, translate its centroid along a polygonal path in the chart, and expand at the end. The compact centroid path has positive boundary clearance, so a common small size suffices. A positive linear path exists without a matrix connectivity assumption: rotate the first column of a matrix to the positive horizontal axis; the remaining upper-triangular matrix has positive diagonal entries, and its straight interpolation to the identity keeps both diagonal entries positive. This supplies the shape path. All the local extensions just constructed are explicitly isotopic to the identity and never require an extension on a region with three boundary components.

Now let h be a positive PL sphere map. If it is not the identity, choose p in the interior of a PL cell with $h(p) \neq p$. Choose a PL disc chart containing $p, h(p)$ in its interior: remove a small triangle about a third point and use Lemma 3.11 on its complementary disc. Move p slightly, if necessary, off the finitely many subdivision edges and their h -preimages; the inequality $h(p) \neq p$ persists. Refine so the chart and h are affine on a sufficiently small triangle B about p and on its image. Take B so small that B and $h(B)$ are disjoint. Their two chart images are positive triangles. The preceding local construction gives a positive map k , supported in that chart interior and isotopic to the identity, with $k(h(B)) = B$. Thus $g = kh$ preserves B setwise. No claim about relocating arbitrary prescribed discs is needed.

Its boundary map b is positive. In a positive circle coordinate, let $\beta : \mathbb{R} \rightarrow \mathbb{R}$ be an increasing lift of b^{-1} with $\beta(x + 1) = \beta(x) + 1$. The maps $\beta_t(x) = (1 - t)x + t\beta(x)$ are strictly increasing with that same periodicity and descend to a circle isotopy. The two discs bounded by ∂B have parametrizations from Lemma 3.11; extend β_t radially over both and glue their common boundary maps. Postcomposing g by its last map fixes ∂B pointwise.

A boundary-fixed unit-disc homeomorphism a has the explicit isotopy

$$A_t(x) = \begin{cases} ta(x/t), & |x| \leq t, \\ x, & t \leq |x| \leq 1, \end{cases} \quad A_0(x) = x. \quad (176)$$

At $|x| = t$ the formulas agree. At $t = 0$ both the moved point and its image have norm at most t , proving continuity. Replacing a by a^{-1} gives the inverse, also continuous. Apply this separately on the two boundary-fixed discs and glue. Concatenating, and reversing where necessary, the relocation, boundary-correction and two Alexander isotopies gives an isotopy from the identity to h . Every constituent preserves orientation. □

Proof of Lemma 3.13

Proof of 3.13. Choose a height bound H larger than the absolute values of both assignments. In a small crossing chart each of the two branch heights extends continuously to the chart by holding it constant in the transverse coordinate of that branch. Shrink the chart so both source and target extensions retain their strict lower/upper ordering. For each page point x define an increasing piecewise-linear map f_x of the height axis which sends, in order, the four source levels $-H, h_{0,U}(x), h_{0,O}(x), H$ to $-H, h_{1,U}(x), h_{1,O}(x), H$, and is identity outside $[-H, H]$. All slopes are positive and depend continuously on x . On a strip meeting only one branch use the same construction with just that one intermediate source and target height. Take these strips outside smaller crossing discs so that no strip meets an unaccounted second branch. These finitely many local maps cover the shadow.

Choose nonnegative continuous weights subordinate to these chart interiors and an identity map on their complement, whose sum is one and whose identity weight is zero on the shadow. Explicitly use distances to the complements of slightly smaller chart interiors, together with the distance to the closed shadow, and divide by their positive sum. The smaller charts still cover the shadow, so the sum never vanishes; extending weighted displacements by zero across a chart boundary is continuous since all displacements are bounded by $2H$. Set $f(x, z)$ to the weighted average of the local $f_x(z)$ and the identity value z . A convex combination of strictly increasing functions is strictly increasing, and all functions equal z for $|z| \geq H$. On a point of the shadow every participating local map sends each actual source height there to its actual target height, so their average does too.

Thus

$$(x, z) \mapsto (x, (1-t)z + tf(x, z)), \quad 0 \leq t \leq 1, \quad (177)$$

is a homeomorphism on every fibre and fixes the product boundary. Continuity of its inverse follows directly: if (x_j, t_j, w_j) converges, its inverse heights are bounded by $\max(H, |w_j|)$; every convergent subsequence has the unique inverse limit by continuity and strict monotonicity, hence the full sequence converges. The maps extend by the identity to the ambient three-sphere and form an ambient isotopy. They transport the original parameter orientation. No tubular smooth-extension hypothesis is used. $\square$

Proof of Theorem 3.14

Proof of 3.14. Add one point at infinity to each oriented plane. The resulting pages are oriented spheres, and the diagrams miss the added point. View compatible lifts in a thin product neighbourhood of the spherical page in S^3 . The hypothesis is exactly a named record isomorphism of Definition 3.9: its sign convention is that of Definition 3.8. Apply Lemma 3.10 to each diagram first, transporting its finite record. Choose circle coordinates pulled back from the subdivided PL shadow arcs; no original smooth parametrization at unmarked points is prescribed. Subdivide these arcs and choose the positive PL extension $\bar{\Phi}$ of the finite occurrence bijection as in Definition 3.9. The construction below concerns this chosen extension, not an arbitrary prescribed map on every circle parameter. For paired visits v, w , the determinant sign of the ordered tangents (u_v, u_w) is the recorded crossing sign if v is the over visit, and its negative if v is the under visit. Thus the named bijection preserves this determinant sign for either ordering. No reversal of traversal, page orientation or crossing signs is allowed.

Write v^+, v^- for the outgoing and incoming rays at a crossing, both pointing away from it. For

$w = \tau(v)$ the positive cyclic order is

$$\begin{cases} (v^+, w^+, v^-, w^-), & \det(u_v, u_w) > 0, \\ (v^+, w^-, v^-, w^+), & \det(u_v, u_w) < 0. \end{cases} \quad (178)$$

Indeed in the first case u_w lies between u_v and $-u_v$ counterclockwise, and in the second $-u_w$ does; the other two rays are antipodes. The named map preserves these orders. It induces a graph map

$$f_\Phi(\gamma(x)) = \gamma'(\overline{\Phi}(x)). \quad (179)$$

This is well defined because equal images of distinct parameters are exactly paired crossing preimages. The inverse argument proves it a homeomorphism of the directed shadow graphs, preserving the named rays and their rotation systems.

Suppose first there is at least one crossing. The shadow is a connected finite four-valent graph. Construct its closed regular neighbourhood N from vertex discs and edge bands. For each boundary circle, the connected interior of N lies on one side; the other disc supplied by Lemma 3.11 is disjoint from N . These complementary discs cannot be nested, since an inner boundary belongs to N and would then lie in a disc disjoint from N . To see that their interiors cover the complement of N , join any point outside N by a path to its connected interior and take the first boundary hit; until that hit the path remains on the opposite side of that boundary circle, so the initial point lies in its complementary disc. Adding the collars proves that the shadow fills the sphere with disc faces. Connectedness of an abstract ribbon alone would not prove this: if it has q crossing vertices and b boundary cycles, capping gives genus $(q + 2 - b)/2$, which need not be zero without the actual sphere embedding.

Attach abstract oriented discs and untwisted bands in the orders (178). Their positive boundary successor is determined by the recorded ray orders and the directed edges joining successive visits. Map these labelled cells to the actual vertex discs and edge bands. This gives an orientation-preserving ribbon identification relative to the entire named spine, including each boundary cycle. Applying it in both diagrams extends (179) to an oriented ribbon map. Each boundary has its corresponding complementary disc. Subdivide its boundary map and extend by the positive triangle-fan construction of Lemma 3.11, then glue. The resulting named PL sphere homeomorphism h preserves ambient orientation, the oriented traversal and the over/under designation separately. These three kinds of orientation data are preserved separately by the construction.

Identify the two oriented spheres. Lemma 3.12 constructs a positive sphere isotopy h_t from the identity to this particular PL map h .

Lift each decorated diagram into a thin product neighbourhood of the sphere, with the over height larger than the under height at each crossing. $h \times \text{id}$ carries the source lift to a compatible target lift. Regard S^3 as the suspension of S^2 ; concretely $[x, s] \mapsto (\sqrt{1-s^2}x, s) \in S^3 \subset \mathbb{R}^4$. Then

$$[x, s] \mapsto [h_t(x), s] \quad (180)$$

is an ambient isotopy of S^3 , including the two suspension points, and agrees with the product map near the equator. The height function of the transported lift is continuous and has the target over/under order. Lemma 3.13 gives an ambient isotopy to the chosen target lift directly at that continuous regularity. Its parameter orientation is unchanged.

If there are no crossings, both shadows are embedded oriented circles. The map $\gamma'\overline{\Phi}\gamma^{-1}$ identifies their boundaries. Name left and right discs by ambient and traversal orientation, using Lemma 3.11; map left to left and right to right with the positive PL boundary extensions. This gives a PL oriented sphere map, to which the same sphere isotopy and fibrewise lift argument applies. No empty record is sent through the crossing-labelled ribbon construction.

The sphere map may send the original point at infinity into a different complementary face. This is harmless: the isotopy constructed above lives in S^3 , not in the plane with infinity fixed. Therefore this proof gives oriented knot equivalence without claiming a plane isotopy that preserves the outer face. $\square$

Proof of Corollary 3.16

Proof of 3.16. The common cyclic word identifies the visits and the successor permutation before smoothing. Exchanging the same selected successors in that permutation gives the same cycles afterwards. This is an explicit correspondence and needs no ambient knot-equivalence theorem. Noninterlacement ensures that each selected pair separates two carriers: cut at one selected pair; each other selected pair lies entirely on one of the two resulting arcs, since it does not interlace that pair. Apply the same argument recursively on those arcs. It gives $|S| + 1$ cycles, with the two reconnections at any selected crossing on distinct cycles. Thus no carrier visits a selected smoothing point twice.

An original vertex appears once on the parent traversal; selected crossing points are distinct and avoid all vertices. The subsegments between successive vertices or crossing visits therefore have positive length. The occurrence μ_j lies on exactly one cycle. Its incoming and outgoing directions at the centre are positive multiples by (2); removing it merely fuses those segments. At another original corner the determinant is a positive multiple of the nonzero original turn determinant. At a selected smoothing corner it is a positive multiple of either $\det(d_e, d_f)$ or $\det(d_f, d_e)$ at that transverse crossing. These exhaust the corners. They prove the stated nonzero turns and absence of antiparallel corners, including in the deletion. A deletion carrier has at least three corners: one nonzero segment cannot close, and a closed two-segment polygon has antiparallel directions. No such deletion carrier occurs.

An unselected crossing is a self-crossing of a carrier exactly when both its visits belong to that cycle. The common successor permutation preserves this condition and the two visits' inherited positive over/under bits. Tracing inherited straight subsegments introduces no other self-intersection: every intersection was already a parent crossing, and a selected crossing lies on two different carriers, as proved above. The nonzero determinants just listed have constant signs on a smaller interval. Replacing either old incident direction by the fused direction multiplies the relevant determinant by a positive scalar. Hence all surviving corner signs and all retained crossing data agree among the configurations in (ii).

For completeness, the rotation comparison uses limits of every affected turn, not equality of finite-parameter angles. For a closed polygonal carrier with nonzero segments and no antiparallel consecutive directions, let $\vartheta_k \in (-\pi, \pi)$ be its principal turns. The successive unit directions satisfy $q_{k+1} = e^{i\vartheta_k} q_k$. Multiplying around the closed cycle gives

$$e^{i\sum_k \vartheta_k} = 1, \quad \frac{1}{2\pi} \sum_k \vartheta_k \in \mathbb{Z}. \quad (181)$$

Principal turns are continuous when the two incident directions are nonzero and not antiparallel. On the carrier through μ_j its distinguished turn tends to zero, and each other principal turn tends to the corresponding deletion turn. On every other carrier all turns have their common deletion limits. Thus each side's signed rotation tends to the deletion's integer rotation. An integer with distance less than 1/2 from that integer must equal it. There are finitely many supports and carriers, so a common sufficiently small interval works for all of them. This proves the signed and absolute rotation equalities in (ii).

It remains to evaluate the selector. Let c be the number of surviving corners in the distinguished carrier. Their signs agree in all three configurations. Its extra turn is right or left according to the named side. The exhaustive table is

surviving turns	W_{right}	W_{left}	W_{del}
mixed	0	0	0
all right	1	0	1
all left	0	$(-1)^{c+1}$	$(-1)^c$

(182)

In the last row $0 - (-1)^{c+1} = (-1)^c$; the other two rows have the same first-minus-second identity immediately. All other carriers retain all their corner signs and counts. This proves (iii). $\square$

Proof of Lemma 3.17

Proof of 3.17. Adjoin a formal i with $i^2 = -1$ and put $T_G = \mathbb{Z}[i][l^{\pm 1}, m^{\pm 1}]$ and $R_G = \mathbb{Z}[i][a^{\pm 1}, z^{\pm 1}]$. Since $1, i$ is a free integer basis of $\mathbb{Z}[i]$, every element of T_G is uniquely $u + iv$ with $u, v \in T$, and the inclusions $T \subset T_G$, $R \subset R_G$ are injective.

Uniqueness over T_G . Let G be a map from oriented link diagrams to T_G , invariant under planar isotopy and the three moves, with $G_{\bigcirc} = 1$ and $lG_{D_+} + l^{-1}G_{D_-} + mG_{D_0} = 0$ on every skein triple. Write $G_D = U_D + iV_D$ with $U_D, V_D \in T$. Both coordinates inherit the invariance, and since the three coefficients l, l^{-1}, m lie in T , comparing coordinates gives the source skein for U and for V separately, with $U_{\bigcirc} = 1$ and $V_{\bigcirc} = 0$. The skein is linear, so U and $F + V$ are both invariant T -valued maps satisfying the source skein with unknot value 1. Literature input 3 gives $U = F$ and $F + V = F$, hence $V = 0$ and $G = F$. No division by 2 is used.

Transport. The assignments

$$\phi : T_G \rightarrow R_G, \quad l \mapsto ia, \quad m \mapsto -iz, \quad \psi : R_G \rightarrow T_G, \quad a \mapsto -il, \quad z \mapsto im, \quad (183)$$

both fixing i , send the Laurent generators to units and therefore define ring homomorphisms. On generators, $\psi(\phi(l)) = i(-il) = l$, $\psi(\phi(m)) = -i(im) = m$, $\phi(\psi(a)) = -i(ia) = a$ and $\phi(\psi(z)) = i(-iz) = z$, so ϕ and ψ are mutually inverse isomorphisms. Let $\tilde{Q}$ be a map as in the statement, regarded as R_G -valued, and put $\tilde{Q}_D = \psi(Q_D) \in T_G$. It is invariant under planar isotopy and the three moves, and $\tilde{Q}_{\bigcirc} = 1$. Since $\psi(a^{-1}) = (-il)^{-1} = il^{-1}$, applying ψ to $aQ_{D_+} - a^{-1}Q_{D_-} - zQ_{D_0} = 0$ gives

$$(-il)\tilde{Q}_{D_+} - (il^{-1})\tilde{Q}_{D_-} - (im)\tilde{Q}_{D_0} = 0, \quad (184)$$

and multiplying by the unit i gives $l\tilde{Q}_{D_+} + l^{-1}\tilde{Q}_{D_-} + m\tilde{Q}_{D_0} = 0$. By the first step, $\tilde{Q} = F$, hence $Q_D = \phi(F_D)$ for every diagram D . Two maps as in the statement are therefore equal in R_G , and by injectivity of $R \subset R_G$ equal in R . $\square$

Proof of Theorem 3.18

Proof of 3.18. Start in $R_G = \mathbb{Z}[i][a^{\pm 1}, z^{\pm 1}]$ with $i^2 = -1$. The elements ia and $-iz$ are units, so they define a Laurent-ring homomorphism and a first evaluation

$$\phi(l) = ia, \quad \phi(m) = -iz, \quad P_D^G = \phi(F_D). \quad (185)$$

Inverting the first unit gives $(ia)^{-1} = -ia^{-1}$. The image of the source skein is therefore

$$iaP_+^G - ia^{-1}P_-^G - izP_0^G = 0. \quad (186)$$

Multiplying by $-i$ gives the skein in (25). For the source initialization, the numerator $l + l^{-1}$ maps to $i(a - a^{-1})$ and the denominator m maps to $-iz$. Their quotient with its preceding minus sign is $(a - a^{-1})/z$. Thus $P_D^G = \delta^{c-1}$ for every UNDER-first diagram. This calculation does not yet assert integral coefficients for arbitrary diagrams.

Choose auxiliary basepoints and a component order. Let $N(D)$ be the crossing count and let $b(D)$ count first encounters which are OVER. The initialization is exactly $b = 0$. A switch at the first bad crossing keeps the parameter circles and their traversal order. It changes that first encounter from OVER to UNDER and leaves every other first-encounter designation unchanged. Thus D^{sw} has complexity $(N, b - 1)$. Its chosen crossing sign is negated, since the two ordered over/under tangents are exchanged.

Oriented smoothing reconnects each incoming strand to the other strand's outgoing end. It removes just the selected crossing and creates no other crossing in its clean disc. A self crossing splits one parameter circle into two; a mixed crossing joins two into one. This can be checked by writing a self-crossing cycle as $(xAyB)$ and replacing it by the two cycles $(xB), (yA)$ before erasing x, y . For a mixed crossing, the cycles $(xA), (yB)$ become $(xB y A)$ before erasing the marks. Empty resulting occurrence words are still circles. Thus the resulting component counts are $c + 1$, or $c - 1 \geq 1$ in the mixed case, and the smoothing always remains in the nonempty domain. Its crossing count is $N - 1$. New orders and basepoints can be chosen because the first complexity coordinate has decreased. These facts justify lexicographic induction on (N, b) , for every finite number of components.

At a positive chosen crossing the skein is $aP_D^G - a^{-1}P_{D^{\text{sw}}}^G = zP_{D^0}^G$. Move the second term to the right and divide by the unit a to obtain

$$P_D^G = a^{-2}P_{D^{\text{sw}}}^G + a^{-1}zP_{D^0}^G. \quad (187)$$

At a negative chosen crossing the positive diagram is D^{sw} , so the skein is $aP_{D^{\text{sw}}}^G - a^{-1}P_D^G = zP_{D^0}^G$. Move the latter two terms across and multiply by a . This gives

$$P_D^G = a^2P_{D^{\text{sw}}}^G - azP_{D^0}^G. \quad (188)$$

For integral descent let $M_c = z^{1-c}\mathbb{Z}[a^{\pm 1}, z^2]$, viewed as an additive subgroup of R_G . The initialization belongs to M_c , since $\delta^{c-1} = z^{1-c}(a - a^{-1})^{c-1}$. The switched term has the same component count, so induction puts it in M_c ; multiplication by any a unit preserves this subgroup. For a self smoothing, induction puts the smoothed value in M_{c+1} , and

$$zM_{c+1} = z z^{-c}\mathbb{Z}[a^{\pm 1}, z^2] = M_c. \quad (189)$$

For a mixed smoothing, induction puts it in M_{c-1} , and

$$zM_{c-1} = z z^{2-c}\mathbb{Z}[a^{\pm 1}, z^2] = z^2M_c \subseteq M_c. \quad (190)$$

Both recurrences therefore put P_D^G in M_c . The standard inclusion $R \hookrightarrow R_G$ is injective (the Gaussian coefficient ring has the free integer basis $1, i$). There is consequently one $P_D \in R$ with image P_D^G . This proves support and integral descent simultaneously. The local source equalities remain equalities after ϕ and, by injectivity, are equalities in R .

For uniqueness, let $Q_D \in R$ satisfy the stated skein and every UNDER-first initialization. Apply the same induction. At $b = 0$ its value is the specified $\delta^{c-1} = P_D$. Otherwise its value is given by the same positive or negative solved recurrence. The switched and smoothed values agree with P by the inner and outer induction respectively. Substitution gives $Q_D = P_D$. No existence or confluence is inferred from this uniqueness argument; existence was the literal source construction.

Finally map R to $\mathbb{Q}(a)$ by fixing the indeterminate a and sending z to $a - a^{-1}$. This target is a field, and $a - a^{-1}$ is a nonzero rational function, so the Laurent specialization is well-defined. The initialized δ maps to 1. The same induction shows that every nonempty diagram specializes to 1. In the positive case the two inductively specialized smaller values are each 1, and

$$a^{-2} + a^{-1}(a - a^{-1}) = a^{-2} + 1 - a^{-2} = 1. \quad (191)$$

In the negative case they are also each 1, and

$$a^2 - a(a - a^{-1}) = a^2 - a^2 + 1 = 1. \quad (192)$$

A zero element of R would specialize to zero. This proves $P_D \neq 0$.

Finally, $D \mapsto P_D$ and $D \mapsto H_D$ are maps from oriented link diagrams into R that take the value 1 on the crossing-free circle, satisfy the skein (25), and are invariant under planar isotopy and the three Reidemeister moves: for P these properties were proved above from Literature input 2, and for H they are the existence and invariance clauses of Literature input 1. Lemma 3.17 identifies two such maps, so $P_D = H_D$ on every actual diagram. This identification consumes the uniqueness transported to R by that lemma from Literature input 3, together with the existence clause of Literature input 1; the uniqueness clause of the latter is not used. It does not derive arbitrary ambient-isotopy invariance from the local construction alone. The support proof itself used only the local construction and the displayed induction. $\square$

Proof of Lemma 3.20

Proof of 3.20. Choose an order on the components of D , and one basepoint outside the crossings on each component. Transfer the order to D' . On a component with crossings, choose its new basepoint in the interval just preceding the image of the old first crossing occurrence. On a crossing-free component choose any point. These choices make the complete traversal orders of crossing occurrences correspond. They are allowed because the source construction is independent of component order and basepoint.

Call a crossing bad if its first encounter in this ordered traversal is the overpass. Write N for the total number of crossings and b for the number of bad crossings. Both numbers agree for the two based diagrams. We use lexicographic induction on (N, b) . If $b = 0$, both diagrams are under-first ascending and have the same number $c \geq 1$ of components. Their two source values are therefore separately μ^{c-1} by the initialization. This includes $N = 0$ and any crossing-free components; no value for an empty link is introduced.

For $b > 0$, choose the first bad crossing in D and its named image in D' . Switching that crossing preserves each traversed component and its basepoint. It swaps the two over/under bits at just that crossing. Hence that crossing ceases to be bad and every other first-encounter designation is unchanged: the switched diagrams have complexity $(N, b - 1)$ and still have isomorphic decorated records. Their values are equal by the inner induction.

Oriented smoothing removes the two occurrences of the chosen crossing. Cut each traversed strand there into an incoming and an outgoing end. The oriented smoothing connects each incoming end to the outgoing end of the other strand. The given bijection preserves these ends and their orientation, so it preserves this precise recombination of successors. All other paired crossings retain their over/under bits and signs. The resulting actual smoothed diagrams thus have isomorphic decorated records, including their component bijection. A self-crossing splits one traversed component into two; a mixed crossing joins two into one. Crossing-free resulting circles are retained as components rather than discarded. The smoothed diagram is always nonempty: in the mixed

case the original component number is at least two, and in the self case the number increases. The new crossing number is $N - 1$.

Choose new component orders and basepoints on one smoothed diagram and transfer them through its induced record bijection as above. The outer induction applies regardless of their new bad-crossing count. It gives equal values for the two smoothed diagrams. No relation between the old and new component orders is needed, and no topological classification is used to supply the record bijection.

It remains to insert these two equalities into the same scalar recursion. The crossing signs are part of the record, so the chosen crossings in D, D' have the same sign. For a positive chosen crossing the source skein reads

$$lF_D + l^{-1}F_{D^{\text{sw}}} + mF_{D^0} = 0. \quad (193)$$

Subtract the last two terms and divide by the unit l . This gives

$$F_D = -l^{-2}F_{D^{\text{sw}}} - l^{-1}mF_{D^0}. \quad (194)$$

For a negative chosen crossing the positive diagram is the switched one, and the source skein instead reads

$$lF_{D^{\text{sw}}} + l^{-1}F_D + mF_{D^0} = 0. \quad (195)$$

Subtract the other two terms and multiply by l , obtaining

$$F_D = -l^2F_{D^{\text{sw}}} - lmF_{D^0}. \quad (196)$$

The identical formula with primes holds in the corresponding sign case. The two switched terms agree by the inner induction and the two smoothed terms by the outer induction. Substitution therefore proves $F_D = F_{D'}$. This completes the induction, including both signs and both smoothing component cases. No division by a polynomial value is performed.

A common Laurent-ring homomorphism preserves equality. Coefficient extraction is an additive map defined also at zero and therefore preserves that resulting equality. These give the last assertions. $\square$

Realizing clean marked joins

Here are sufficient planar constructions. Choose a point in a face adjacent to I_A , off the diagram, and regard the page as an oriented sphere. Use an orientation preserving chart omitting that point. The marked interval is now adjacent to the unbounded face; the finite decorated data have not changed, so Lemma 3.21 preserves its scalar. If a polygonal presentation is desired, the finite crossing-disc and disjoint-strip construction of Lemma 3.10 gives one without changing those data; its ambient conclusion is not used here. Carry the interval and its endpoint collars with the local page maps; this preliminary change of presentation is not claimed to fix their original coordinates.

In the unbounded face choose a simple polygonal path from a point just beside the interval to outside a rectangle containing the diagram. Such a finite path exists: the points reachable from a given point by finite polygonal paths in an open connected set form an open and relatively closed subset. A polygonal path can be made simple by a small perturbation and removal of its finitely many loops. At its initial end use the clean interval disc. Its remaining compact image is disjoint from the diagram and hence has positive clearance. A sufficiently thin ribbon about the path therefore misses the diagram; at its finitely many corners use disjoint small discs. Two parallel sides of this ribbon extend the cut ends to the exterior rectangle without crossing one another or

the diagram. No old crossing is changed and no new one is added. The remaining diagram is thus presented in a rectangle with the two oriented cut ends accessible from its boundary.

Shrink this rectangle by a positive affine map into the clean disc of I_B after cutting its smaller subinterval. Match the two ends to the two collars of B with disjoint arcs, orienting the inserted traversal from the incoming collar to the outgoing collar. There is no reversal of A : label its incoming and outgoing cut ends accordingly, and use an orientation preserving planar placement. The exterior of that clean disc in B is literally unchanged. This realizes the stated record concatenation. The old A presentation, the new one and any other actual clean realization are compared only by Lemma 3.21. In particular moving the omitted face has not been called a relative ambient isotopy fixing an arbitrary exterior. Any finite list of other marks can be kept out of the local crossing/join discs by shrinking them first; their cyclic positions and collars survive.

Proof of Theorem 3.23

Proof of 3.23. Use induction on the sum $N = N(A) + N(B)$, with an inner induction on the sum $b = b(A) + b(B)$ of bad crossings for chosen based orders. In EACH factor put the marked component first and base it at its marked gap. Choose arbitrary orders and bases for the remaining components.

If $b = 0$, traverse the joined component starting just before the A portion, then its B portion. Traverse the remaining A components in their chosen order, then the remaining B components in their chosen order. For crossings of A , the relative order of all visits is the same as in its factor traversal; visits of B inserted between some of them do not change which of the two visits to an A crossing comes first. The identical observation holds for B . There are no crossings with one visit in each factor. Thus the entire joined diagram is UNDER-first. Its source base gives

$$P_{J(A,B)} = \delta^{c(A)+c(B)-2} = \delta^{c(A)-1} \delta^{c(B)-1} = P_A P_B. \quad (197)$$

This is a scalar computation on an actual ascending diagram, not a claim identifying its marked arcs with a trivial tangle.

If $b > 0$, choose a bad crossing in one factor, say A . Its disc is disjoint from the mark and from the joining operation. Switching it gives A^{sw} with the same component number, same marked interval and smaller b . In $J(A, B)$ it is precisely the corresponding crossing switch. Smoothing gives A^0 , of crossing number $N(A) - 1$, with the same interval on one unique component. In the joint diagram, smoothing has exactly the record of $J(A^0, B)$: both operations are recombinations at disjoint incoming/outgoing ends. All other successor relations and all other crossings are unchanged. If the selected crossing involves the marked component, a self-smoothing places the interval on exactly one of its two output components, and a mixed smoothing joins that component to another while retaining the interval. If the selected crossing involves no strand of the marked component, that component and its interval are unchanged: a self-smoothing splits one unmarked component into two, and a mixed smoothing merges two unmarked components. In all four cases every output component is retained, including all crossing-free components. These are the exhaustive self/mixed and marked/unmarked cases; nonemptiness was checked in Theorem 3.18.

Use the actual smoothed diagram as the realization of $J(A^0, B)$, or apply Lemma 3.21 to a newly chosen clean realization. Inner induction gives $P_{J(A^{\text{sw}}, B)} = P_{A^{\text{sw}}} P_B$; outer induction gives $P_{J(A^0, B)} = P_{A^0} P_B$. Apply the appropriate solved skein to the joint diagram and substitute these two equalities. Its right side is respectively $(a^{-2} P_{A^{\text{sw}}} + a^{-1} z P_{A^0}) P_B$ or $(a^2 P_{A^{\text{sw}}} - a z P_{A^0}) P_B$. The factor in parentheses is P_A by the same solved skein in A . This proves the identity. If the selected bad crossing is in B , the identical argument with the two factor names interchanged applies. For

each outer induction case new auxiliary choices can again put its marked component first, because the source value is independent of those choices. $\square$

Proof of Lemma 3.25

Proof of 3.25. Use finite polygonal page models preserving all transverse crossing signs; the clean crossing-disc and strip construction cited above does not change any signed intersection. Write the oriented edges of component 2 as $[v_j, v_{j+1}]$. Choose an auxiliary point o off the finitely many lines and intersections that would make a fan triangle degenerate or a radial edge nongeneric against component 1. A finite union of proper lines and points cannot fill the plane, so such o exists. Include further finite forbidden lines through a vertex of component 1 and a vertex of component 2 to exclude radial-edge hits at the former vertices. Orient the triangle chain $[o, v_j, v_{j+1}]$ by this order. Its coefficient relative to the usual positive plane orientation is its determinant sign.

For one positively oriented triangle, a closed transverse path has as many entries into its interior as exits, counted with sign. At an entry the ordered tangent pair (path, triangle boundary) has the opposite determinant sign from an exit: the triangle lies consistently on the left of its positive boundary. Thus its algebraic intersection sum with that boundary is zero. Reversing triangle orientation negates the sum and still gives zero. Summing over the finitely many fan triangles, each radial edge occurs once in each orientation and cancels. Their boundary chain is exactly component 2. Hence the intersection sum of components 1 and 2 is zero. This remains valid for self-intersecting components: the argument counts traversed occurrences, not distinct points of an embedded boundary for either component.

At each mixed crossing the decorated sign is either the fixed-order determinant sign or its negative, according to which component is over. Their total differs from the zero fixed-order sum by twice an integer. Its half is therefore integral. If the over-component is constant, all signs use the same one of those two conventions and the total is zero. This proves the assertions using the original mixed crossings of this coambient pair, not a replacement split pair. $\square$

Proof of Lemma 3.27

Proof of 3.27. First fix a connected component A of the interlacement graph and a chord b outside it. Cutting the circle at the endpoints of b gives two open intervals. Since b interlaces no chord of A , each chord of A has both endpoints in one of them. If chords of A occupied both intervals, an interlacement path in A would have an adjacent pair in different intervals. Such chords cannot alternate. Thus all endpoints of A lie in just one interval, which means both endpoints of b lie in one cyclic gap between successive endpoints of A .

If b, b' are interlacing chords outside A , they must occupy the same gap of A : distinct gaps are disjoint intervals and cannot contain alternating endpoints. Applying this along an interlacement path shows that every other connected component B has all its endpoints in one gap of A . This argument needs connectedness of A as well as B ; noninterlacement of individual chords alone would not suffice.

Choose a base gap for the full cyclic record and let A be the block containing its first occurrence. The other blocks lie in gaps of A . Within a nonempty such linear gap take the block B containing its first occurrence. Every other block inside the interval lies either in an internal gap of B or after its last occurrence: it cannot precede its first occurrence, and the one-gap property excludes a block straddling two such gaps. Recurse on the internal gaps, and then on the suffix after that last occurrence. The same description for the outer cyclic gap is made by cutting at the chosen base gap. Each step removes at least one nonempty connected block from the pending list, so the

resulting finite nested forest, attached to root A , terminates.

Start with the actual C_A . For each gap attach the recursively constructed actual diagrams for that gap in their inherited order, using small disjoint crossing-free intervals in that gap. Each clean join concatenates just the chosen cyclic blocks and introduces no new crossing. Multiple insertions can use successively smaller intervals; there are only finitely many. The construction of marked joins above supplies actual planar realizations at every step, including nested and exterior gaps. Thus the final oriented named record is literally the full given record, with no assumption that an arbitrary restricted word is realizable. Every leaf was an already supplied actual diagram.

Repeated use of Theorem 3.23 gives the product value for the constructed diagram. Lemma 3.21 equates it with any actual diagram having the same full data. Every old crossing is present once, with its old sign, and no joining arc has a crossing, so writhe adds. In the canonical application each sign is positive; its sum is $\sum_H |H|$. If there are no owned blocks, the grouped value remains the separately defined algebraic empty product 1 and empty sum 0. No empty link is introduced and no knot is asserted to represent that empty product by this lemma. $\square$

Proof of Lemma 3.30

Proof of 3.30. Fix a labelled weakly generic tuple P and, by Lemma 1.17, a relative neighbourhood N of P in $\mathcal{W}_n$ on which the visible signature is constant. On N the original vertices and the crossings carry fixed names. For a transverse crossing the two edge parameters solve a two-by-two linear system with nonzero determinant and are continuous functions of the parent vertices; after shrinking N their strict order on each edge and their membership in the open edge intervals persist. The marked successor data, the independent sets and the cycle lists of Lemma 3.7 are therefore the same for every tuple in N , and every carrier has a fixed finite corner list — original vertices and selected crossing points — whose coordinates depend continuously on the parent.

At P each such list lies in the regular locus of its own arity, by Lemma 3.7(ii). There are finitely many independent sets and carriers; for each corner list choose a ball about its value at P inside that regular locus, which is open, and shrink N so that the continuous image of N lies in the ball. A ball is convex, so Lemma 1.20(ii) applied to straight paths in it makes the rotation of every carrier constant on N . All corner determinants are nonzero at P by Lemma 3.7(ii); shrink N once more so that their signs are constant. Hence the uniform/mixed classification, every r_Q , every m_Q (the number of unselected crossings with both visits on the cycle, fixed with the cycle lists) and $\ell(P)$ are constant on N .

The retained double points of a carrier are exactly the unselected crossings named by its cycle list, in a fixed traversal order. At each of them the two strand vectors are positive multiples of the corresponding parent edge vectors, whose nonzero determinant keeps its sign on N ; so the over strand of the positive lift is the same named strand throughout N , and every crossing sign is $+1$. The named decorated record of each actual positive carrier diagram — oriented cyclic successor, crossing pairing, over/under bits and signs, including the one-component record of a carrier without crossings — is thus the same at every tuple of N . Lemma 3.20 gives equality of the source polynomials of corresponding carrier diagrams, the common Laurent substitution of Theorem 3.18 gives equality of their P -polynomials, and that theorem identifies these with their H -polynomials. Since d_Q is constant, every coefficient $c(Q)$ is constant on N , and so is the finite sum of Definition 3.29. This proves local constancy on the labelled locus, at simultaneous silent zeros as well as at simple ones: no factorization of a path into wall crossings was used.

Cyclic relabelling renames the crossings, the independent sets and the traced oriented carrier diagrams without changing their turns, rotations, positive lifts or coefficients, or $\ell(P)$; so $C(\sigma P) = C(P)$. The union of the cyclic translates of a neighbourhood on which C is constant is a saturated

open set on which C has that value, so C is locally constant on the quotient as well. A locally constant function is constant on a connected set, since the preimage of a value and its complement are both open; this gives constancy on every visible chamber.

On $\mathcal{U}_n$ the definition agrees with Definition 3.28 term by term. If C' is another locally constant function on $\mathcal{W}_n$ with that restriction, then at any $P \in \mathcal{W}_n$ the intersection of neighbourhoods on which C and C' are constant meets $\mathcal{U}_n$, by the density assertion of Lemma 1.17, and at a generic point of it $C = C'$; hence $C'(P) = C(P)$. $\square$

Proof of Lemma 3.42

Proof of 3.42. Omitting spectator strands, label one left endpoint L and three right endpoints 1, 2, 3 from top to bottom. Put

$$A = l_2\sigma_1, \quad A' = l_1\sigma_2, \quad C_{\text{top}} = l_1, \quad C_{\text{bottom}} = l_2. \quad (198)$$

Both A and A' pair L with 2 and 1 with 3. The L -2 strand is the through-strand; the other contains the cusp. In A the through-strand is over, whereas in A' the cusp strand is over. Each strand encounters the one crossing once. Exchanging over and under at that crossing makes their local ordered crossing records identical, including the boundary attachments. After gluing any same actual exterior, the full named records are identical. Thus their polynomial values are the two crossing choices at one ordinary skein site. The move before that over/under exchange is not asserted to be page isotopy.

Let $t = 1$ for the orientation from L to 2, and $t = -1$ otherwise. Let $u = 1$ for the orientation from 1 to 3, a downward cusp, and $u = -1$ otherwise. In A , the cusp's upper arm has x direction $-u$; in A' , the cusp's lower arm has x direction u . The crossing determinant rule gives signs $-tu$ and tu , respectively:

(t, u)	$\text{sign}(A)$	$\text{sign}(A')$	Compatible smoothing	Local D in A, A', C
$(+, +)$	-1	$+1$	C_{top}	1
$(+, -)$	$+1$	-1	C_{bottom}	0
$(-, +)$	$+1$	-1	C_{bottom}	1
$(-, -)$	-1	$+1$	C_{top}	0

Oriented smoothing joins each incoming strand to the other outgoing strand. If $t = u$, it pairs 1 with 2 and L with 3: the top cusp and through-strand. If $t = -u$, it pairs L with 1 and 2 with 3: the bottom cusp and through-strand. The new cusp direction is u in either case, from 1 to 2 or from 2 to 3. This proves the last column without any hypothesis on how the exterior joins the two strands. Smoothing may split or join parameter components, but cannot create an empty diagram. Only the one compatible smoothing is used, never both unoriented smoothings.

Let E_+, E_-, E_0 be these actual ordinary skein diagrams. Their writhes are $w_0 + 1, w_0 - 1, w_0$, where w_0 is the common exterior writhe. Solving the skein relation separately in the two directions gives

$$P_+ = a^{-2}P_- + a^{-1}zP_0 \quad (199)$$

and

$$P_- = a^2P_+ - azP_0. \quad (200)$$

The degree of a nonzero sum is at most the maximum of the summand degrees. Multiplication by a^k shifts degree by k , and multiplication by z does not change it. Equation (199) yields $d_+ \leq \max\{d_- - 2, d_0 - 1\}$. Adding $w_0 + 2$ gives

$$d_+ + 1 + (w_0 + 1) \leq \max\{d_- + 1 + (w_0 - 1), d_0 + 1 + w_0\}. \quad (201)$$

Equation (200) instead yields $d_- \leq \max\{d_+ + 2, d_0 + 1\}$. Adding w_0 gives

$$d_- + 1 + (w_0 - 1) \leq \max\{d_+ + 1 + (w_0 + 1), d_0 + 1 + w_0\}. \quad (202)$$

The table and unchanged exterior give the same downward-cusp count in all three diagrams. Subtracting (201) and, respectively, (202) from this common number proves

$$\begin{aligned} B(A) &\geq \min\{B(A'), B(C)\}, \\ B(A') &\geq \min\{B(A), B(C)\}. \end{aligned} \quad (203)$$

Here C is the compatible smoothing. It has one fewer crossing and the same number of cusps, so $s(C) = s(A) - 1 = s(A') - 1$; the principal crossing-change branch preserves s . Reflected and right-cusp templates follow by relabeling their actual attachments and arrows in this local calculation, not from a global reflection equality for P . $\square$

Details of the finite front-word reduction

For precision, the finite cases and index transcription used in this source input are as follows. Write the word at its rightmost left cusp as $Xl_m Y$, where Y contains only crossings and right cusps. The two arm strings are $(\sigma_{m-1} \cdots \sigma_{m-N_1})(\sigma_{m+1} \cdots \sigma_{m+N_2})$. They use disjoint crossing positions and commute. The complete cases are listed below; each row includes both arm versions where applicable.

First letter of Y	Finite procedure	Progress or terminal
Crossing outside the strings and their next neighbors	Commute into X	Fewer right-hand singularities
Crossing at a next neighboring position	Extend that arm string	N_1 or N_2 increases, bounded by active strand count
Crossing at an occupied outer string endpoint	Move the cusp along its string using skein interchanges	Expose type II and remove two crossings
Crossing strictly inside a string	Type III, then commutations	Move the crossing into X
Middle crossing, exactly one nonempty string	Expose type II	Remove two crossings
Middle crossing, no strings	Apply the $l_i \sigma_i$ shortcut	Lower s by Lemma 3.40
Middle crossing, both strings nonempty	Cusp interchange, type III and commutations	Move the crossing into X
Right cusp beyond an arm's next neighbor	Commute into X with the correct two-strand index shift	Fewer right-hand singularities and active strands
Right cusp at a next neighbor	Move the selected cusp along the arm string	Expose a zigzag, then delete it
Right cusp at a string endpoint	Expose $\sigma_i r_i$	Crossed-cusp shortcut
Right cusp strictly inside a string	Commute to expose type II	Remove two crossings
Middle right cusp, neither string nonempty	Standard circle without mixed crossings	Delete, or stop at the last component
Middle right cusp, exactly one string nonempty	Type I	Remove one crossing and two cusps

First letter of Y	Finite procedure	Progress or terminal
Middle right cusp, both strings nonempty	Cusp interchange, then commuting	Expose type II and remove two crossings

Right-cusp absorption lowers active strand count by two; crossing absorption lowers the right-hand crossing count. String extension is bounded by active strand count and cannot continue without eventually reaching another row. These are the source's nested finite descents. Only deletion directions of types I and II occur here; type III, commutations and cusp-skein changes preserve s . Stabilized cusp cases use the direct shortcut of Lemma 3.40 (display (35); $\Delta s = -1$), so no temporary increase of s is needed.

Two impossible printed ranges $m - N_1 > i > m$ on Rutherford's pp. 11–12 are transcribed here as the interior range $m - N_1 < i < m$, as the neighboring cases and actual strands require. The duplicated $m + 2$ in the increasing string on p. 11 continues with $m + 3$, with the corresponding two-strand shift on p. 12. The malformed literal strings are not imported as identities: the actual strand operations and Ng's Figure 1 specify the moves used here. Ng cites the published numbering Lemma 3.3; the retained Rutherford arXiv version numbers this proof Lemma 3.2.

No Rutherford ruling or Kauffman value, maximizing representative, or earlier general characterization of Legendrian equivalence is consumed. Literature input 4 is precisely the local constructive portion of those proofs. It is not an application of the literal statement of Ng's Theorem 1 to an ambient link invariant. The local polynomial/count certificates above and the following degree induction are separate arguments.

Proof of Theorem 3.43

Proof of 3.43. Use strong induction on $s(F)$; the quantity $d(F) = \deg_a P_{S(F)}$ of display (32) does not depend on the rounding $S(F)$, by Lemma 3.34. Every nonempty component has at least one left and one right cusp. Thus the base $s = 2$ is a standard circle, and Lemma 3.41 gives $B = 0$. Any union of standard front circles encountered as a terminal base also has $B = 0$. An empty object is never evaluated.

Suppose the assertion holds for all fronts with smaller s than F . Take the finite principal chain of Literature input 4, stopped at its first strict decrease or at a standard-circle base. Every preterminal principal front has s at most $s(F)$. At each cusp-skein step, its compatible smoothing therefore has strictly smaller s than F , and strong induction gives nonnegative B for it. The final principal front has nonnegative B by the same induction or by its explicit standard-circle base.

Traverse the chain backwards. Commutations and local front moves preserve B by Lemmas 3.36, 3.37, 3.38 and 3.39. A deletion cannot increase B in the forward direction, by Lemmas 3.40 and 3.41; nonnegativity after deletion therefore implies nonnegativity before it. At a cusp-skein step, the later principal branch and its smaller smoothing branch both have nonnegative B . The appropriate inequality of Lemma 3.42, display (203), gives nonnegative B for the earlier principal branch. These are all possible operations. Finite backward propagation yields $B(F) \geq 0$, completing strong induction. Substituting (32) gives (37).

Nonvanishing is used before every degree, including both skein branches and the split product. Cancellation in sums was allowed. The proof covers arbitrary positive component counts and all orientations. $\square$

Proof of Lemma 3.45

Proof of 3.45. Write $T = (x, y, z)$, $a = \alpha(T') > 0$, and $c = \sqrt{2a}$. The global frame $E_1 = \partial_x + y\partial_z$, $E_2 = \partial_y$ spans $\ker \alpha$, and $d\alpha(E_1, E_2) = 1$. Set

$$F(\theta, u, v) = T(\theta) + c(\theta)(uE_1(T(\theta)) + vE_2(T(\theta))). \quad (204)$$

At the core its derivative has independent columns T', cE_1, cE_2 : applying α to a relation first kills the T' coefficient. The inverse function theorem and compactness give a common radius on which every derivative is invertible. This radius can be decreased so that F is injective: otherwise two distinct points with disc radii tending to zero have equal images; limiting circle parameters coincide because T is embedded, placing both points in one local inverse chart, a contradiction. Thus F embeds a product neighbourhood.

Put $\beta = F^*\alpha/a(\theta)$. At the core $\beta = d\theta$ and

$$d\beta(\partial_u, \partial_v) = a^{-1}d\alpha(cE_1, cE_2) = c^2/a = 2. \quad (205)$$

The derivative of a^{-1} vanishes on this vertical pair. Consequently $\alpha_t = (1-t)\alpha_0 + t\beta$ satisfies $(\alpha_t \wedge d\alpha_t)(\partial_\theta, \partial_u, \partial_v) = 2$ on the core. Compactness of $[0, 1] \times S^1$ supplies a common smaller neighbourhood where every α_t is contact and $\alpha_t(\partial_\theta) > 0$.

Let $\nu = \beta - \alpha_0$ and define

$$X_1 = \partial_u - \frac{\alpha_t(\partial_u)}{\alpha_t(\partial_\theta)}\partial_\theta, \quad X_2 = \partial_v - \frac{\alpha_t(\partial_v)}{\alpha_t(\partial_\theta)}\partial_\theta, \quad (206)$$

$$D_t = d\alpha_t(X_1, X_2), \quad V_t = \frac{\nu(X_1)X_2 - \nu(X_2)X_1}{D_t}. \quad (207)$$

The X_i form a basis of $\ker \alpha_t$; contactness makes $D_t \neq 0$. Evaluation on each X_i gives

$$(\iota_{V_t}d\alpha_t)|_{\ker \alpha_t} = -\nu|_{\ker \alpha_t} \quad (208)$$

The vector field V_t vanishes on the core, since ν does. Choose $\rho > 0$ so that $S^1 \times \overline{D}_\rho$ lies strictly inside the common neighbourhood where V_t is smooth for every $t \in [0, 1]$. On this closed product its norm is at most $C\sqrt{u^2 + v^2}$ uniformly in t, θ . Along a trajectory, $(u^2 + v^2)' \leq 2C(u^2 + v^2)$, hence $r(t) \leq e^{Ct}r(0)$. Choosing the initial disc radius below $\rho e^{-C}/2$ keeps the flow in a compact inner product neighbourhood for $0 \leq t \leq 1$. Smooth ODE existence, continuation and uniqueness give a flow Φ_t , fixing the core and invertible onto its image.

The covector $\nu + \iota_{V_t}d\alpha_t$ annihilates $\ker \alpha_t$, so it is $\mu_t\alpha_t$ for a smooth scalar μ_t . Since $\alpha_t(V_t) = 0$, Cartan's differentiation identity gives

$$\begin{aligned} \frac{d}{dt}\Phi_t^*\alpha_t &= \Phi_t^*(\nu + \mathcal{L}_{V_t}\alpha_t) \\ &= (\mu_t \circ \Phi_t)\Phi_t^*\alpha_t. \end{aligned} \quad (209)$$

Solving this scalar equation gives $\Phi_1^*\beta = \exp(\int_0^1 \mu_t \circ \Phi_t dt)\alpha_0$. For $H = F \circ \Phi_1$, the additional factor $a \circ \text{pr}_\theta \circ \Phi_1$ is also positive, proving (38), including coorientation and core orientation.

Choose an integer $N > 0$ with $b = N^{-1/2} < \delta$. The embedded circle

$$L_0(\theta) = (\theta, b \cos N\theta, -b \sin N\theta) \quad (210)$$

has $\alpha_0(L'_0) = 1 - Nb^2 = 0$; its degree-one first coordinate makes it an embedding. Hence $L = H \circ L_0$ is Legendrian. Fix $\kappa \neq 0$ and a small $\epsilon > 0$ with $b + \epsilon < \delta$, and put

$$B(\theta, s) = (\theta + \kappa s, (b - s) \cos N\theta, -(b - s) \sin N\theta), \quad -\epsilon < s \leq b. \quad (211)$$

Radius $b - s$ and the first coordinate prove injectivity. The radial derivative proves immersion for $s < b$; at $s = b$ it remains independent of ∂_θ , and the Cartesian formula is smooth there. Direct differentiation yields

$$B^* \alpha_0 = (1 - N(b - s)^2) d\theta + \kappa ds. \quad (212)$$

Thus the annulus is transverse to the contact planes, even along $s = 0$. Its circles are negative transverse for $s < 0$, Legendrian at $s = 0$, and positive transverse for $0 < s \leq b$. A sufficiently small positive circle is therefore the positive pushoff in the source convention. The whole family from that circle to $s = b$ is a transverse isotopy to $T(\theta + \kappa b)$, which is the same oriented transverse knot. Finally $H \circ B$ on $S^1 \times [0, b]$ is an embedded oriented annulus with boundary $L - T$. To prove the stronger ambient-isotopy conclusion needed by the polynomial, we give such an isotopy explicitly. Write $v(\theta) = b(\cos N\theta, -\sin N\theta)$, choose $\epsilon_0 = (\delta - b)/4$, and put $R_1 = b + \epsilon_0$, $R_2 = b + 2\epsilon_0$. Thus $b < R_1 < R_2 < \delta$. For $\eta(s) = e^{-1/s}$ when $s > 0$ and $\eta(s) = 0$ when $s \leq 0$, set

$$\chi(s) = \frac{\eta(R_2^2 - s)}{\eta(R_2^2 - s) + \eta(s - R_1^2)}. \quad (213)$$

Every positive-side derivative of η is $e^{-1/s}$ times a polynomial in $1/s$, tending to zero at zero; hence η is smooth. The denominator is positive because $R_1 < R_2$. Thus χ is smooth, equals one for $s \leq R_1^2$ and zero for $s \geq R_2^2$. The smooth product vector field

$$V(\theta, w) = (0, -\chi(|w|^2)v(\theta)) \quad (214)$$

is well defined on the circle since N is integral. Squared radius makes it smooth at $w = 0$; its support lies in $S^1 \times \overline{D}_{R_2}$. The embedding H has an open image U and a smooth inverse there by the inverse function theorem. Push V forward by H on U and extend by zero outside U . This field X is globally smooth: its support lies in the compact set $K = H(S^1 \times \overline{D}_{R_2}) \subset U$, and it vanishes on a neighbourhood of every point outside K .

The field X is bounded, so a trajectory moves a distance at most $M|t|$ on any finite time interval. Smooth ODE continuation gives its global flow Ψ_t for every real t . Uniqueness gives the smooth inverse Ψ_{-t} ; thus these are ambient diffeomorphisms. They fix every point outside K . Their derivative determinants start positive and never vanish, so they preserve ambient orientation. Extending them by the identity at infinity gives an ambient isotopy of the oriented three-sphere, with the same compact support.

For $0 \leq t \leq 1$, the path $(\theta, (1 - t)v(\theta))$ stays at radius at most $b < R_1$, where $\chi = 1$, and its time derivative is exactly V . The chain rule and ODE uniqueness therefore give

$$\Psi_t(L(\theta)) = H(\theta, (1 - t)v(\theta)), \quad \Psi_1(L(\theta)) = T(\theta). \quad (215)$$

The parameter θ is unchanged, so the endpoint circle orientation is exactly that of T . The core need not be fixed: its own track is $\Psi_t(T(\theta)) = H(\theta, -tv(\theta))$ on this interval. There is thus no collapse of two distinct curves by a diffeomorphism. This is an ordinary ambient isotopy, not a contact isotopy. Indeed the model form on its tracked knot has value $1 - Nb^2(1 - t)^2 = 2t - t^2 > 0$ for $0 < t \leq 1$, whereas the initial knot is Legendrian. The preceding twisted annulus, not this ambient extension, remains the positive-pushoff and transverse-isotopy producer. $\square$

Proof of Theorem 3.48

Proof of 3.48. Write $L = (x, y, z)$, so $z' = yx'$. For finitely many compactly supported Hamiltonians, let Φ_a be the composition of their time- a_i flows. At $a = 0$, the a_i derivative of $\Phi_a \circ L$ is

$X_{H_i} \circ L$. The path $s \mapsto \Phi_{sa}$ is an ambient contact isotopy by Lemma 3.47; its images are exactly Legendrian embeddings, without an approximation or an area-closing correction.

First remove simultaneous zeros of (x', x'') . At such a point θ_0 , immersion and $z' = yx' = 0$ imply $y'(\theta_0) \neq 0$. With $y_0 = y(\theta_0)$, choose Hamiltonians equal near $L(\theta_0)$ to

$$H_2 = -\frac{1}{2}(y - y_0)^2, \quad H_3 = -\frac{1}{6}(y - y_0)^3, \quad (216)$$

and multiply them by a smooth bump equal to one there. Their infinitesimal x changes are $y - y_0$ and $(y - y_0)^2/2$. Thus their derivative columns for (x', x'') at θ_0 are $(y', y'')^T$ and $(0, (y')^2)^T$, with determinant $(y')^3 \neq 0$. Finitely many such minor neighbourhoods cover the initial compact zero set. On their compact complement $|(x', x'')|$ has a positive minimum. Continuity therefore supplies a small parameter ball in which every possible zero retains one of the nonzero minors. Apply Lemma 3.46 with $d = 1, q = 2$ to choose a small parameter avoiding all these zeros. If the initial zero set is empty, this step is unnecessary. Call the result L_1 . Every zero of x'_1 is simple, so compactness makes their number finite. At each of them $y'_1 \neq 0$ by immersion and Legendrianity.

We next supply a uniform collar of the parameter diagonal. Cover S^1 by finitely many interiors of closed intervals of two kinds: regular intervals where x'_1 has fixed nonzero sign, and critical intervals where x''_1 and y'_1 have fixed nonzero signs and the endpoint signs of x'_1 are opposite. On the first kind x_1 is strictly monotone. On a critical interval there is one critical point θ_c . For two points on opposite sides with equal x coordinate X , the inverse branches and $dz = y dx$ give

$$z_1(\theta_+) - z_1(\theta_-) = \int_{x_1(\theta_c)}^X [y_1(\theta_+(v)) - y_1(\theta_-(v))] dv \neq 0. \quad (217)$$

The integrand has a fixed strict sign by monotonicity of y_1 ; the integration interval has nonzero length. The formula also handles a maximum by reversing the integration sign. The front is therefore injective on each interval. Retaining the finitely many derivative and endpoint-sign margins preserves this argument for every sufficiently C^2 -close Legendrian embedding. A Lebesgue number $\delta > 0$ of the finite interval cover gives uniform front injectivity for every pair of distinct parameters at circular distance less than δ . This Lebesgue number follows from compactness: a sequence of pairs of diameter tending to zero and contained in no cover interval would converge to a point inside one of those open intervals.

Let K_2 be the compact ordered-pair set with distance at least δ , and K_3 the compact ordered-triple set with every pairwise distance at least δ . Write $p_a = (x_a, z_a)$. The bad configurations are zeros of

$$C(\theta, \eta, a) = (x'_a(\theta), p_a(\theta) - p_a(\eta)) \in \mathbb{R}^3, \quad (218)$$

$$R(\theta, \eta, \tau, a) = (p_a(\theta) - p_a(\eta), p_a(\theta) - p_a(\tau)) \in \mathbb{R}^4. \quad (219)$$

At a zero of C for $a = 0$, use disjoint small balls about the two distinct spatial points. The preceding H_2 at the first point changes $x'(\theta)$ by $y'_1(\theta) \neq 0$ but changes neither x, z there. At the second point q , Hamiltonians equal locally to $H^x = -(y - y_q)$ and $H^z = 1$ have independent (x, z) velocities $(1, y_q)^T$ and $(0, 1)^T$. Their supports avoid the first point. Hence C has parameter rank three. At a zero of R , choose disjoint balls at the three distinct spatial points and use the same two Hamiltonians at the second and third points only. The two independent blocks give parameter rank four.

Choose finitely many such Hamiltonians covering both initial compact zero sets. A zero-free initial map has a positive norm minimum and needs no Hamiltonians. The nonzero minors persist near those zero sets; the positive norm on their compact complements excludes new zeros there for

small parameters. Shrink the parameter ball also to retain the preceding uniform local-injectivity margins. Apply Lemma 3.46 simultaneously with $(d, q) = (2, 3)$ and $(3, 4)$, and choose a parameter outside both bad projections. The resulting L_2 has no cusp on another branch and no triple point.

At any remaining double point both parameters are regular. Their y values differ, since otherwise all three spatial coordinates coincide, contradicting embeddedness. The two front tangents have determinant

$$x'(\theta)x'(\eta)(y(\eta) - y(\theta)) \neq 0. \quad (220)$$

Thus every double point is transverse and isolated by the inverse function theorem applied to $p(\theta) - p(\eta)$. Their ordered-pair set is closed in compact K_2 , so there are finitely many: an infinite set would accumulate at another zero, contradicting isolation. All properties so far persist under sufficiently small C^2 Legendrian perturbations. Indeed the local-injectivity margins persist, and the now nowhere-zero maps $C, R, (x', x'')$ have positive norm minima on their respective compact domains.

We finally make each cusp germ exactly semicubical, without a singularity-normal-form assumption. At an x -critical point, $y' \neq 0$ allows y as a local parameter. Write $x = f(y)$ and put

$$u = y - y_0, \quad A = \frac{1}{2}f''(y_0) \neq 0, \quad q(y) = f(y_0) + Au^2, \quad \psi(y) = \int_{y_0}^y (f(v) - q(v)) dv. \quad (221)$$

Here $\psi = O(u^4)$ and its first three derivatives have orders $O(u^3), O(u^2), O(u)$. Where $H = \psi(y)$ the contact flow is exactly

$$(x, y, z) \mapsto (x - s\psi'(y), y, z + s(\psi(y) - y\psi'(y))). \quad (222)$$

At time one, $x = q(y)$ and Legendrianity gives

$$x = x_0 + Au^2, \quad z = z_0 + Ay_0u^2 + \frac{2}{3}Au^3. \quad (223)$$

The invertible affine page coordinates $X = (x - x_0)/A$ and $Z = 3(z - z_0 - y_0(x - x_0))/(2A)$ give $(X, Z) = (u^2, u^3)$.

Multiply $\psi(y)$ by a smooth spatial cutoff equal to one on an inner ball about the cusp and supported in a ball of radius $O(\epsilon)$ meeting no other portion of the embedded knot. Its derivatives of order k are $O(\epsilon^{-k})$, so the derivatives of H through order three are $O(\epsilon^{4-k})$. The vector field and its first two derivatives are consequently $O(\epsilon)$ on the support. The time-one flow is C^2 -close to the identity: differentiating its ODE gives $J' = DX_H J$ and $M' = DX_H M + D^2 X_H [J, J]$, with initial values $I, 0$, and the elementary exponential estimate bounds $J - I, M$ by $O(\epsilon)$. An inner subarc has displacement $O(\epsilon^3)$ and a gap of order ϵ to the boundary of the region where the cutoff equals one. Its whole flow therefore uses the exact displayed formula. Finitely many disjoint cusp balls and sufficiently small radii retain all previous genericity margins, including those in the transition regions. This gives the required exact cusp germs without new crossings.

Concatenate the constructed contact isotopies, using smooth time reparametrizations flat at the joins. Each of them is a composition of flows of Lemma 3.47, so $\Phi_s^* \alpha = c_s \alpha$ with $c_s > 0$ the product of their conformal factors. If B is the chosen thin pushoff annulus and K its positive parallel circle, then $\Phi_s \circ B$ remains a transverse annulus and

$$\alpha((\Phi_s \circ K)') = (c_s \circ K) \alpha(K') > 0. \quad (224)$$

Its central circle is the Legendrian isotope and its positive parallel circle is a positive pushoff thereof. Thus the chosen pushoff moves through positive transverse embeddings. The ambient isotopy also preserves the oriented topological knot type. No contact-isotopy extension or approximation-uniqueness theorem is used. $\square$

Proof of Lemma 3.49

Proof of 3.49. The denominator in (42) is bounded away from zero on the compact parameter torus. It remains so uniformly in any compact smooth family of disjoint pairs. All derivatives and differentiations under the integral are therefore ordinary smooth calculus on a compact domain. The integrand can equivalently be written

$$\frac{(C_1(u) - C_2(v)) \cdot (C_1'(u) \times C_2'(v))}{|C_1(u) - C_2(v)|^3}. \quad (225)$$

Indeed radial derivative terms disappear in the determinant, and the derivative in u has the sign from $-C_1'$. Exchanging the two curves negates both the difference vector and the cross product, leaving this scalar unchanged after exchanging the integration variables. This proves symmetry, including the source's possible ordering of the pushoff first.

For a smooth family $G(u, v, s)$ into the unit sphere set

$$A = G \cdot (G_u \times G_v), \quad B = G \cdot (G_s \times G_v), \quad C = G \cdot (G_u \times G_s). \quad (226)$$

The vectors G_u, G_v, G_s all lie in the two-dimensional plane orthogonal to G . Hence every scalar triple product of these three vectors vanishes. On expanding the three derivatives, the two terms containing G_{uv} cancel because the cross product is antisymmetric; the remaining terms match by equality of mixed derivatives. Explicitly,

$$\begin{aligned} \partial_s A &= G \cdot (G_{us} \times G_v + G_u \times G_{vs}), \\ \partial_u B + \partial_v C &= G \cdot (G_{su} \times G_v + G_s \times G_{vu} + G_{uv} \times G_s + G_u \times G_{sv}) \\ &= \partial_s A. \end{aligned} \quad (227)$$

Integrating a full period in u kills $\partial_u B$ by the fundamental theorem of calculus; integrating a full period in v kills $\partial_v C$. Thus the derivative of (42) is zero. This is a direct calculation, not an appeal to invariance of degree or to an ambient-isotopy theorem.

We verify the normalization used in a generic diagram. Let ω be the outward sphere area form divided by 4π . In positively oriented Cartesian coordinates (X, Y, Z) on the unit sphere, put

$$\lambda_N = -\frac{X dY - Y dX}{4\pi(1 - Z)} \quad \text{on } S^2 \setminus \{N\}, \quad N = (0, 0, 1). \quad (228)$$

This is smooth also at the south pole. In polar coordinates about N , it is $-(1 + \cos \theta)d\phi/(4\pi)$; differentiation gives $\sin \theta d\theta \wedge d\phi/(4\pi) = \omega$. On the positively oriented boundary of a cap of angular radius h , its integral is $-(1 + \cos h)/2$, tending to -1 as $h \rightarrow 0$.

Suppose N is a regular value of the particular map G under consideration. Its preimage is a finite set: the inverse function theorem makes it discrete, and it is closed in a compact torus. Choose disjoint inverse charts there. Compactness of their complement then gives a small cap whose full inverse image lies in those charts. Delete the inverse cap discs from the torus. On what remains, $G^*\omega = d(G^*\lambda_N)$. The boundary integral formula gives the integral as the sum of the integrals around the deleted discs with negative boundary orientation. A disc whose local orientation sign is σ contributes $\sigma(1 + \cos h)/2$. The omitted disc integrals tend to zero: the pullback form is smooth and their areas tend to zero in the fixed inverse charts. Consequently

$$\int_{S^1 \times S^1} G^*\omega = \sum_{p \in G^{-1}(N)} \sigma(p). \quad (229)$$

The empty-preimage case uses the primitive on the entire torus and has both sides zero. No assertion that an arbitrary value is regular was used. A positive rotation of the coordinates proves the same formula for any specified regular pole, including the opposite pole.

For completeness, the boundary integral formula just used requires only planar Green's formula. Cover the compact punctured torus by finitely many interior or boundary coordinate patches and choose smooth bump functions supported in the patches whose sum is positive there; division by that sum gives a finite partition of unity. Apply the planar formula to each supported one-form. For an interior patch the two integrals of partial derivatives vanish by the fundamental theorem and Fubini. In a boundary patch straighten the smooth boundary to a coordinate graph; the same one-variable integration leaves exactly its oriented boundary integral. The derivatives of the partition functions cancel because their sum is one. This proves the required formula without a topological degree theorem, general transversality theorem or source assumption.

Now let ν point toward the observer, and orient the projection plane so that its positive basis followed by ν is positive in $\mathbb{R}^3$. In a generic projection, a crossing between the components is exactly a preimage of ν or $-\nu$ under G ; the projected tangents are independent, so both poles are regular values, with the empty-preimage case allowed. At $G = \nu$, the second component is over. Up to a positive common factor the projected derivatives of G are $-C'_1, C'_2$. Their determinant in the tangent plane oriented by ν is

$$\det_{\nu}(-C'_1, C'_2) = \det_{\nu}(C'_2, C'_1), \quad (230)$$

the overpass-first crossing sign. At $G = -\nu$, the first component is over and the sphere tangent orientation is reversed. Its sign is therefore

$$\det_{-\nu}(-C'_1, C'_2) = \det_{\nu}(C'_1, C'_2), \quad (231)$$

again the overpass-first sign. Formula (229) at each pole says that each of these two signed crossing sums equals ℓ . Adding them and dividing by two proves the claimed mixed crossing formula. The argument concerns the generic diagrams actually used in the source computations, and does not invoke a theorem asserting genericity of every projection.

It remains to supply a uniform framed pair, rather than assume that pointwise small pushoffs assemble into a family. For the general framing assertion put $n_s = \partial_u C_s \times v_s$ and

$$F(s, u, r, t) = (s, C_s(u) + rv_s(u) + tn_s(u)). \quad (232)$$

At $(r, t) = (0, 0)$ its derivative is invertible: the last three columns have determinant $|\partial_u C_s \times v_s|^2 > 0$. The inverse function theorem, continuity and compactness give a common small closed normal disc on which every derivative remains invertible. For the interval endpoints use a smooth local extension in s ; the same estimates restrict back to $[0, 1]$. There is a still smaller common radius on which every fixed- s normal chart is injective. Otherwise choose collisions at normal radii tending to zero. Compactness gives subsequences with $s_j \rightarrow s_*$, $u_j \rightarrow u_*$, and $u'_j \rightarrow u'_*$. Equality of their limits gives $C_{s_*}(u_*) = C_{s_*}(u'_*)$; embeddedness gives $u_* = u'_*$ on the circle. Both colliding points then lie in one inverse chart of the four-dimensional map F near $(s_*, u_*, 0, 0)$, a contradiction. This argument includes both nearby and remotely separated parameter pairs.

Thus for a common $\epsilon_0 > 0$, the normal charts embed the closed disc of radius $2\epsilon_0$ at every s . Restriction to the circles with normal coordinates $(0, 0)$ and $(\epsilon, 0)$, $0 < \epsilon \leq \epsilon_0$, gives disjoint oriented embedded pushoff pairs varying smoothly in s . Interpolating between two positive radii stays in the same disc. The already proved paired invariance therefore gives both family and radius

independence. This proves the general framing assertion, including homotopies of normal fields without any ambient extension theorem.

In particular, for a Legendrian knot choose a smooth contact field w independent of its tangent. Such a field exists explicitly: if $L' = A(\partial_x + y\partial_z) + B\partial_y$, take $w = -B(\partial_x + y\partial_z) + A\partial_y$; $A^2 + B^2 > 0$. The homotopy $v_\theta = \cos \theta w + \sin \theta \partial_z$, $0 \leq \theta \leq \pi/2$, never becomes tangent to L . For $\theta > 0$ its contact-form evaluation is $\sin \theta > 0$; at zero independence follows from the choice of w . The general framing assertion therefore justifies the replacement by the z -direction pushoff in the source's front calculation of tb , without an assumption about homotopic framings.

For $C_s = T_s$ and $v_s = \partial_y$, independence of the two columns follows by applying $dz - y dx$ to a linear relation: it is positive on $\partial_u T_s$ and zero on ∂_y . Write $T_s = (x_s, y_s, z_s)$. Let $a_0 > 0$ be the minimum of $z'_s - y_s x'_s$ and let M bound $|x'_s|$ on the same compact cylinder. Decrease ϵ_0 until $\epsilon_0 M < a_0/2$ (no restriction from this inequality if $M = 0$). Then

$$(dz - y dx)_{T_s + \epsilon \partial_y}(\partial_u T_s) = (z'_s - y_s x'_s) - \epsilon x'_s > a_0/2 > 0. \quad (233)$$

The uniform normal-chart construction has already proved that these are disjoint embedded pairs with the required orientations. Fixing any $0 < \epsilon \leq \epsilon_0$ now gives the smooth family to which (227) applies. Interpolating between two such positive radii stays inside the same normal chart and proves radius independence as well. Finally ∂_y is a globally nonzero section of $\ker(dz - y dx)$, exactly the source's standard-space framing. Its local linking computations use the mixed-crossing formula proved above. This fixes the convention in (43) and proves the required smooth-family self-linking invariance without importing it as a separate theorem. $\square$

Proof of Lemma 3.51

Proof of 3.51. Disjoint clean supports and positive charts. Restrict one exact-germ interval so that u is a coordinate throughout. The affine page chart

$$X = \frac{x - x_0}{A}, \quad Z = \frac{3(z - z_0 - y_0(x - x_0))}{2A} \quad (234)$$

sends the front to (u^2, u^3) . Its linear determinant is $3/(2A^2) > 0$ for either sign of A , so it preserves page orientation. The cubic coordinate is injective, hence the full local arc is injective even at its cusp.

Choose a smaller closed parameter interval about the cusp inside the chart. The cusp image is absent from the compact image of the complement of that interval's interior: local injectivity and the no-other-branch hypothesis exclude every other preimage, including on different components. Its distance from that complement is positive. Thus a sufficiently small open page neighbourhood V meets only the exact-chart arc. The finitely many cusp images are distinct; shrink their neighbourhoods to make them pairwise disjoint and disjoint from every double point. Choose $b > 0$ so small that the closed subarc $|u| \leq b$ lies inside V and the exact chart.

Let ρ be a smooth real cutoff equal to one near zero, with support strictly inside $(-b, b)$. The image of its compact support has positive distance d from the closed complement of V . Put $M = \max |u\rho(u)|$ on that support; $M > 0$ because $\rho = 1$ near zero. Choose $\epsilon > 0$ with

$$|A|\epsilon M \sqrt{1 + y_0^2} < d/2. \quad (235)$$

Make these choices at each of the finitely many cusps. No continuous selection over a further family or generic projection theorem is used.

The complete rounding family. In the normalized chart set

$$X_\lambda(u) = u^2 + \lambda\epsilon u\rho(u), \quad Z_\lambda(u) = u^3, \quad y_\lambda(u) = y_0 + u. \quad (236)$$

Inverting (234) shows that the physical changes are

$$\Delta x = A\lambda\epsilon u\rho(u), \quad \Delta y = 0, \quad \Delta z = Ay_0\lambda\epsilon u\rho(u). \quad (237)$$

The cutoff vanishes on collars of the chart ends, so every derivative of the change vanishes there. Pasting to the unchanged L therefore gives a jointly smooth family on all parameter circles and the closed λ interval. The formula even extends smoothly as a map to a slightly larger time interval; embeddedness outside $[0, 1]$ is not claimed.

For $u \neq 0$ one has $dZ_\lambda/du = 3u^2 > 0$. At $u = 0$, $dX_\lambda/du = \lambda\epsilon$, since $\rho(0) = 1$. Thus the two projected derivatives are never simultaneously zero for $\lambda > 0$, regardless of the size or sign of a cutoff derivative. Composing with the smooth coordinate u preserves this when $du/d\theta < 0$ as well. At $\lambda = 0$ the only local projected singularity is the original cusp. Spatial immersion holds for every λ because $dy_\lambda/du = 1$ in each modified chart and the rest of L is unchanged.

Throughout the whole exact chart, including cutoff collars and unchanged tails, the coordinate Z_λ is exactly u^3 . For $u < v$,

$$v^3 - u^3 = (v - u)(u^2 + uv + v^2) > 0. \quad (238)$$

Indeed the second factor is $(u + v/2)^2 + 3v^2/4$ and can vanish only when $u = v = 0$, excluded here. Hence distinct parameters anywhere in that chart have different projected points at every λ . This proves injectivity, not merely a tangent test at the centre or at an endpoint.

Equations (235) and (237) keep every moved point inside V , since its projected displacement norm is at most $|A|\epsilon M\sqrt{1 + y_0^2} < d/2$. No parameter outside the exact chart projects into V . Therefore a moved point cannot meet a remote parameter. The different cusp neighbourhoods are disjoint, so different modifications cannot meet each other. Together with strict local injectivity these facts exclude every new projected crossing for every $\lambda \in [0, 1]$, including pairs on different components. All old double points are outside the chosen supports. Their exact branches, tangent determinants and y heights remain unchanged, so transversality, crossing signs and O/U choices persist. There are no other coincidences by hypothesis. For positive λ the result is therefore precisely a finite regular generic diagram.

Any spatial coincidence would project to a new coincidence or to an old double point. The former were excluded; at the latter the two unchanged y heights are distinct. Spatial injectivity follows for the whole family. Combined with the immersion already proved, compactness of the finite parameter union and the Hausdorff target give an embedding: the continuous bijection to its image has continuous inverse. Parameters and component labels have never changed.

Each modified portion is one regular embedded oriented arc in a clean cusp neighbourhood, agreeing with the old arc in its endpoint collars and containing no crossing. It is exactly a clean cusp smoothing. There is no assertion that it stays in the contact plane. For example, where $\rho = 1$, differentiating the physical formulas gives

$$x'_\lambda = 2Au + A\lambda\epsilon, \quad y_\lambda = y_0 + u, \quad (239)$$

$$z'_\lambda = 2Ay_0u + 2Au^2 + Ay_0\lambda\epsilon. \quad (240)$$

Subtracting $y_\lambda x'_\lambda$ from the second line yields

$$z'_\lambda - y_\lambda x'_\lambda = -A\lambda\epsilon u. \quad (241)$$

This can have either sign and vanishes at the centre. It cannot justify Legendrian or positive-transverse invariance, nor transport self-linking. $\square$

Proof of Lemma 3.54

Proof of 3.54. Use Lemma 3.51 on the exact germs (45). It supplies a jointly smooth family L_λ of spatial embeddings, $L_0 = L$, fixed outside finitely many clean cusp neighbourhoods. Its endpoint L_1 has an ordinary finite regular generic diagram D_ϵ , with all old crossings and every component retained. The construction keeps $y = y_0 + u$ and, in the positive affine page chart, uses $X_\lambda = u^2 + \lambda \epsilon u \rho(u)$, $Z_\lambda = u^3$. The determinant of this chart is $3/(2A^2) > 0$, so it has not changed the crossing convention. Lemma 3.51 proves injectivity on the whole chart, including cutoff collars, and excludes remote intersections for every parameter value; an endpoint tangent check alone would not do so.

The two actual diagrams $S(F)$ and D_ϵ have identical named decorated records. Each clean cusp replacement inserts no crossing visit and keeps the same oriented attachments. All crossing pairings, cyclic successors, signs, O/U bits and crossing-free component circles agree. Corollary 3.52 (through Lemma 3.20) therefore gives

$$P_{S(F)} = P_{D_\epsilon}. \quad (242)$$

This is scalar record equality, not a claimed classification of arbitrary relative arc embeddings.

Let $\eta : [0, 1] \rightarrow [0, 1]$ be smooth and increasing, with its endpoint values 0, 1 and every positive-order derivative zero at both endpoints. For example normalize the integral of the bump $\exp(-1/(s(1-s)))$ on $(0, 1)$, extended by zero. Put

$$H_s = \begin{cases} L_{1-\eta(2s)}, & 0 \leq s \leq \frac{1}{2}, \\ G_{\eta(2s-1)}, & \frac{1}{2} \leq s \leq 1. \end{cases} \quad (243)$$

Both branches equal L at the join. Every positive-order time derivative vanishes there; mixed parameter derivatives have the same property, while pure circle derivatives agree with those of L . Thus the chain rule gives joint smoothness, including at the join. Every slice is one of the supplied or proved embeddings. The literal endpoints are L_1, T , with their given parameter orientations, not unidentified nearby knots.

Literature input 1 retains the original full global HOMFLY interface. The supplied smooth family H_s is a smooth isotopy of the actual oriented embedded circles. To make the ambient premise explicit, extend its velocity to a compactly supported ambient field as follows. At each (s, u) choose a local normal chart for the embedded circle, smoothly also in s , by the inverse function theorem; finitely many charts cover the compact parameter cylinder. Restrict every chart to one common normal radius on which a spatial point at time s has at most one normal-bundle preimage. Such a radius exists by the argument of the uniform normal chart in the proof of Lemma 3.49, applied to the normal bundle of the jointly smooth family H_s of embeddings of the finite union C : two distinct normal-bundle points at a common time with a common image and normal lengths tending to zero would accumulate, by compactness of the cylinder, at one point of the cylinder — the slice H_{s_*} being an embedding of C — near which a single normal chart is injective, a contradiction. Below that radius a point of the tracked curve lies in a chart only as its own foot point. In each chart extend $\partial_s H_s(u)$ constantly in its two normal coordinates. Choose smooth spatial-time cutoffs supported in those charts whose sum is positive on the tracked curve, and divide by their sum near that curve. Their weighted sum equals $\partial_s H_s(u)$ on the curve, because every local extension has that same value there. Multiply by a further cutoff equal to one near the whole compact tracked cylinder and zero outside a compact neighbourhood. This gives a jointly smooth bounded spatial vector field, with bounded derivatives on the compact time interval. At the time endpoints extend the given flattened family smoothly to a small interval; the same construction restricts to $[0, 1]$.

Smooth ODE existence and continuation give its flow for the full compact time interval. Uniqueness and reverse-time evolution give an ambient diffeomorphism at each time. The chain rule and

uniqueness imply that the flow carries the parametrized H_0 to H_s . It fixes the complement of a compact set and preserves ambient orientation, since its derivative determinant starts positive and never vanishes. Thus its endpoint gives an actual ambient isotopy between the two oriented links. The retained global source premise yields $H_{D_\epsilon} = H_{D_T}$. Theorem 3.18 identifies these original values with P on both actual diagrams, proving $P_{D_\epsilon} = P_{D_T}$. Combining this equality with (242) proves (46). This step explicitly consumes the global L1 premise, Literature input 1; it is not a source-free replacement or an application of the local LM theorem beyond its stated domain.

Only the endpoints were assumed to have ordinary generic diagrams. The cusped diagram at the join causes no problem because H_s is a spatial embedding throughout. The rounding is ordinary, not Legendrian or positive transverse: where $\rho = 1$ its contact-form evaluation in the u coordinate is $-A\lambda\epsilon u$. No self-linking number is transported along this family. $\square$

Proof of Theorem 3.57

Proof of 3.57. Writing $\alpha = dz - y dx$ gives $d\alpha = dx \wedge dy$ and

$$\alpha \wedge d\alpha = dz \wedge dx \wedge dy = dx \wedge dy \wedge dz. \quad (244)$$

The source observer is on the negative- y side looking in the positive y direction, so smaller y is over. For a Legendrian front this is also the smaller-slope rule $y = dz/dx$. The normal toward the observer is $\nu = -\partial_y$ and $\det_{xyz}(\partial_x, \partial_z, \nu) = 1$. For projected over and under tangents (a, c) and (b, d) ,

$$\det \begin{pmatrix} a & b & 0 \\ 0 & 0 & -1 \\ c & d & 0 \end{pmatrix} = ad - bc = \det_{xz}((a, c), (b, d)). \quad (245)$$

This identifies each source crossing sign with the sign convention used here. Summing the source transverse-front writhe formula, whose sl is the number of (43) by Remark 3.50, proves (49). Absence of downward tangencies alone is not a converse transverse-lift theorem: the crossing-height inequalities will be checked separately in the carrier-floor proof.

With the same vertical coordinate, Ng's downward cusp count is D . Substitute the first two formulas of (47) into the third:

$$\begin{aligned} tb - r &= w - \frac{D + U}{2} - \frac{D - U}{2} \\ &= w - \frac{2D}{2} = w - D = sl_{Ng}(F). \end{aligned} \quad (246)$$

This identifies Ng's pushoff as Etnyre's positive pushoff, fixing the possible exchange of names in other conventions.

Fix T with the stated finite ordinary generic diagram D_T . Lemma 3.45 gives the original Legendrian helix L and a chosen positive pushoff transversely isotopic to T . Let Φ_s be the ambient contact isotopy of Theorem 3.48; write $L_T = \Phi_1 \circ L$ and let F_T be its finite generic front, with the exact cusp germs (45). The isotopy carries the chosen pushoff through positive transverse embeddings, so $T_+(L_T)$ is transversely isotopic to the same T . Lemma 3.49 preserves self-linking along this smooth positive transverse family. The preceding cusp calculation identifies its value with $sl_{Ng}(F_T)$. The front F_T lies on the domain of Definition 3.33: Theorem 3.48 gives finitely many semicubical cusps with the germ (45), whose $x''(0) = 2A \neq 0$, transverse double points, no triple point and no cusp on another branch, and Legendrianity $z' = yx'$ makes the front tangent vanish wherever $x' = 0$, so that every such point is one of these cusps and no regular arc has a vertical tangency. Lemma 3.56, applied to F_T , bounds it by $-\max \deg_a P_{S(F_T)} - 1$.

Let Ψ_t be the ordinary ambient flow of (215). The maps

$$\Theta_t = \Psi_t \circ \Phi_{1-t} \circ \Phi_1^{-1}, \quad 0 \leq t \leq 1, \quad (247)$$

give the supplied family $G_t = \Theta_t \circ L_T$. Indeed $\Phi_1^{-1} \circ L_T = L$, so $G_t = \Psi_t \circ \Phi_{1-t} \circ L$; at zero it equals L_T and at one it equals the parametrized T . The two smooth compactly supported ambient flows make G jointly smooth, and each slice is an embedding with nonzero parameter derivative. They transport the given circle orientation. The generic front theorem supplies all other cusp and crossing hypotheses of Lemma 3.54; the specified ordinary diagram D_T is the endpoint hypothesis already included in the present statement. Consequently that lemma gives $P_{S(F_T)} = P_{D_T} = P_T$. Substitution proves (50).

The polynomial equality explicitly uses the original global HOMFLY source premise in Literature input 1. The positive transverse family alone transports self-linking; the distinct ordinary flow and cusp rounding are used only for the polynomial. No maximum over representatives, mirror, variable change or new owner assumption enters this composition. $\square$

Proof of Lemma 3.59

Proof of 3.59. For (i), uniform continuity gives $0 = s_0 < \dots < s_m = 1$ such that $T(s) \cdot T(s_{j-1}) > 0$ on every $[s_{j-1}, s_j]$. After choosing $\theta(0)$, continue it there by

$$\theta(s) = \theta(s_{j-1}) + \text{atan2}(\det(T(s_{j-1}), T(s)), T(s_{j-1}) \cdot T(s)).$$

The second argument is positive, so the added angle lies in $(-\pi/2, \pi/2)$; the formulas match at the subdivision points and construct a continuous lift without a covering-space theorem. Since $T(1) = T(0)$, its endpoint increment belongs to $2\pi\mathbb{Z}$. Two lifts differ continuously by a value in $2\pi\mathbb{Z}$, hence by one constant.

If the increment is $2\pi k$, extend the lift by $\theta(s+n) = \theta(s) + 2\pi nk$. Every interval of length one then has the same increment, proving seam independence. An orientation-preserving reparametrisation of the circle has an increasing representative $\phi: \mathbb{R} \rightarrow \mathbb{R}$ with $\phi(s+1) = \phi(s) + 1$; the lift $\theta \circ \phi$ has the same increment.

For a homotopy $T(s, t)$ fix t_0 . Uniform continuity gives a neighbourhood of t_0 on which $T(s, t) \cdot T(s, t_0) > 0$ for every s . There is therefore a unique continuous relative angle

$$\alpha(s, t) = \text{atan2}(\det(T(s, t_0), T(s, t)), T(s, t_0) \cdot T(s, t)) \in (-\pi/2, \pi/2).$$

If θ_0 lifts $T(\cdot, t_0)$, then $\theta_0 + \alpha(\cdot, t)$ lifts $T(\cdot, t)$. Periodicity gives $\alpha(1, t) = \alpha(0, t)$, so the endpoint increment is locally, hence globally, constant in t . A regular homotopy supplies such a homotopy of normalised tangents. For the oppositely oriented curve, $-T_\gamma(1-s)$ has lift $\theta(1-s) + \pi$, whose increment is the negative of the original one.

For (ii), the identity and all its polygonal consequences are exactly Lemma 1.20: polygonal rotation was defined there by the sum of principal angles, whose integrality and continuity were proved there.

For (iii), concatenate constant lifts on the straight pieces with the prescribed corner lifts. Their total increment is $\sum_i \vartheta_i$, so (ii) proves the rounding assertion.

For the replacement assertion, use the same lift on the unchanged complementary arc. Shifting that complementary lift at a join changes neither its increment nor the computation; subtraction cancels it and leaves $\Delta_c - \Delta_b$.

Finally follow the tangent path of c and then the tangent path of b backwards. This is a closed direction loop of increment $\Delta_c - \Delta_b$. The maps

$$\nu_t(z) = \frac{A_t z}{|A_t z|}$$

give a homotopy of that direction loop. Clause (i) keeps its increment constant, proving the last assertion without invoking degree. $\square$

Proof of Lemma 3.60

Proof of 3.60. The junction, constructed. Since $0 < |\vartheta_i| < \pi$, the unit directions $u = \delta_{i-1}/|\delta_{i-1}|$ and $v = \delta_i/|\delta_i|$ are neither parallel nor antiparallel. Write $A_0 = q_i - \varepsilon u$ and $A_1 = q_i + \varepsilon v$ for the two points at distance ε from q_i along the incident edges, so that the displacement to be realized is

$$A_1 - A_0 = \varepsilon(u + v),$$

a nonzero vector along the bisector of the convex angle at q_i . Fix once and for all the *transition profile*

$$\varphi(t) = \frac{f(t)}{f(t) + f(1-t)}, \quad f(t) = \begin{cases} e^{-1/t}, & t > 0, \\ 0, & t \leq 0, \end{cases}$$

which is exhibited rather than posited, and which has the four properties every use below reads. It is C^∞ on $[0, 1]$: f is C^∞ on $\mathbb{R}$ with all derivatives vanishing at 0, and the denominator $f(t) + f(1-t)$ is positive on $[0, 1]$, being $f(1) > 0$ at $t = 0$ and $f(1) > 0$ at $t = 1$ and a sum of two nonnegative terms not both zero in between. It has $\varphi(0) = 0$ and $\varphi(1) = 1$. It satisfies $\varphi(1-t) = 1 - \varphi(t)$, by inspection of the formula. And it is *strictly increasing on* $(0, 1)$:

$$\varphi'(t) = \frac{f'(t)f(1-t) + f(t)f'(1-t)}{(f(t) + f(1-t))^2} > 0 \quad (0 < t < 1),$$

both terms of the numerator being positive there, since $f > 0$ and $f'(t) = t^{-2}e^{-1/t} > 0$ on $(0, 1)$. Every derivative of φ vanishes at 0 and at 1, because every derivative of f vanishes at 0 and the denominator is smooth and nonvanishing.

Strict increase is not a convenience. An earlier revision asked only that φ be nondecreasing, and a nondecreasing profile may be constant on a subinterval; the tangent direction would then be constant on a whole sub-arc of the junction, so a direction attained there would be attained at a continuum of parameters rather than at isolated ones, and every count of the points carrying a prescribed tangent direction — the counts this section takes later — would be infinite. With φ strictly increasing on $(0, 1)$ the tangent angle $\theta_u + \vartheta_i \varphi(t/\ell)$ is strictly monotone across the junction, and each direction in the swept arc is attained exactly once. Let θ_u be the angle of u and, for a length $\ell > 0$, let

$$\gamma_\ell(s) = A_0 + \int_0^s (\cos(\theta_u + \vartheta_i \varphi(t/\ell)), \sin(\theta_u + \vartheta_i \varphi(t/\ell))) dt, \quad s \in [0, \ell].$$

Then γ_ℓ is a unit-speed C^∞ arc. Its tangent equals u at $s = 0$ and v at $s = \ell$, and agrees with those constants *to infinite order* at those two parameters, every derivative of φ vanishing there; it does not equal them on a neighbourhood, and cannot, since the angle is strictly monotone on $(0, \ell)$. Infinite-order agreement at the endpoint is what the junction needs: it is exactly the condition under which replacing the corner by the arc leaves a C^∞ regular curve, since all one-sided derivatives match

those of the straight edge. The tangent angle moves strictly monotonically through exactly ϑ_i and no more, and its derivative $\vartheta_i \varphi'(s/\ell)/\ell$ is nonzero on $(0, \ell)$, which is (b).

Its endpoint is $A_0 + \ell m(\varphi)$ where $m(\varphi) = \int_0^1 (\cos(\theta_u + \vartheta_i \varphi), \sin(\theta_u + \vartheta_i \varphi)) dt$. The symmetry $\varphi(1-t) = 1 - \varphi(t)$ makes the angles $\theta_u + \vartheta_i \varphi(t)$ and $\theta_u + \vartheta_i \varphi(1-t)$ reflections of each other in the bisector, so $m(\varphi)$ points along $u + v$; and $|m(\varphi)| > 0$ because all the directions integrated lie in a closed angular interval of width $|\vartheta_i| < \pi$. Taking $\ell = \varepsilon |u + v|/|m(\varphi)|$ makes the endpoint exactly A_1 .

The arc lies in the triangle $A_0 q_i A_1$, in coordinates. Put $\alpha = |\vartheta_i|/2 \in (0, \pi/2)$ and $s = \operatorname{sgn} \vartheta_i$, let w be the unit vector along $u + v$ and z the unit vector with $\det(w, z) = 1$, and take q_i as origin, writing a point as (x_w, x_z) in that frame. Then

$$u = \cos \alpha w - s \sin \alpha z, \quad v = \cos \alpha w + s \sin \alpha z,$$

because u and v are unit vectors whose sum is along w and whose angle is $|\vartheta_i| = 2\alpha$, turning from u to v by ϑ_i . Hence $A_0 = -\varepsilon u$ has coordinates $(-\varepsilon \cos \alpha, s\varepsilon \sin \alpha)$ and $A_1 = \varepsilon v$ has $(\varepsilon \cos \alpha, s\varepsilon \sin \alpha)$, and the closed triangle with vertices A_0, q_i, A_1 is exactly

$$T_i = \{ (x_w, x_z) : |x_w| \leq \cot \alpha \cdot (s x_z), \quad 0 \leq s x_z \leq \varepsilon \sin \alpha \}, \quad (248)$$

the first inequality being the pair of cone conditions at the apex q_i and the second the chord through A_0 and A_1 : a point $\lambda(-u) + \mu v$ with $\lambda, \mu \geq 0$ has $x_w = (\mu - \lambda) \cos \alpha$ and $s x_z = (\lambda + \mu) \sin \alpha$, so $|x_w| \leq (\lambda + \mu) \cos \alpha = \cot \alpha \cdot (s x_z)$ with equality exactly on the two edges, and $\lambda + \mu \leq \varepsilon$ is the chord.

The junction arc satisfies all three inequalities of (248). Write its unit tangent as $T(\sigma) = \cos \beta(\sigma) w + s \sin \beta(\sigma) z$ with $\beta(\sigma) = -\alpha + 2\alpha \varphi(\sigma/\ell)$, which is the tangent angle constructed above read in this frame: at $\sigma = 0$ it is u ($\beta = -\alpha$), at $\sigma = \ell$ it is v ($\beta = \alpha$), and β is strictly increasing with values in $[-\alpha, \alpha]$. *The two cone inequalities.* Let n_1 be the unit normal to u with $\langle n_1, v \rangle > 0$. The line $\mathbb{R}u$ contains both q_i and A_0 , and T_i lies in $\{\langle \cdot, n_1 \rangle \geq 0\}$. Along the arc, $\frac{d}{d\sigma} \langle \gamma, n_1 \rangle = \langle T, n_1 \rangle \geq 0$, because $T = au + bv$ with $a, b \geq 0$ — a direction at angle $\beta \in [-\alpha, \alpha]$ from w lies in the closed cone of u and v — so $\langle T, n_1 \rangle = b \langle v, n_1 \rangle \geq 0$; and $\langle \gamma(0), n_1 \rangle = \langle -\varepsilon u, n_1 \rangle = 0$. Hence $\langle \gamma, n_1 \rangle \geq 0$ throughout. Symmetrically, with n_2 the unit normal to v with $\langle n_2, A_0 \rangle > 0$, one has $\langle T, n_2 \rangle = a \langle u, n_2 \rangle \leq 0$, so $\langle \gamma, n_2 \rangle$ is nonincreasing and equals $\langle \varepsilon v, n_2 \rangle = 0$ at $\sigma = \ell$; hence it is ≥ 0 throughout. Those two are the first inequality of (248). *The chord inequality.* $s x_z(\sigma) = \varepsilon \sin \alpha + \int_0^\sigma \sin \beta(t) dt$, and $\int_0^\sigma \sin \beta \leq 0$ for every $\sigma \in [0, \ell]$: β is negative on $(0, \ell/2)$ and positive on $(\ell/2, \ell)$, and the symmetry $\varphi(1-t) = 1 - \varphi(t)$ makes $\beta(\ell-t) = -\beta(t)$, so the integral decreases to its minimum at $\ell/2$ and increases back to 0 at ℓ . Hence $s x_z \leq \varepsilon \sin \alpha$, which is the second. The lower bound $s x_z \geq 0$ is implied by the two cone inequalities, the apex being the origin.

Its bisector coordinate has derivative $dx_w/d\sigma = \cos \beta(\sigma) \geq \cos \alpha > 0$. Thus that coordinate is strictly increasing and the junction arc is injective; no self-intersection can be hidden inside its rounding disc.

So the arc lies in T_i , and T_i lies in the closed disc D_i of radius ε about q_i : the disc is convex and contains the three vertices, $|A_0 - q_i| = |A_1 - q_i| = \varepsilon$. An earlier revision argued this by saying the arc “cannot leave the region cut off by the two tangent lines”, which names the conclusion rather than proving it and leaves the chord side unaddressed.

For (a) and (c), let $\mathcal{E}_i$ be the union of the closed edges not incident to q_i , and define the four clearances separately:

$$\begin{aligned} \eta_v &= \min_i \min_{k \neq i} |q_k - q_i|, & \eta_e &= \min_i \operatorname{dist}(q_i, \mathcal{E}_i), \\ \eta_\ell &= \min_i |\delta_i|, & \eta_X &= \min_{i, x \text{ a double point}} |x - q_i|. \end{aligned}$$

Set

$$\varepsilon_0(L) = \frac{1}{3} \min\{\eta_v, \eta_e, \eta_\ell, \eta_X\},$$

and require $\varepsilon < \varepsilon_0(L)$; the index i is bound by the outer minimum in every term. Each minimum over an empty index set is $+\infty$ — which is how the last one is read when L has no double point. Every listed quantity is positive: the corners are finitely many and distinct; a corner is at positive distance from a non-incident closed edge because the edges are compact and a corner on one would contradict the hypothesis that no corner lies on a non-incident edge; the edges have positive length; and a double point lies in the relative interiors of two edges, so it is distinct from every corner. With this ε the discs D_i are pairwise disjoint, each D_i meets no edge not incident to q_i , each D_i contains no double point, and each D_i meets the two incident edges in the two sub-segments of length ε at q_i , and each contains the whole modification made at its corner, the replacement arc lying within distance ε of q_i . That is clause (e). Hence the rounding changes the curve only inside those discs, where it meets nothing else, so no double point is created or destroyed and the strands through each surviving double point are unchanged as arcs, with their orientations; the over/under assignment is inherited, so the crossing signs and the writhe are unchanged. That is (a) and (c). For (d), the total turning of L_ε is the sum of the turnings of its arcs, which is $\sum_i \vartheta_i$ because the tangent is constant along the straight parts; and the rotation number of a C^1 regular closed curve is $\frac{1}{2\pi}$ times its total turning (Lemma 3.59(iii)), which is also the value assigned to L by Lemma 1.20. $\square$

Proof of Lemma 3.61




Figure 14: The local transverse-lift replacement: the forbidden downward turn is replaced by a curl with a positive crossing, as in Lemma 3.61.




Figure 15: The curl and its parallel push-off, followed by the straightened ribbon. The displayed positive twist illustrates the local framing contribution in the curl construction.

Proof of 3.61. Work first in the rational model on $-2 \leq t \leq 2$,

$$b(t) = \left(\frac{3(t^2+7)}{11}, -3t \right), \quad c(t) = (t^2 - 1, t - t^3).$$

The endpoints and endpoint tangent rays agree up to positive scale: $b(\pm 2) = c(\pm 2) = (3, \mp 6)$ and $b'(\pm 2) = \frac{3}{11}c'(\pm 2)$. The old arc turns strictly positively and has exactly one tangent pointing

straight down, at $t = 0$:

$$\det(b'(t), b''(t)) = \frac{18}{11} > 0, \quad b'(0) = (0, -3).$$

The replacement turns strictly negatively and has exactly one tangent pointing straight up, at $t = 0$:

$$\det(c'(t), c''(t)) = -2(1 + 3t^2) < 0, \quad c'(0) = (0, 1).$$

In each arc the first derivative has vanishing first component only at $t = 0$, so neither arc has any other vertical tangent; this is (ii) in the model. The replacement runs between the same endpoint rays along the complementary, clockwise tangent path, so its compatible lift increment is the old one minus 2π ; Lemma 3.59(iii) gives (iv) for rot. It has exactly one double point,

$$c(-1) = c(1) = (0, 0), \quad c'(-1) = (-2, -2), \quad c'(1) = (2, -2),$$

and putting the later branch over the earlier one makes it negative, since a crossing sign is the sign of the determinant of the overpassing tangent followed by the underpassing tangent and $\det(c'(1), c'(-1)) = -8 < 0$; that is (iii), and with it (iv) for w . Being a single Reidemeister-I monogon, the replacement does not change the knot type, which is (i).

To insert the model at p we first record the projective fact it needs, in the form the insertion argument requires.

Ordered-ray lemma. Let (r_1, r_2, r_3) and (s_1, s_2, s_3) be triples of rays from the origin such that r_2 lies in the open convex cone spanned by r_1, r_3 and s_2 lies in the open convex cone spanned by s_1, s_3 , with r_1, r_3 and s_1, s_3 independent, and suppose in addition that for positive representatives

$$\operatorname{sgn} \det(\rho_1, \rho_3) = \operatorname{sgn} \det(\sigma_1, \sigma_3), \quad (249)$$

a condition independent of the representatives chosen. Then there is a linear map A with $\det A > 0$ carrying each r_k into s_k . *Proof:* if both determinants in (249) are negative, exchange the names 1 and 3 in both triples simultaneously; this preserves the correspondence $r_k \mapsto s_k$ and makes both determinants positive. The linear map A_0 with $A_0\rho_1 = \sigma_1$, $A_0\rho_3 = \sigma_3$ then has positive determinant and carries r_1, r_3 to s_1, s_3 . It carries the open cone of r_1, r_3 onto that of s_1, s_3 , so A_0r_2 is a ray σ in the latter cone. Writing $s_2 = \mathbb{R}_{>0}(x\sigma_1 + y\sigma_3)$ and $\sigma = \mathbb{R}_{>0}(x'\sigma_1 + y'\sigma_3)$ with $x, y, x', y' > 0$, the map $D = \operatorname{diag}(x/x', y/y')$ in the basis (σ_1, σ_3) has positive determinant, fixes s_1 and s_3 , and carries σ to s_2 . Take $A = DA_0$. $\square$

An earlier revision of this lemma omitted (249) and said the positive determinants could be arranged “after swapping the names once if necessary”. A swap performed in one triple only destroys the correspondence it is supposed to produce, and the peer refuted that step; the hypothesis is now stated and the swap is simultaneous. In the application below the condition is not inferred from the two cones: both determinants are computed and are positive.

A positive-turn chart, and the cuts. Parametrise the embedded arc by arclength as $\gamma(s)$ with $\gamma(0) = p$ and unit tangent $\gamma'(0) = u$. Because the tangent turns strictly positively, after shrinking the chart there is a continuous strictly increasing lift θ with

$$\gamma'(s) = \cos \theta(s) u + \sin \theta(s) Ju, \quad \theta(0) = 0, \quad |\theta(s)| < \frac{\pi}{2},$$

where J is the positive quarter turn. Put $v = -Ju$, so that (v, u) is a positively oriented orthonormal basis, and set

$$\xi(s) = \langle \gamma(s) - p, v \rangle, \quad \eta(s) = \langle \gamma(s) - p, u \rangle,$$

so that $\xi'(s) = -\sin \theta(s)$ and $\eta'(s) = \cos \theta(s)$. Hence ξ increases strictly for $s < 0$ and decreases strictly for $s > 0$, with a strict maximum $\xi(0) = 0$, while η increases throughout. For every level

sufficiently close to $\xi(0)$ from below, strict monotonicity and the intermediate value theorem give unique cuts $s_- < 0 < s_+$ with $\xi(s_-) = \xi(s_+)$, and both tend to 0 as the level tends to $\xi(0)$. Write $p_{\pm} = \gamma(s_{\pm})$. Equal transverse coordinates and monotone η give

$$p_+ - p_- = \ell u, \quad \ell > 0, \quad \ell \rightarrow 0.$$

Two side facts follow at once and are what the earlier revision lacked: the central arc $\gamma((s_-, s_+))$ lies strictly in $\xi > \xi(s_-)$, and both retained tails lie strictly in $\xi < \xi(s_-)$.

The affine fit, quantitatively. In the positively oriented model basis $v_0 = (-1, 0)$, $u_0 = (0, -1)$ one has $q = \frac{4}{11}$ and

$$b'(-2) \in \mathbb{R}_{>0}(qv_0 + u_0), \quad b'(0) \in \mathbb{R}_{>0}u_0, \quad b'(2) \in \mathbb{R}_{>0}(-qv_0 + u_0).$$

Write the two target endpoint tangents in the basis (v, u) as $T_- = \gamma'(s_-) = av + bu$ and $T_+ = \gamma'(s_+) = -cv + du$; then $a = -\sin \theta(s_-) > 0$, $c = \sin \theta(s_+) > 0$ and $b, d = \cos \theta(s_{\mp}) > 0$, and $a, c \rightarrow 0$, $b, d \rightarrow 1$ as the cuts approach p . The two outer determinants are $\det(qv_0 + u_0, -qv_0 + u_0) = 2q > 0$ and $\det(T_-, T_+) = ad + bc > 0$, so (249) holds; it is computed, not inferred. Define

$$H = \frac{2}{\frac{b}{a} + \frac{d}{c}} = \frac{2ac}{bc + ad}, \quad x = \frac{H}{q}, \quad y = \frac{\frac{b}{a} - \frac{d}{c}}{q(\frac{b}{a} + \frac{d}{c})} = \frac{bc - ad}{q(bc + ad)},$$

and let A be the linear map with $Au_0 = u$ and $Av_0 = xv + yu$. Then

$$A(qv_0 + u_0) = \frac{H}{a} T_-, \quad A(-qv_0 + u_0) = \frac{H}{c} T_+, \quad \det A = x = \frac{2ac}{q(bc + ad)} > 0,$$

each by substitution, and

$$1 - (qy)^2 = \frac{4abcd}{(bc + ad)^2} > 0,$$

so $|y| < 1/q$ uniformly, while $x \rightarrow 0$ as the cuts approach p : the map A stays bounded. This is the quantitative statement that the earlier revision replaced by an unquantified “fixed unit-normalized part of A ”.

Since $b(2) - b(-2) = 12u_0$, put $B = \frac{\ell}{12}A$ and $\Phi(z) = p_- + B(z - b(-2))$. Then $\Phi(b(-2)) = p_-$, $\Phi(b(2)) = p_- + \ell Au_0 = p_- + \ell u = p_+$, and the endpoint tangent rays match with positive scale. Because $\ell \rightarrow 0$ while A stays bounded, the diameter of $\Phi(c([-2, 2]))$ tends to 0, so the cuts may be chosen so that the whole inserted arc lies in a preassigned disc Δ about p whose intersection with the diagram is contained in the embedded chart — the disc is fixed first and the cuts afterwards.

Only the model's own double point is created. In the basis (v_0, u_0) the transverse coordinate of $c(t)$ is $1 - t^2$, which equals -3 at both endpoints, so the model's transverse excess over its endpoint chord is $4 - t^2 > 0$ for $-2 < t < 2$. Since $Bv_0 = \frac{\ell}{12}(xv + yu)$ with $x > 0$ and $Bu_0 = \frac{\ell}{12}u$, the physical transverse excess of the inserted arc over the chord $[p_-, p_+]$ is exactly

$$\frac{\ell x}{12}(4 - t^2) > 0.$$

So the inserted interior lies strictly in $\xi > \xi(s_-)$ while both retained tails lie strictly in $\xi < \xi(s_-)$: they meet only at the two joins. Together with the shrinking of the previous paragraph, which keeps the inserted arc away from every nonlocal part of the diagram, this *proves* rather than assumes that the only double point created is the model's own.

The replacement, and why the curve stays C^∞ regular. Replace the sub-arc of F between p_- and p_+ by $\Phi \circ c$, modified near its two endpoints as follows. At each join the two one-sided unit

tangents already agree, $\Phi c'(\pm 2)$ being a positive multiple of $\Phi b'(\pm 2)$ and that of the old tangent, so the concatenation is C^1 ; it need not be C^∞ , and the collar that repairs that is constructed here rather than described. An earlier revision said only that one may “interpolate the tangent angle by the same device”, which does not say what curve results, and interpolating angles moves the endpoint.

Before choosing the two collars, take disjoint small neighbourhoods of p_- and p_+ meeting no part of the retained old curve or fitted model outside their respective local endpoint intervals. These neighbourhoods exist by compactness: each endpoint has a unique preimage in the joined curve, the only model double point is away from the endpoints, and the transverse-excess argument has excluded a model–tail intersection. Delete small endpoint parameter intervals; each remaining compact image has positive distance from that endpoint. Take a ball smaller than those distances. As its length tends to zero, the region between the two local graphs used for a collar shrinks to the endpoint, so both collars can be chosen inside these balls. This also separates them from the monogon and from one another.

The collar, as a graph. Work at the join p_- first; the join at p_+ is constructed after property (3) below, with its own signs printed rather than left to “roles exchanged”. In the chart coordinates (ξ, η) of the positive-turn chart — $\xi = \langle \gamma - p, v \rangle$ with $v = -Ju$, $\eta = \langle \gamma - p, u \rangle$ — the old arc is a graph $\xi = f(\eta)$, since $\eta' = \langle \gamma', u \rangle = \cos \theta > 0$ there. Its slope carries a sign that is load-bearing below, so it is computed: $\xi' = \langle \gamma', v \rangle = -\sin \theta$, whence

$$f'(\eta) = \frac{\xi'}{\eta'} = -\tan \theta.$$

For $\eta < \eta(p)$ the angle θ is strictly negative, θ being strictly increasing with $\theta(0) = 0$, so $f' > 0$ there — not $f' < 0$, which is what the earlier revision’s $f' = \tan \theta$ gave. Let $\eta_- = \eta(p_-) < \eta(p)$. The inserted arc leaves p_- with the same unit tangent, so near its start it is also a graph $\xi = g(\eta)$ with $g(\eta_-) = f(\eta_-)$ and $g'(\eta_-) = f'(\eta_-)$; and because the model turns strictly negatively its tangent angle only decreases from $\theta(s_-) < 0$, so by the same formula

$$g' > f' > 0$$

on that stretch. Two consequences are used below: $g > f$ on $(\eta_-, \eta_- + \lambda]$, the two agreeing at η_- with $g' > f'$; and at every point of the collar the smaller slope is $f'(\eta)$, still positive. The bound is pointwise of necessity: $f' = -\tan \theta$ decreases along the collar, θ increasing, so $0 < f'(\eta) < f'(\eta_-)$ there and an endpoint bound from the left end is unavailable. Fix $\lambda > 0$ small enough that $[\eta_-, \eta_- + \lambda]$ lies both below $\eta(p)$ and inside the stretch on which the inserted arc is a graph, and set, on $[\eta_-, \eta_- + \lambda]$,

$$h(\eta) = (1 - \varphi_\lambda(\eta))f(\eta) + \varphi_\lambda(\eta)g(\eta), \quad \varphi_\lambda(\eta) = \varphi\left(\frac{\eta - \eta_-}{\lambda}\right),$$

with φ the transition profile of Lemma 3.60. The modified curve runs along the old arc up to η_- , along the graph of h across the collar, and along the inserted arc from $\eta_- + \lambda$ on. Three properties, each proved from the formula.

(1) *It is C^∞ and regular, with positive speed.* h is a C^∞ function, and every derivative of φ vanishes at 0 and 1, so h agrees with f to infinite order at η_- and with g to infinite order at $\eta_- + \lambda$: the two joins are C^∞ . Parametrized by η , the collar has velocity $h'(\eta)v + u$ in the frame (v, u) , of norm $\sqrt{1 + h'(\eta)^2} \geq 1$, so its speed is positive and it is regular; being a graph, it is embedded.

(2) *No tangent of the collar equals u , and no smallness condition is needed.* The tangent is parallel to u exactly where $h' = 0$, and

$$h' = (1 - \varphi_\lambda)f' + \varphi_\lambda g' + \frac{\varphi'((\eta - \eta_-)/\lambda)}{\lambda} (g - f).$$

Every term is nonnegative, and the sum of the first two is positive: it is the convex combination $(1 - \varphi_\lambda)f' + \varphi_\lambda g'$ of two positive numbers — each *term* separately can vanish, φ_λ being 0 at the left end and 1 at the right — while the third term is $\varphi' \cdot (g - f)/\lambda \geq 0$, since $\varphi' \geq 0$ and $g - f \geq 0$ there, the two agreeing at η_- with $g' > f'$. Hence, pointwise,

$$h'(\eta) \geq (1 - \varphi_\lambda)f'(\eta) + \varphi_\lambda g'(\eta) \geq f'(\eta) > 0,$$

the middle inequality because $g' > f'$, so h' never vanishes and the collar has no tangent equal to u ; and none equal to $-u$ either, a graph over η having $\eta' > 0$ throughout. An earlier revision, working from the wrong slope sign, had both f' and g' negative, needed a Taylor bound on $g - f$ to control the third term, and chose λ small to make the sum negative. With the frame computed the third term is on the same side as the other two and no choice of λ is required.

(3) *It creates no double point and changes no turning.* Pointwise $f \leq h \leq g$ on the collar — the value is a convex combination and $g \geq f$ there, which is the separation the corrected slope chain supplies — so the collar lies in the region between the two graphs; that region is inside the disc Δ and inside the embedded chart, where the only strands of the diagram are the old arc — whose collar stretch is removed — and the inserted arc. The collar is embedded by (1) and meets those two only at its endpoints, so no double point is created in it, and the model's own double point, at $|t| = 1$, is interior to the inserted arc and untouched. For the turning: the collar's tangent angle is continuous, stays in $(-\pi/2, 0)$ — the slope h' being positive and finite, so the tangent angle θ has $-\tan \theta = h' > 0$ with $\cos \theta > 0$ — and takes at its ends exactly the angles of the two arcs it joins, so the lift of the tangent along the modified curve is the concatenation of the old lifts with an interpolating stretch of the same endpoint values; the total tangent winding is unchanged by its insertion. An earlier revision stopped at C^1 and the document then carried a lemma to recover smoothness. This sentence is what the downstream consumers of this lemma read, the front construction below taking tangent directions of the curve, and an intermediate revision of this proof deleted it while rewriting the tail; it is restored here.

The collar at p_+ . The same three properties are now proved at the other join, with its own signs printed: here the *incoming* arc is the inserted one and the *outgoing* arc is the old one, and both slopes are negative. Let $\eta_+ = \eta(p_+) > \eta(p)$. On a stretch $[\eta_+ - \lambda_+, \eta_+]$ chosen above $\eta(p)$ and inside the stretches on which both arcs are graphs, write $\xi = g(\eta)$ for the inserted arc and $\xi = f(\eta)$ for the old arc extended backward from p_+ ; the two agree at η_+ with common slope $m_+ = -\tan \theta_+ < 0$, the angle $\theta_+ = \theta(s_+)$ lying in $(0, \frac{\pi}{2})$. Just left of the join the inserted arc turns strictly negatively into its endpoint, so its angle exceeds θ_+ there, while the old arc turns positively, so its angle lies below θ_+ ; by slope $= -\tan \theta$, decreasing in θ ,

$$g'(\eta) < m_+ < f'(\eta) < 0 \quad \text{on } [\eta_+ - \lambda_+, \eta_+],$$

the outer inequality $f' < 0$ because $\theta > 0$ above $\eta(p)$. Integrating $g' - f' < 0$ backward from the common value at η_+ gives $g > f$ on $[\eta_+ - \lambda_+, \eta_+]$. Set

$$h = (1 - \chi_{\lambda_+})g + \chi_{\lambda_+}f, \quad \chi_{\lambda_+}(\eta) = \varphi\left(\frac{\eta - (\eta_+ - \lambda_+)}{\lambda_+}\right),$$

with the same transition profile φ : the modified curve runs along the inserted arc up to $\eta_+ - \lambda_+$, along the graph of h across the collar, and along the old arc from η_+ on, the two joins C^∞ by the flat endpoint jets of φ exactly as in (1). For the derivative,

$$h' = (1 - \chi_{\lambda_+})g' + \chi_{\lambda_+}f' + \frac{\varphi'((\eta - \eta_+ + \lambda_+)/\lambda_+)}{\lambda_+} (f - g),$$

and the sign of $f - g$ in the cutoff term is exactly what the negativity of the two slopes alone does not control: here $f - g \leq 0$ by the separation just integrated, and $\varphi' \geq 0$, so the cutoff term is nonpositive, while the convex part is a convex combination of two negative numbers. Hence, pointwise,

$$h'(\eta) \leq (1 - \chi_{\lambda_+})g'(\eta) + \chi_{\lambda_+}f'(\eta) \leq f'(\eta) < 0,$$

the middle inequality because $g' < f'$, and the bound pointwise for the mirrored reason: f' varies along the collar and no endpoint bound is available. So h' never vanishes, and this collar — a graph over η with $\eta' > 0$ throughout — has no tangent equal to u and none equal to $-u$; its tangent angle stays in $(0, \frac{\pi}{2})$, with $-\tan \theta = h' < 0$ and $\cos \theta > 0$, and takes at its ends exactly the angles of the two arcs it joins, so the lift concatenation of (3) applies to it verbatim and the total tangent winding is unchanged. Pointwise $f \leq h \leq g$, the value being a convex combination and $g \geq f$ there, so the collar lies in the region between the two graphs, inside the disc Δ and the embedded chart, and it is embedded and meets the two arcs only at its endpoints: no double point is created, by the argument of (3) unchanged.

The four conclusions, separately. (ii): $c'(t) = (2t, 1 - 3t^2)$ is parallel to u_0 only at $t = 0$, where $c'(0) = -u_0$; since A is invertible with $Au_0 = u$, the inserted arc has no tangent u and exactly one tangent $-u$, and the retained tails have neither, their angles lying in $(-\frac{\pi}{2}, \frac{\pi}{2})$ away from 0.

(iii): if $c(s) = c(t)$ then $s = \pm t$, and the second coordinate forces $2t(1 - t^2) = 0$, so the only pair of distinct parameters is $\{-1, 1\}$: exactly one double point. With the later branch over the earlier one its sign is that of $\det(c'(1), c'(-1)) = -8 < 0$, and $\det B > 0$ preserves it.

(iv) for rot: Let R be the rotation carrying the positive orthonormal basis (v_0, u_0) to (v, u) and put $C = R^{-1}A$. Relative to (v_0, u_0) , C has matrix $\begin{pmatrix} x & 0 \\ y & 1 \end{pmatrix}$ with $x > 0$, so $C_t = \begin{pmatrix} 1-t+tx & 0 \\ ty & 1 \end{pmatrix}$ joins I to C inside $\text{GL}^+(2, \mathbb{R})$. If R_t rotates through t times any fixed angle representing R , first C_t and then $R_t C$ give a path A_t from I to $A = RC$ in that group. Apply $z \mapsto A_t z / |A_t z|$ to the closed direction loop obtained by following the tangent path of c and the tangent path of b backwards. Lemma 3.59(iii) therefore preserves their compatible endpoint lift-increment difference. The tangent of b follows the short positive arc between the endpoint rays; the tangent of c follows the complementary clockwise arc, since $\det(c', c'') = -2(1 + 3t^2) < 0$ and it passes through $-u_0$ at $t = 0$; so their relative winding is -1 . The old central arc also follows the short positive arc, its lifted sweep $\theta(s_+) - \theta(s_-)$ lying in $(0, \pi)$. Hence Lemma 3.59(iii) says that the replacement changes the closed curve's rotation by exactly -1 . For w : the removed arc was embedded and carried no crossing, so the writhe changes by the one new negative crossing, that is by -1 .

The polynomial conclusion, without a knot-category detour. The construction just completed also proves that F' remains in the lemma's input class. It replaces one oriented arc by one C^∞ regular oriented arc, leaves the component count unchanged, and creates exactly one transverse double point in a disc meeting no other strand. Thus its double points remain finite and transverse and no triple point is created.

Let F° be the oriented diagram obtained from F' by deleting the unique monogon in Δ by a Reidemeister-I move. The move is available: the preceding separation argument shows that the monogon is disjoint from all other strands, and the crossing calculation shows it is the unique new double point. It replaces that monogon by one crossing-free regular arc in Δ , so F° is again a one-component immersed-circle diagram with finite transverse double points and no triple points. Reidemeister-I invariance of the local polynomial P (Theorem 3.18, from Literature input 2) gives $P_{F'} = P_{F^\circ}$.

Outside Δ , the diagrams F° and F coincide. Inside Δ , each has one crossing-free oriented arc joining the same two boundary points in the same traversal direction. Match every marked preimage outside that local traversal interval with the identical marked preimage of the other

diagram. The local intervals contain no marked preimage, so this bijection preserves cyclic order, crossing pairs, over/under positions and signs: it is a record isomorphism. Both actual diagrams have one component and finite transverse crossings. Lemma 3.20 therefore gives $P_{F^\circ} = P_F$ directly, and hence $P_{F'} = P_F$. Equality outside Δ proves the remaining clause of (i). $\square$

Proof of Theorem 3.62

Proof of 3.62. Clause (R). Reversing all arrows of a diagram preserves every crossing sign: the sign is $\text{sgn det}(u_{\text{over}}, u_{\text{under}})$, and reversing both strands replaces both vectors by their negatives, leaving the determinant unchanged. Hence the writhe is preserved, and reversal sends a skein triple (D_+, D_-, D_0) of diagrams to the skein triple $(-D_+, -D_-, -D_0)$, the oriented smoothing of the reversed diagram being the reverse of the oriented smoothing. Consider the two maps $D \mapsto P_{-D}$ and $D \mapsto P_D$ on all oriented link diagrams, with values in $R = \mathbb{Z}[a^{\pm 1}, z^{\pm 1}]$. Both take the value 1 on the crossing-free circle, both satisfy the skein of Theorem 3.18, the first by the preceding sentence, and both are invariant under planar isotopy and the three Reidemeister moves, since P is (Theorem 3.18) and reversal carries each such move and isotopy to one of the same kind. Lemma 3.17 identifies the two maps: $P_{-D} = P_D$ for every diagram D . Since P_D is the polynomial H_D of the oriented link presented by D (Theorem 3.18), an oriented knot and its reverse have the same polynomial. For a polygonal underlying curve every principal turn changes sign and Lemma 3.59(ii) negates rotation. For a C^1 regular curve, if θ lifts its normalised tangent, then $\theta(1-s) + \pi$ lifts the reversed tangent with the negative endpoint increment, as in Lemma 3.59(i).

Clause (A). The transition profile is fixed once and for all in the proof of Lemma 3.60. At each corner, ε and the two incident unit directions determine that profile's junction, and the endpoint condition determines its length ℓ ; outside the rounding discs the curve is L , and the diagram data are inherited. Thus the named construction returns one specified pair at every allowed ε . This asserts uniqueness of the fixed construction's record, not uniqueness among all possible smoothings.

Clause (B). Reversing the orientation of L negates every principal turn and the signed rotation and preserves every self-crossing sign, so perform the allowed normalization once. In the uniform case $\text{rot}(L) = R \geq 1$ by Lemma 1.21(i); in the one-dissent case the same holds by Lemma 1.21(ii).

Take $\varepsilon_1 = \varepsilon_0(L)$, the clearance of Lemma 3.60. Choose $u \in S^1$ so that neither u nor $-u$ is an edge direction and, in the one-dissent case, neither lies in the negative turn's closed swept arc A , of length $\alpha := |\vartheta_-| < \pi$. Such u exists: in the one-dissent case $A \cup (-A)$ has length at most $2\alpha < 2\pi$, while in the uniform case there is no such arc; the additional forbidden set of edge directions and their antipodes is finite. This choice depends only on L . Fix an argument ϕ of u .

Fix $\varepsilon \in (0, \varepsilon_1)$, write $(L_\varepsilon, D_\varepsilon) = \text{Round}(L, D, \varepsilon)$, and let T be its unit-tangent map. Parameterize the rounded traversal by $s \in [0, 1]$, with $s = 0$ in the interior of a straight part; because $u, -u$ are not edge directions, its tangent is congruent to neither ϕ nor $\phi + \pi$. Concatenating straight-part arguments with the junction lifts supplied by Lemma 3.60(b) gives a continuous $\theta: [0, 1] \rightarrow \mathbb{R}$, starting at an argument $\theta(0)$ of $T(0)$. Since $T(1) = T(0)$, write $\theta(1) = \theta(0) + 2\pi m$ for an integer m .

Each straight portion contributes zero tangent increment, and each rounding junction contributes its principal angle by Lemma 3.60(b). Thus $2\pi m = \sum_i \vartheta_i = 2\pi \text{rot}(L) = 2\pi R$, by Lemma 1.20. Thus

$$\theta(1) = \theta(0) + 2\pi R. \quad (*)$$

Fix $\beta = \phi$ or $\phi + \pi$. Because $u, -u$ are not edge directions, no level $\beta + 2\pi k$ is met on a straight part or at a junction endpoint; by the choice of u , none is met on the decreasing junction. Hence floor jumps occur only in increasing junctions and never elsewhere. Lemma 3.60(b) makes each

occurrence unique with positive angular derivative. Put

$$x = \frac{\theta(0) - \phi}{2\pi}, \quad y = \frac{\theta(0) - (\phi + \pi)}{2\pi}.$$

For u , (*) and the integer floor identity give, separately,

$$\#T^{-1}(u) = \lfloor x + R \rfloor - \lfloor x \rfloor = R.$$

For $-u$, the same argument at the other levels gives

$$\#T^{-1}(-u) = \lfloor y + R \rfloor - \lfloor y \rfloor = R.$$

There are finitely many junctions, so the preimages are finite; their positive derivatives prove the assertion for arbitrary ε .

Clause (C). Fix $\varepsilon \in (0, \varepsilon_1)$ with ε_1 the clearance of clause (B) and pass to the rounding record $(L_\varepsilon, D_\varepsilon) = \text{Round}(L, D, \varepsilon)$ of clause (A) — one curve and one diagram, not a member of an unnamed family. Lemma 3.60(c),(d) preserves all crossings and their oriented decorated data, the writhe and the rotation number. The rounded diagram and the original polygonal diagram have the same named record, including the single original component, so Lemma 3.20 gives equal P values. By clause (B) there is a direction u met by the tangent of L_ε exactly R times, each time positively, and likewise for $-u$. Rotate coordinates so that u is the downward vertical. The rounded diagram then has exactly R downward vertical tangencies and R upward ones.

Switch every crossing. The result is a diagram $\overline{D}$ of the mirror knot, with the same underlying plane curve — hence the same tangencies — and with writhe $-w$; every crossing of $\overline{D}$ is negative. Apply Lemma 3.61 at each of the R downward vertical tangencies; they are isolated, lie in the interiors of positively turning rounding arcs (clause (B)), and each rounding junction sweeps an angle of magnitude below π , so it contains at most one of these downward tangencies. The rounding discs returned by Lemma 3.60(e) are pairwise disjoint and meet no crossing. Choose each curl disc inside its corresponding rounding disc. Thus all R replacements have disjoint supports and leave one another intact. Call the result T . By Lemma 3.61, $P_T = P_{\overline{D}}$, T has no downward vertical tangency, and it has writhe

$$w(T) = -w - R.$$

Every crossing of T is negative: the old ones by the switch, the R new ones by Lemma 3.61(iii).

We now construct the transverse lift rather than import a front criterion. Write the smooth regular parametrisation of T as $s \mapsto (x(s), z(s))$, put

$$v(s) = \sqrt{x'(s)^2 + z'(s)^2}, \quad y_0(s) = -\frac{x'(s)}{v(s) + z'(s)},$$

and use the convention that, in the xz projection, the branch with smaller y -coordinate is drawn over. Since T has no downward vertical tangency, $v + z' > 0$ everywhere: equality would force $x' = 0$ and $z' < 0$. Thus y_0 is smooth and periodic, and direct algebra, with no division by x' , gives

$$z' - y_0 x' = z' + \frac{x'^2}{v + z'} = z' + \frac{v^2 - z'^2}{v + z'} = v > 0. \quad (*)$$

It remains only to impose the crossing order. At a crossing branch there is no vertical tangency, so put $m = z'/x'$. The strict inequality in (*) permits $y < m$ on a rightward branch and $y > m$ on a leftward branch. For every crossing choose admissible constants $c_0 < c_U$ at its over- and underpassing preimages. Such a choice is automatic if the overpass points right: after choosing

any admissible c_U , take $c_O < \min\{m_O, c_U\}$. If both point left, choose any $c_O > m_O$ and then $c_U > \max\{m_U, c_O\}$. The only remaining case has $x'_O < 0 < x'_U$. Every crossing of T is negative, and here

$$0 > \det(t_O, t_U) = x'_O x'_U (m_U - m_O).$$

Since $x'_O x'_U < 0$, this says $m_O < m_U$; choose $m_O < c_O < c_U < m_U$. These are exactly the two allowed half-lines and the required over/under order.

There are finitely many crossing preimages. Around each one s_j , choose pairwise disjoint cyclic parameter intervals (α_j, β_j) so small that $z' - c_j x' > 0$ throughout. Using the explicit flat transition profile φ displayed in the proof of Lemma 3.60, define a smooth bump on the circle by

$$\psi_j(s) = \begin{cases} \varphi((s - \alpha_j)/(s_j - \alpha_j)), & \alpha_j \leq s \leq s_j, \\ \varphi((\beta_j - s)/(\beta_j - s_j)), & s_j \leq s \leq \beta_j, \\ 0, & \text{otherwise.} \end{cases}$$

All endpoint derivatives are zero, so this is smooth, equals one at s_j and has support in that interval. Set

$$y = y_0 + \sum_j \psi_j(c_j - y_0).$$

The supports are disjoint. On the j -th one,

$$z' - yx' = (1 - \psi_j)(z' - y_0x') + \psi_j(z' - c_jx') > 0,$$

and outside them this is (*). Hence y is smooth and periodic and the lift $s \mapsto (x(s), y(s), z(s))$ is positively transverse to $dz - y dx$.

Finally the lift is regular because its xz projection is immersed. If two lifted points agreed, their projected points would agree; the only distinct parameters with that property are the two preimages of a listed double point, and there $c_O < c_U$. Thus the lift is injective, hence, by compactness of the parameter circle, an embedded transverse knot K_T . Its front and its smaller- y over/under assignments are exactly the diagram T .

By Theorem 3.57, $\text{sl}(K_T) = w(T) = -w - R$. Since $P_T = P_{\overline{D}}$, Theorem 3.57 applied to K_T with its actual finite regular generic diagram T gives

$$\max \deg_a P_{\overline{D}} \leq -\text{sl}(K_T) - 1 = w + R - 1.$$

For every oriented link diagram D let $\overline{D}$ be its mirror, the diagram with every crossing switched, and put $\mathcal{M}(D) = P_{\overline{D}}(a, z)$. Mirroring exchanges D_+ with D_- and fixes D_0 , so the skein of Theorem 3.18 on the mirrored triple, multiplied by -1 , is

$$a^{-1}\mathcal{M}(D_+) - a\mathcal{M}(D_-) = -z\mathcal{M}(D_0), \quad \mathcal{M}(\bigcirc) = 1.$$

Let ι be the involutive automorphism $(a, z) \mapsto (a^{-1}, -z)$ of the Laurent ring. If a map Q on oriented link diagrams satisfies these two conditions and is invariant under planar isotopy and the three Reidemeister moves, then $\iota \circ Q$ satisfies the skein and the unknot value with the same invariance, because ι sends the coefficients $a^{-1}, -a, -z$ to $a, -a^{-1}, z$. Two such maps Q, Q' therefore have $\iota \circ Q = \iota \circ Q'$ by Lemma 3.17, hence $Q = Q'$. Both $\mathcal{M}$ and $D \mapsto \iota(P_D) = P_D(a^{-1}, -z)$ satisfy the two conditions and are invariant — the first because P is and mirroring carries a planar isotopy or a Reidemeister move to one of the same kind, the second because P is — hence $P_{\overline{D}}(a, z) = P_D(a^{-1}, -z)$ for every diagram D . As a Laurent-ring automorphism, the substitution negates every a -exponent and cannot cancel a nonzero extreme coefficient. Hence $\max \deg_a P_{\overline{D}} = -\min \deg_a P_D$,

and the displayed bound becomes $\min \deg_a P_D = -\max \deg_a P_{\overline{D}} \geq 1 - w - R$, which is (51). Finally, the support of f_D is contained in the support of P_D , so the same bound holds for it when it is nonzero. $\square$

Proof of Lemma 3.67

Proof of 3.67. The self-crossings of a carrier are exactly its owned undominated labels, by Lemma 3.5(iii). If two such labels interlace, their alternating four visits cannot lie on two different carriers, by that lemma's noncrossing assertion. Applying this along an interlacement path puts every connected block on one owner.

Fix one nonempty block H . Starting at S , choose any label $c \in U(S) \setminus H$ and adjoin it to the support. It is adjacent to no selected label, so the enlarged support is independent. At any such step, with current support T , the new residual set is exactly

$$U(T \cup \{c\}) = U(T) \setminus (\{c\} \cup N_{G_P}(c)). \quad (250)$$

No original label outside H in $U(S)$ is adjacent to a label of H , since H is a connected component of the induced graph. Thus every step preserves all labels of H and removes at least the chosen label. The sets only shrink. After at most $|U(S) \setminus H|$ steps there is an independent support $S_H \supseteq S$ with $U(S_H) = H$.

Lemma 3.5 supplies actual geometric carriers for S_H . The preceding connectedness argument puts all labels of H on one of them. No other self-crossing label survives, so its actual positive diagram is a required D_H . Successor splitting retains the cyclic order inherited from the original traversal. The surviving crossing germs have not changed. Therefore the complete record of D_H is exactly the original record restricted to H , with its signs, pairing and over/under choices. This statement concerns an actually constructed diagram; it does not assume arbitrary Gauss subwords are realizable. Any other choices produce the same named record, so Lemma 3.20 proves independence of P_H .

For the original owner A , the connected components of its own interlacement graph are precisely the blocks owned by it: its self-crossing labels are their union and there are no graph edges between different blocks. The preceding construction has supplied an actual diagram for every one of these restricted records. Lemma 3.27 now constructs clean marked joins with the full record of D_A and proves the product in (53). Every retained crossing is positive and appears once, proving the crossing-count identity. If there are no owned blocks, D_A is an actual crossing-free one-circle diagram. Its local polynomial is 1 by the initialization in Theorem 3.18, and its crossing count is zero. No empty link is evaluated. $\square$

Proof of Lemma 3.68

Proof of 3.68. An undominated label interlacing c would have the same owner A by Lemma 3.67. The hypothesis therefore makes $\{c\}$ a singleton block of $G_P[U(S)]$. Since $c \in U(S)$, $S' = S \cup \{c\}$ is independent. Equation (250) and $N_{G_P}(c) \cap U(S) = \emptyset$ give $U(S') = U(S) \setminus \{c\}$.

Smoothing c splits A into two actual daughter carriers Λ_1, Λ_2 . Each other block owned by A lies wholly in one of the two intervals cut by c . Indeed none of its chords interlaces c , and a connected interlacement path cannot connect chords internal to different intervals. Successor splitting therefore assigns that whole block to exactly one daughter. Blocks owned by other carriers remain there. The singleton block's actual one-crossing diagram has polynomial 1 by Lemma 3.22

and positive writhe 1. Applying Lemma 3.67 before and after the split consequently gives

$$P_A = P_{\Lambda_1} P_{\Lambda_2}, \quad m_A = m_{\Lambda_1} + m_{\Lambda_2} + 1. \quad (251)$$

The two new smoothing turns are nonzero and opposite as real principal angles. Every old turn of A lies on exactly one daughter. Each daughter has at least three corners by Lemma 3.5, including its one new corner, so it has at least two inherited turns. One daughter is uniform with the old sign and the other has exactly one dissent. Summing principal angles gives $r_A = r_{\Lambda_1} + r_{\Lambda_2}$. Lemma 1.21, with orientation reversal if the old sign is negative, puts all three rotation integers on that old sign's ray. Hence

$$|r_A| = |r_{\Lambda_1}| + |r_{\Lambda_2}|. \quad (252)$$

Put $f_A = [z^0]P_A$ and $g_i = [z^0]P_{\Lambda_i}$. The knot support in Theorem 3.18 has no negative z exponents, so the first identity in (251) gives $f_A = g_1 g_2$. If $f_A = 0$, the desired coefficient is zero without assigning a degree to it. Otherwise both g_i are nonzero in the Laurent integral domain. Theorem 3.64 applies to the actual uniform and one-dissent daughters: they are subpolygons of the decomposition S' of P (Lemma 3.5(i)), the theorem's domain, and their turn patterns are the ones it requires. It gives $\mindeg_a g_i \geq 1 - m_{\Lambda_i} - |r_{\Lambda_i}|$ separately. Every exponent of the product is at least the sum of these bounds; coefficient cancellation cannot create a smaller exponent. Thus

$$\begin{aligned} \mindeg_a f_A &\geq 2 - (m_{\Lambda_1} + m_{\Lambda_2}) - (|r_{\Lambda_1}| + |r_{\Lambda_2}|) \\ &= 2 - (m_A - 1) - |r_A| \\ &= (1 - m_A - |r_A|) + 2 = d_A + 2. \end{aligned} \quad (253)$$

The coefficient at degree d_A is zero, which is exactly $c(A)$. □

Proofs for Section 4

Proof of Proposition 4.2

Proof of 4.2. First consider a continuous path of labelled generic polygons. By Proposition 1.11, the crossing set and the order of its visits along each directed edge are constant. Concatenating those edgewise lists identifies the oriented cyclic traversal records, their crossing pairings and interlacement graphs. It therefore gives a bijection of all independent sets S , preserving $|S|$.

Fix one such support S . Include the original vertices as marked points on the traversal circle and perform the same exchanges of outgoing successors at its selected visits. By Lemma 3.5, with the incoming-visit convention, this gives a fixed finite collection of oriented carrier cycles. Corresponding original vertices and smoothing corners have the same places on those cycles. An unselected crossing is retained as a self-crossing of a cycle exactly when both its visits belong to that cycle, so every carrier's retained crossing pairing and number m_Q are constant.

Each crossing point varies continuously: its two edge directions have nonzero determinant, and solving the two-by-two intersection equations gives continuous parameters in the relative interiors. Thus the corners of each named carrier form a continuous ordered polygonal tuple. Lemma 3.5 puts it in the regular locus: its segments are nonzero and no consecutive directions are antiparallel. Its signed rotation is consequently constant by Lemma 1.20(ii). At an original corner the turn sign is the corresponding original turn; at a smoothing corner it is the sign of the determinant of the inherited ordered edge directions. These determinants are continuous and nonzero along the path. All carrier corner signs, corner counts and uniformity conditions therefore remain constant, as does the original left-turn count $\ell(P)$.

At a retained self-crossing the two carrier directions are positive multiples of the same ordered parent directions throughout. Their nonzero determinant has constant sign. The physical over strand in the positive lift therefore stays over, and the crossing sign stays positive. Corresponding carrier lifts have a cyclic-order-, pairing-, sign- and over/under-preserving bijection of crossing visits. The printed scalar named-record Lemma 3.20, followed by the common substitution and Theorem 3.18, gives equal H polynomials, including for a carrier with no self-crossings. Since m_Q and $|\text{rot}(Q)|$ are constant, the exponent $d_Q = 1 - m_Q - |\text{rot}(Q)|$ is constant. Hence the coefficient $c(Q) = [a^{d_Q} z^0] H_Q^+$ is constant for every carrier, whether uniform or mixed. Every term and the prefactor of Definition 3.28 are now constant, proving constancy along the labelled path.

To pass to unlabelled chambers, first note that cyclically shifting the vertex labels preserves C directly from its definition. It merely renames the same directed traversal, crossings and independent supports; the corresponding carrier images and orientations, corner counts, rotations, positive lifts and coefficients are unchanged. Openness of the labelled generic locus gives each point a path-connected neighbourhood inside it. The path argument makes C locally constant there. Its cyclic invariance makes it descend to a locally constant function on the quotient by cyclic relabelling: the preimage of any value is an open saturated subset, so its image is open in the quotient topology. A locally constant function on a connected chamber is constant. This proves the stated assertion without assuming a lift of an arbitrary quotient path. $\square$

Proof of Proposition 4.3

Proof of 4.3. By Lemma 1.31(E),(C), the remote crossings and their cyclic visit order agree on both sides and at the centre. Its proof also gives the following geometric facts at the centre: vertices are distinct, all edges and all original turns are nonzero, no vertex lies on a nonincident closed edge, every actual crossing is transverse and interior, and distinct crossings have distinct points. In (E) the sole vertex-line incidence lies outside its closed base segment; in (C) its zero triple contains no

adjacent pair and hence gives no vertex-edge incidence. Thus the same finite crossing records and independent supports are identified through the entire small interval.

For each common support S , exchange the same successors to identify its carriers. Original vertices and selected smoothing points remain distinct, so the subsegments between successive marks have positive lengths at the centre and on a smaller interval. At an original corner the two directed segments are positive multiples of consecutive parent directions with nonzero turn determinant. At a smoothing corner they are positive multiples of the two transverse directions at its selected crossing, again with nonzero determinant. Omitting unselected marks only joins positive subsegments in the same direction. These observations verify directly, including at the nongeneric parent centre, that every carrier tuple is continuous, has nonzero segments, and has no antiparallel corner. They do not apply a generic-polygon lemma outside its hypothesis.

It follows from Lemma 1.20(ii) that each carrier's signed rotation is the same on both sides. Every original and smoothing turn sign is constant, and so are $\ell(P)$, uniformity and all corner counts. The common successor cycles preserve which unselected crossing visits share one carrier, hence all m_Q . At every retained crossing its direction determinant is nonzero and continuous through zero; therefore its positive over/under designation is also constant. Apply the scalar named-record equality of Lemma 3.21 and Theorem 3.18 to corresponding carrier diagrams on the two generic sides. This gives equal H polynomials; the equal m_Q and $|\text{rot}(Q)|$ give equality of their corner coefficients. The independent supports, their uniformity, $|S|$ and the common prefactor now identify the two sums term by term. There are finitely many supports and carriers, so all interval shrinkings can be made simultaneously.

Only carrier geometry was used at the centre; the state sums in (56) are evaluated on the generic sides. Proposition 4.2 makes them independent of the chosen representatives in the respective side chambers. $\square$

Proof of Theorem 4.4

Proof of 4.4. We first rewrite the actual state sum using only the corner signs. For a carrier Q with N_Q corners, define

$$W(Q) = \begin{cases} 1, & \text{all its turns are right,} \\ (-1)^{N_Q}, & \text{all its turns are left,} \\ 0, & \text{its turns are mixed.} \end{cases} \quad (254)$$

If a support S is uniform, the total number of left carrier corners is $\ell(P) + |S|$: every original vertex belongs to exactly one carrier with its original turn, and each selected crossing creates exactly one left and one right smoothing corner, by Lemma 3.5. It follows that the product of the selectors is $(-1)^{\ell(P) + |S|}$. If S is not uniform, at least one selector is zero. Consequently the definition of C gives, on every generic polygon,

$$C(P) = \sum_{S \in \text{Ind}(G_P)} \prod_{Q \text{ carrier of } S} W(Q) c(Q). \quad (255)$$

This is an identity for the coefficients of the individual actual carrier diagrams in Definition 3.28; no factorization of one diagram's polynomial into polynomials of other pieces has been used.

Put $D = P(0) \setminus j$. Lemma 1.30 proves that D is generic and identifies its crossings with those on both sides, preserving cyclic order, pairing and positive over/under data. Hence it identifies the three interlacement graphs and all independent supports, including supports whose uniformity differs between the configurations. For a fixed support, Corollary 3.16 gives corresponding oriented

carriers with the same retained crossing records and the same signed rotations. Lemma 3.21 and Theorem 3.18 therefore identify the H polynomials of their actual positive lifts. The numbers m_Q and $|\text{rot}(Q)|$ agree, so the exponents d_Q agree and each corresponding coefficient $c(Q)$ is the same in all three configurations. This applies also to mixed carriers and to crossing-free carriers.

Exactly one named carrier passes through the extra vertex μ_j . Call it Q_* and call its corresponding deletion carrier Q_*^D . All other corresponding selectors agree. The selectors of the distinguished carrier satisfy

$$W(Q_*^{\text{right}}) - W(Q_*^{\text{left}}) = W(Q_*^D) \quad (256)$$

by Corollary 3.16(iii). Explicitly, if the surviving corners are mixed, all three entries are zero; if all are right, the entries are $(1, 0, 1)$; if all k surviving corners are left, the entries are $(0, (-1)^{k+1}, (-1)^k)$. In the last case the first minus the second is $-(-1)^{k+1} = (-1)^k$. The deletion carrier has at least three corners, so these alternatives do not involve an empty corner list.

For the fixed common support S , let B_S be the common product of all carrier coefficients, and let U_S be the common product of all selectors other than the distinguished one (with empty product 1). Its right-minus-left contribution to (255) is

$$B_S U_S (W(Q_*^{\text{right}}) - W(Q_*^{\text{left}})) = B_S U_S W(Q_*^D). \quad (257)$$

The right side is exactly this support's contribution to the deletion sum. No division by B_S , U_S or a coefficient is performed; zeros cause no exception. Summing over the finite common support set proves (57). Proposition 4.2 supplies the well-defined side values. The only state sum on the right is that of the generic deletion, not a value assigned to the flat centre. $\square$

Proof of Theorem 4.6




Figure 16: The local positive-lift skein picture at a vertex-edge wall. Oriented smoothing separates the halves Q_L and Q_R in the wall term.

Proof of 4.6. We use the actual carrier selector sum of Lemma 3.31. The geometric domain is exactly the simple SM vertex-edge wall in the statement. Lemma 1.31 supplies a clean contact disc, stable old visits and the stated crossing changes. Lemma 1.34 supplies the actual generic ordered halves, each with at least three vertices. No additional guarded-factor or parallel-direction hypothesis is imposed.

In this proof write $\mathbf{m} = \mu_M$, $\mathbf{a} = \mu_a$ and $\mathbf{b} = \mu_{a+1}$ for the positions along the germ, suppressing its parameter. Evaluations at a specified side use that side's positions. The cyclic index M in the theorem remains unchanged. A local diagram is always the actual positive carrier diagram; its polynomial $P_A = H_A^+$ and write m_A are those in the state sum. The blocks and their owners are those of Definition 3.66; their products are proved in Lemma 3.67. Equality of named scalar records uses Lemma 3.20, including crossing-free components. Put $W(S) = \prod_L \text{wt}(L)$ and let

$\vartheta(u, v)$ denote the real principal angle from a nonzero vector u to v . All such local angles below are strictly between $-\pi$ and π .

Sliding: relocation and the exact selector difference.

Put $r = \mathbf{b} - \mathbf{a}$, $u_{\text{in}} = \mathbf{m} - \mu_{M-1}$ and $u_{\text{out}} = \mu_{M+1} - \mathbf{m}$ at the wall. The two neighbours are on opposite sides, while both tangents follow the oriented traversal through the contact. Therefore

$$\text{sgn det}(r, u_{\text{in}}) = \text{sgn det}(r, u_{\text{out}}) \neq 0. \quad (258)$$

By continuity from the relative-interior contact, each incident line meets the relative interior of the finite remote segment on a sufficiently small event interval. Exactly one incident segment reaches that line intersection on each side; the two sides exchange which one. Denote the crossing on the chosen side P_- by x_- and its relocated mate on P_+ by x_+ . Removing that label within the contact collar leaves the same cyclic word. Equation (258) also fixes which visit is the remote overpass. Thus the named map $\varphi(x_-) = x_+$, identity on old labels, preserves cyclic order, pairing, over/under, signs and the interlacement graph. It transports supports, undominated blocks, owners, actual carrier diagrams and writhes.

For a support avoiding this crossing, all successor cycles pass regularly through the wall and all corner determinants stay nonzero. Their rotation integers and selectors are fixed, and so are their coefficient reads. This sector cancels term by term.

For a support containing x_- , the two visits cut the traversal into the two half intervals. Every other selected chord is internal to one interval, by independence. Conversely any two half supports together with x_- form a support. This gives a bijection with explicit inverse

$$\{S \in \text{Ind}(G_{P_-}) : x_- \in S\} \longleftrightarrow \text{Ind}(G_{\lambda_1}) \times \text{Ind}(G_{\lambda_2}), \quad S \longmapsto (S_1, S_2). \quad (259)$$

Smoothing x_- closes the two intervals separately. All subsequent selected smoothings are internal, and labels interlacing x_- are dominated. Consequently

$$G_{P_-}[U(S)] = G_{\lambda_1}[U(S_1)] \sqcup G_{\lambda_2}[U(S_2)]. \quad (260)$$

The named maps identify every residual diagram and every owner. The only changed direction list, on the carrier containing the short edge, is one of

$$(r, u_{\text{in}}, u_{\text{out}}) \longrightarrow (r, u_{\text{out}}), \quad (u_{\text{in}}, u_{\text{out}}, r) \longrightarrow (u_{\text{in}}, r). \quad (261)$$

The two incident tangents lie in one open angular half-plane bounded by r . Thus their principal turns satisfy the real equalities

$$\vartheta(r, u_{\text{in}}) + \vartheta(u_{\text{in}}, u_{\text{out}}) = \vartheta(r, u_{\text{out}}), \quad (262)$$

$$\vartheta(u_{\text{in}}, u_{\text{out}}) + \vartheta(u_{\text{out}}, r) = \vartheta(u_{\text{in}}, r). \quad (263)$$

There is no hidden multiple of 2π . The chamber and half carriers have equal signed, and hence absolute, rotations. Writing $C_{\pm}(S)$ for the complete coefficient product on the two sides, and $C_i(S_i)$ for its half versions, we get

$$C_-(S) = C_+(\varphi S) = C_1(S_1)C_2(S_2). \quad (264)$$

The half contact signs $\eta_1 = \text{sgn det}(r, u_{\text{out}})$ and $\eta_2 = \text{sgn det}(u_{\text{in}}, r)$ are opposite. In P_- the short-edge refinement is on the half with contact sign s , and in P_+ on the one with sign $-s$. Let h_η

be that half-carrier weight, c_η its refined weight, and $\tau = \text{sgn det}(u_{\text{in}}, u_{\text{out}})$. The refinement adds the turn τ , so the full selector table is

$$c_\eta = \begin{cases} -\eta h_\eta, & \tau = \eta, \\ 0, & \tau = -\eta. \end{cases} \quad (265)$$

Indeed a mixed original carrier stays mixed; if τ disagrees with its contact sign it has both signs; if all signs are right the extra right corner preserves weight 1; and if all are left the extra left corner reverses $(-1)^{\#\text{corners}}$. These are all cases. Let Q_{sp} be the product of the other carrier weights, without assuming it nonzero. For the row in (259),

$$W_1 W_2 = Q_{\text{sp}} h_s h_{-s}, \quad W_- = Q_{\text{sp}} c_s h_{-s}, \quad W_+ = Q_{\text{sp}} h_s c_{-s}. \quad (266)$$

If $\tau = s$, then $c_s = -s h_s$ and $c_{-s} = 0$; if $\tau = -s$, then $c_s = 0$ and $c_{-s} = s h_{-s}$. Either way

$$W_+ - W_- = s W_1 W_2. \quad (267)$$

Multiply by (264) and sum over the finite bijection. Distributivity proves the stated product law for sliding, with no spectator cancellation.

Bigon words and the two-newborn sector.

Let P_0 be the side where M and its two neighbours lie on the same side of the remote line, and P_2 the other side, with two newborn crossings x, y . Set

$$s_0 = \text{sgn det}(r, \mathbf{m} - \mathbf{a})|_{P_0}, \quad \delta_{\text{dir}} = \begin{cases} +1 & P_- = P_0, \\ -1 & P_- = P_2. \end{cases} \quad s = \delta_{\text{dir}} s_0. \quad (268)$$

Write $\epsilon = 1$ when the newborn chords interlace and 0 otherwise. Starting at the first visit of x , the two complete local words are

$$\begin{array}{l|l} \epsilon = 0 & x_0 y_0 A y_1 x_1 B \\ \epsilon = 1 & x_0 y_0 A x_1 y_1 B \\ P_0 & A B. \end{array} \quad (269)$$

Each old label is internal to A , internal to B , or has one visit in each. The last set N is the common old neighbourhood of x, y . An old support T is called eligible when $T \cap N = \emptyset$. If T meets N , both newborns are dominated and neither extension by a newborn is a support. Deleting their four visits identifies the complete smoothing successor, owners, corners, rotations, carrier diagrams and coefficient reads. The entire ineligible returned sector is therefore zero.

For eligible T , restriction to the intervals gives the bijection

$$T \longleftrightarrow (T_1, T_2) \in \text{Ind}(G_{\lambda_1}) \times \text{Ind}(G_{\lambda_2}). \quad (270)$$

Internal labels in the two intervals have disjoint induced graphs; the halves close just these intervals. Conversely the union of two half supports is independent and misses N , proving the inverse.

If $\epsilon = 0$, $T \cup \{x, y\}$ is a support. Smoothing its four newborn visits makes one local three-corner contact triangle and otherwise exactly the successors of T_1, T_2 . Mixed old labels are dominated; each unmixed residual label, owner and decorated diagram matches its half. For $s_0 = +1$ the triangle has three right turns; for $s_0 = -1$ it has three left turns. Its complete data are

$$\text{rot} = -s_0, \quad R = 1, \quad P = 1, \quad w = 0, \quad [a^0 z^0]1 = 1, \quad \text{wt} = s_0. \quad (271)$$

Every two-newborn support arises this way. Put $J = sC(\lambda_1)C(\lambda_2)$. Finite distributivity and (268) show that the directed two-newborn sector is J when $\epsilon = 0$. When $\epsilon = 1$ it is empty. Hence, in both branches,

$$B = (1 - \epsilon)J. \quad (272)$$

It remains to show that the directed sum R_{ret} of all other noncancelling rows equals ϵJ .

Contact owners and one coambient smoothing diagram.

Fix an eligible newborn-free support T . Smooth its selected chords inside A . Choosing an innermost chord closes its inner subarc and splices it out of the remaining boundary strand. Induction leaves one open boundary successor and internal cycles. The same holds in B . The two open boundary successors and contact stretches form one full contact carrier L^* ; closing them at the cut gives ordered half contact carriers L_1, L_2 . All internal cycles are spectators.

For $i = 1, 2$, put $G_i = G_{\lambda_i}$ and $U_i = V(G_i) \setminus (T_i \cup N_{G_i}(T_i))$. Define the owner-restricted sets

$$O_i = \{c \in U_i : \text{the half owner of } c \text{ is } L_i\}. \quad (273)$$

For an internal label, the successor correspondence gives full owner L^* if and only if its half owner is L_i . Every surviving mixed old label belongs to $M_T = N \setminus N_{G_{P_2}}(T)$ and has owner L^* . Thus the full high contact residual set is precisely

$$O_1 \sqcup O_2 \sqcup M_T \sqcup \{x, y\}. \quad (274)$$

The labels $U_i \setminus O_i$ are on spectator carriers and are never inserted into a contact polynomial.

Let L_L, D_L be the low full contact carrier and its actual positive carrier diagram, and L_H, D_H its high counterpart. For this newborn-free support every selected crossing persists away from the contact disc. Its successor carrier through the contact follows the two original incident segments through M on both sides. The edge pieces have positive lengths at the centre, and all its corner turns have nonzero determinants there. Thus this carrier is a continuous regular family, including at the wall. Lemma 1.20(ii) makes its rotation locally constant. Consequently

$$\text{rot}(L_L) = \text{rot}(L_H) = \text{rot}(L^*). \quad (275)$$

The positive crossing $q = x$ of D_H has oriented smoothing equal to one actual ordered two-component diagram D_A . Component 1 contains the A -boundary successor and corresponds to L_1 ; component 2 contains the B -boundary successor and corresponds to L_2 . Retain every mixed crossing of this one ambient diagram.

If $\epsilon = 1$ define $D_A^{\text{post}} = D_A$. If $\epsilon = 0$, component 1 contains the permitted positive self R-I curl y ; delete exactly that curl, retaining the two ordered tags and every mixed crossing, to define D_A^{post} . Set

$$D_i = \text{component}_i(D_A^{\text{post}}), \quad Q_i = P_{D_i}, \quad w_i = w(D_i), \quad i = 1, 2. \quad (276)$$

The successor words identify these as the actual carrier diagrams of the ordered half contact carriers; their scalar identification is Lemma 3.20. The R-I deletion is in the supplied clean curl disc. The local LM R-I equality applies to the entire two-component diagram and to its component-1 restriction; component 2 is unchanged. It retains both component tags and every mixed crossing, so the coambient mixed signed sum and hence linking number are unchanged. Calling the latter ℓ , the full self/mixed crossing partition is

$$\begin{array}{l|l} \text{component 1 self crossings} & O_1 \cup (\{y\} \text{ if } \epsilon = 0) \\ \text{component 2 self crossings} & O_2 \\ \text{mixed crossings} & M_T \cup (\{y\} \text{ if } \epsilon = 1). \end{array} \quad (277)$$

In particular

$$P_{\text{component}_i(D_A)} = Q_i, \quad (278)$$

$$\text{lk}(D_A) = \text{lk}(D_A^{\text{post}}) = \ell, \quad (279)$$

$$w(\text{component}_1(D_A)) = w_1 + 1 - \epsilon, \quad w(\text{component}_2(D_A)) = w_2. \quad (280)$$

Here $\text{lk}(D_A)$ means the linking of its named ordered components, not linking assigned to two separately embedded restrictions.

Turn and selector bookkeeping.

Send r to the positive horizontal ray. If necessary reflect the picture to place both neighbour vectors above it, with arguments $\alpha, \beta \in (0, \pi)$. Direct line intersection gives

$$\epsilon = 0 \iff \alpha < \beta, \quad \epsilon = 1 \iff \beta < \alpha. \quad (281)$$

The half contact turns are β and $\pi - \alpha$; the full contact turn is $\beta - \alpha - \pi$ in the noninterlacing case and $\beta - \alpha + \pi$ in the interlacing case. Both half signs are s_0 ; the full sign is $-s_0$ or s_0 , respectively. Reflection reverses all signs together and transports these statements back. All noncontact turns partition between the halves. If Q_{sp} is the common product of internal spectator selectors, the actual factors are

$$W_T = Q_{\text{sp}} \text{wt}(L^*), \quad W_1 W_2 = Q_{\text{sp}} \text{wt}(L_1) \text{wt}(L_2). \quad (282)$$

No factor is cancelled in these identities.

For interlacing the contact signs agree. If any noncontact sign disagrees, both the full weight and the relevant half product are zero. Otherwise all three carriers are uniform with the same sign. For right turns each weight is 1 and $s_0 = -1$; for left turns the full corner count is the sum of the half counts minus one, giving a relative minus sign and $s_0 = +1$. Thus, including all zero cases,

$$\text{wt}(L^*) = -s_0 \text{wt}(L_1) \text{wt}(L_2), \quad W_T = -s_0 W_1 W_2. \quad (283)$$

If $W_T \neq 0$, (282) also shows that the spectator product and all three contact weights are nonzero, so the just-described uniform same-sign case applies. Multiplication by a possibly zero Q_{sp} proves the second identity without division.

For noninterlacing, a uniform full carrier gives each half one opposite contact turn, while two uniform halves give the full carrier one opposite turn. Thus at least one of their selectors is zero:

$$W_T W_1 W_2 = 0. \quad (284)$$

Every carrier used here has at least three corners. Two polygon corners would force overlapping straight segments, a polygon and smoothing corner would force vertex–edge incidence, and two smoothing corners would contradict transversality of the crossing straight branches. A one-corner closed straight path is likewise impossible. This also applies to the daughter carriers used below.

Subtract the explicit full contact turn from the sum of the two half turns, and add the unchanged noncontact turns. Dividing by 2π gives the signed identity

$$\text{rot}(L_1) + \text{rot}(L_2) - \text{rot}(L^*) = \begin{cases} 0 & \epsilon = 1, \\ s_0 & \epsilon = 0. \end{cases} \quad (285)$$

The second line contains a full-turn correction; unsigned angles would not give this formula.

Universal skein extraction and the interlacing branch.

For the fixed eligible row bind all polynomial and exponent objects:

$$F_L = P_{D_L}, \quad F_H = P_{D_H}, \quad F_A = P_{D_A}, \quad (286)$$

$$f_L = [z^0]F_L, \quad f_H = [z^0]F_H, \quad f_i = [z^0]Q_i, \quad (287)$$

$$w_L = w(D_L), \quad w_H = w(D_H), \quad R_i = |\text{rot}(L_i)|, \quad (288)$$

$$R_L = |\text{rot}(L_L)|, \quad R_H = |\text{rot}(L_H)|, \quad (289)$$

$$k_L = 1 - w_L - R_L, \quad k_H = 1 - w_H - R_H, \quad k_i = 1 - w_i - R_i, \quad (290)$$

$$\Omega_L = [a^{k_L} z^0]F_L, \quad \Omega_H = [a^{k_H} z^0]F_H, \quad \omega_i = [a^{k_i}]f_i, \quad (291)$$

$$K = k_1 + k_2. \quad (292)$$

These coefficient maps are total even when any Laurent polynomial is zero. The two incident directions in this bigon branch have opposite determinant signs against the remote direction. Their positive decorations therefore put the remote strand over at exactly one newborn crossing. Switching the positive contact crossing q makes the same strand over at both crossings, with opposite crossing signs. The isolated contact disc contains no other strand; this is the actual ordinary R-II template, not merely an opposite-sign pair. Delete it using local LM invariance. The remaining diagram has the same complete decorated record as D_L , including every old retained crossing; Lemma 3.20 gives its value F_L . The actual oriented smoothing is the tagged D_A . Thus the skein relation at q first gives $aF_H - a^{-1}F_L = zF_A$, then

$$F_H = a^{-2}F_L + a^{-1}zF_A. \quad (293)$$

The high writhe is $w_H = w_L + 2$ and (275) gives $R_H = R_L$, hence $k_H = k_L - 2$. Extracting $[a^{k_L-2} z^0]$ in (293) gives $\Omega_H = \Omega_L + [a^{k_L-1} z^{-1}]F_A$. Therefore

$$\Omega_H - \Omega_L = [a^{k_L-1} z^{-1}]F_A. \quad (294)$$

No selector or degree hypothesis is needed for this equality.

Assume first $\epsilon = 1$. The crossing partition (277) puts y and all retained old mixed crossings between the two components of D_A ; none belongs to Q_i or w_i . Lemma 4.5 and the writhe count yield

$$[z^{-1}]F_A = a^{-2\ell}(a - a^{-1})f_1 f_2, \quad w_L = w_1 + w_2 + 2\ell - 1. \quad (295)$$

The second equality follows also by counting w_H : the self crossings contribute $w_1 + w_2$, the mixed crossings of D_A contribute 2ℓ , and the smoothed crossing q contributes one; subtract the two newborn crossings to pass from high to low.

Let U_{sp} be the product of all spectator coefficient reads for this newborn-free row. Before using absolute-rotation or degree identities tied to this row, test W_T and U_{sp} . If $W_T = 0$, the two full terms vanish and (283) makes the half selector product zero. If $U_{\text{sp}} = 0$, both full terms and the half product contain that same spectator zero. Either case is finished. Assume now $W_T U_{\text{sp}} \neq 0$. The uniform same-sign implication proved in the turn calculation above and (285) imply

$$R_L = R_1 + R_2. \quad (296)$$

Substitution of (295) into the definition of k_L , followed by (296), gives

$$\begin{aligned} k_L &= 1 - (w_1 + w_2 + 2\ell - 1) - (R_1 + R_2) \\ &= (1 - w_1 - R_1) + (1 - w_2 - R_2) - 2\ell = K - 2\ell. \end{aligned} \quad (297)$$

In (294), the $a^{1-2\ell}$ term of the mixed row requests degree $k_L + 2\ell - 2 = K - 2$ in $f_1 f_2$; its $-a^{-1-2\ell}$ term requests degree $k_L + 2\ell = K$. Consequently

$$\Omega_H - \Omega_L = [a^{K-2}](f_1 f_2) - [a^K](f_1 f_2). \quad (298)$$

If either $f_i = 0$, both reads and $\omega_1 \omega_2$ are zero; no minimum degree is assigned. Otherwise both half contact carriers are uniform, and each is a subpolygon of a decomposition of its generic half λ_i (Lemma 1.34(ii), Lemma 3.5(i); the hypothesis discharge below), the domain of Theorem 3.64, which gives $\mindeg_a f_i \geq k_i$. In the Laurent integral domain the product has degree at least $k_1 + k_2 = K$. Its degree $K - 2$ coefficient is zero. In the degree K convolution any nonzero pair has exponents at least k_1, k_2 , so the only possible pair is exactly (k_1, k_2) . Thus

$$\Omega_H - \Omega_L = -\omega_1 \omega_2. \quad (299)$$

Both one-newborn sectors vanish by their selectors. To check this directly, take $\mathbf{m} = (0, 0)$, $r = (1, 0)$ on the line $y = c > 0$, and neighbour vectors $(u, h), (v, k)$ with $h, k > c$. Interlacing means $hv - uk > 0$. The corner determinants are

$$\det((-u, -h), (v, k)) = hv - uk > 0, \quad \det(r, (-u, -h)) = -h < 0, \quad \det((v, k), r) = -k < 0. \quad (300)$$

The first turn is at M ; the second or third is the smoothing turn when x or y is selected. Eligibility leaves the intervening collar free of selected visits, so the opposite turns lie on the same daughter carrier and make it mixed. Reflection reverses all three signs and proves the other orientation as well.

Let $\text{Term}_H(T), \text{Term}_L(T)$ denote complete newborn-free support contributions, before multiplying by δ_{dir} . In the nonzero branch (299) gives

$$\text{Term}_H(T) - \text{Term}_L(T) = -W_T U_{\text{sp}} \omega_1 \omega_2. \quad (301)$$

The same equation holds in the already discharged zero-factor branches, with both sides zero. Substitute $W_T = -s_0 W_1 W_2$, multiply by δ_{dir} , and use $s = \delta_{\text{dir}} s_0$. The directed returned share is $s W_1 W_2 U_{\text{sp}} \omega_1 \omega_2$. The owner-restricted decomposition puts each internal spectator into this half product exactly once. Summing over (270) gives

$$R_{\text{ret}} = s C(\lambda_1) C(\lambda_2) = J. \quad (302)$$

This proves the required returned-sector identity for $\epsilon = 1$.

Noninterlacing: full-twist, singleton and separate-block rows.

Now $\epsilon = 0$. The two-newborn sector already equals J by (272); we show every remaining returned row is zero.

The same-block alternative.

In D_A the remaining newborn y is a positive self curl on component 1. Its deletion gives the exact Q_1, Q_2 of (276), not polynomials with spectator factors inserted. Counting self and mixed crossings gives

$$w_L = w_1 + w_2 + 2\ell. \quad (303)$$

Indeed w_H includes the additional q and the self curl y , for a total of $w_1 + w_2 + 2\ell + 2$. If the full selector or an actual spectator multiplier is zero, both newborn-free terms vanish and the row stops before absolute-rotation identities. Otherwise the full carrier is uniform with sign $\tau = -s_0$. Every half inherits its old turns and has exactly one new contact turn of sign s_0 , hence exactly one

dissent. The at-least-three-corner observation ensures inherited turns exist. Lemma 1.21, reversing all orientations when $\tau = -1$, puts the full and both half rotations on the τ ray. Multiplying (285) by τ yields

$$R_1 + R_2 - R_L = \tau s_0 = -1. \quad (304)$$

Let $\Delta_R = R_1 + R_2 - R_L$. The slot definitions give separately

$$\begin{aligned} K - k_L &= 2 - w_1 - w_2 - R_1 - R_2 - (1 - w_L - R_L) \\ &= 1 + 2\ell - \Delta_R = 2\ell + 2. \end{aligned} \quad (305)$$

Combining the universal extraction (294) with the same two-component row of Lemma 4.5, its two requested degrees are now $k_L + 2\ell - 2 = K - 4$ and $k_L + 2\ell = K - 2$. Therefore

$$\Omega_H - \Omega_L = [a^{K-4}](f_1 f_2) - [a^{K-2}](f_1 f_2). \quad (306)$$

If either Laurent row f_i is zero, both reads are zero immediately. Otherwise apply Theorem 3.64 to the two one-dissent half contact carriers, subpolygons of a decomposition of the generic side (Lemma 3.5(i); the hypothesis discharge below); it puts their nonzero product row at degree at least K . Both displayed coefficients are below this floor. Thus the returned newborn-free row is zero. The auxiliary half selectors may themselves be zero: they are not multipliers of this full term, and the floor requires their one-dissent turn patterns, not uniform selectors.

For eligible T in the noninterlacing branch, $T \cup \{x\}$ is a support since x is adjacent to neither T nor y . Every old neighbour of y is also a neighbour of selected x , and hence is dominated. Thus $\{y\}$ is an isolated block; the statement with x, y exchanged is identical. If a one-newborn selector is zero its term vanishes. Otherwise its owner is uniform, and Lemma 3.68 makes the required coefficient zero, including the zero-polynomial case. Every one-newborn term is zero.

The different-block alternative.

If x, y lie in different blocks of the newborn-free high row, both components must be singletons. Otherwise an undominated old common neighbour gives a path of length two joining them; the common-neighbour property follows directly from (269). Each singleton has polynomial 1 and writhe 1. Every other high block is old and is matched to its low block by erasing the finger. For $S = T \cup \{x, y\}$,

$$U(S) = U(T) \setminus \{x, y\}. \quad (307)$$

The triangle and half-factorization construction in the two-newborn calculation above now distributes every old contact block once, with its actual owner, between L_1, L_2 . It follows that

$$f_H = f_L = f_1 f_2, \quad w_H = w_1 + w_2 + 2, \quad w_L = w_1 + w_2. \quad (308)$$

If the common full selector or a spectator multiplier is zero, both newborn-free terms vanish before the row-tied absolute rotations are used. Otherwise the same noninterlacing turn computation gives (304). Thus

$$\begin{aligned} k_L &= 1 - (w_1 + w_2) - R_L \\ &= K - 1 + (R_1 + R_2 - R_L) = K - 2, \end{aligned} \quad (309)$$

$$k_H = k_L - 2 = K - 4. \quad (310)$$

If $f_1 = 0$ or $f_2 = 0$, (308) makes both full rows zero without a degree argument. Otherwise the actual one-dissent half floors give product degree at least K . Both slots $K - 2, K - 4$ are below it, so

both reads vanish separately. The one-newborn rows already vanished by Lemma 3.68. The same-block and different-block alternatives exhaust the residual components containing the newborns. Therefore

$$R_{\text{ret}} = 0 \quad (\epsilon = 0). \quad (311)$$

Complete hypothesis discharge and assembly.

Every floor polynomial above is that of an actual positive carrier one-circle diagram supplied by Lemma 3.5 and Lemma 3.67, and every carrier to which Theorem 3.64 is applied is a subpolygon of a decomposition of a generic polygon (Lemma 3.5(i)), the theorem's domain. For the two half contact carriers that polygon is the generic half (Lemma 1.34(ii)), at either value of ϵ : the traversal circle of λ_1 , respectively λ_2 , is the interval A , respectively B , closed at the cut $\mathbf{m} = \mu_M$ (Definition 1.33: λ_1 follows $E_M, \dots, E_{a-1}$ and closes with $[\mu_a, \mu_M] \subset E_a$; λ_2 opens with $[\mu_M, \mu_b] \subset E_a$, follows $E_b, \dots, E_{M-2}$ and closes with E_{M-1}); T_i is a support of λ_i all of whose selected visits are internal to that interval (the bijection (270)); the oriented reconnections of Definition 3.2 at those visits act inside the interval and leave the cut untouched, so closing the interval at the cut closes the boundary successor of T into the carrier of the decomposition T_i of λ_i through the corner $\mathbf{m}$ — which is how L_i was formed — with the internal cycles as the other carriers of T_i (Lemma 3.5(i)); and the self-crossings of L_i are the owner-restricted set O_i of (273) (Lemma 3.5(iii)). When $\epsilon = 1$ the enlarged supports $T \cup \{x\}$, $T \cup \{y\}$ and $T \cup \{x, y\}$ of the generic side do not take the place of this discharge: $T \cup \{x, y\}$ is not a support, the newborn chords interlacing, and the daughter of $T \cup \{x\}$ or of $T \cup \{y\}$ that carries the turn at $\mathbf{m}$ also carries the opposite smoothing turn of (300) and is mixed, while a live half contact carrier is uniform. In the noninterlacing case the two half contact carriers are the carriers of the newborn-free support T enlarged by the smoothed newborns — of $T \cup \{x, y\}$ when $\epsilon = 0$, when the third carrier is the crossing-free local contact triangle and the deletion of the curl y from component 1 of D_A is that smoothing with the triangle discarded — which is what the closure at the cut and the successor words of (276) record; the daughters of a singleton row are the carriers of the support enlarged by the singleton (Lemma 3.68); and a crossing-free carrier is a carrier of its own row's decomposition. The construction in Theorem 3.64 supplies the needed clean rounding and curl discs; no separate arc-contraction or arbitrary-front lift is presumed. Its turn hypotheses are exhausted as follows: live interlacing rows have two uniform half contact carriers; live noninterlacing rows have two one-dissent half contact carriers; singleton rows have a uniform and a one-dissent daughter; and a crossing-free carrier uses its separately proved $P = 1, w = 0$ branch.

A zero selector or spectator which actually multiplies the current full or target term stops that row before a row-dependent absolute-rotation or degree identity. A zero Laurent row is checked before every minimum-degree or coefficient-floor argument. Independently valid geometric rotation identities do not require nonzero polynomials, and coefficient extraction itself is always total. In particular an auxiliary one-dissent half or daughter can have zero canonical selector: it is not a multiplier in the current row, and its actual turn pattern is precisely a permitted floor hypothesis. No spectator, selector, polynomial or coefficient has been divided out anywhere in the proof.

For a bigon, (272), (302) and (311) give, with all factors already bound,

$$B + R_{\text{ret}} = (1 - \epsilon)J + \epsilon J = sC(\lambda_1)C(\lambda_2). \quad (312)$$

The sliding calculation above gives the same result for sliding. The nonzero neighbour-side signs are either equal or opposite; these alternatives are exclusive and exhaustive. Hence the vertex–edge law holds on its full stated domain. The three positions of every old label—internal to A , internal to B , or mixed—retain all exterior crossing and threading data. All support bijections have the

explicit inverse union operation, and every actual spectator occurs once. There is no restriction to an unthreaded drawing or to a preferred choice of the initial side. $\square$

Proof of Theorem 4.7

Proof of 4.7. Let the cusp be at vertex j of an n -gon, $n \geq 4$, with tail-labelled edges $E_i = [\mu_i, \mu_{i+1}]$. Choose a generic polygon P sufficiently near the wall on its no-loop side. In cusp case A, where $\mu_{j+1}(0)$ lies strictly between $\mu_{j-1}(0)$ and $\mu_j(0)$, put $(c_1, c_2) = (j, j+1)$. In case B, where $\mu_{j-1}(0)$ lies strictly between $\mu_j(0)$ and $\mu_{j+1}(0)$, put $(c_1, c_2) = (j-1, j)$. In both cases the corners c_1, c_2 are consecutive, with intervening edge E_{c_1} .

By Lemma 1.32(ii), their turns on P satisfy

$$\tau_{c_1}(P) = -\tau_{c_2}(P), \quad \tau_{c_1}(P), \tau_{c_2}(P) \in \{-1, 1\}. \quad (313)$$

The cusp is empty, meaning that the two newborn crossing visits on the loop side are cyclically adjacent. Lemma 1.32(iii) therefore gives that E_{c_1} has no crossing on P . This conclusion uses the lemma's specified-short-arc argument; adjacency alone is not silently assigned to one of the two complementary traversal arcs.

Fix any decomposition $S \in \text{Ind}(G_P)$. In the traversal circle, take the closed directed arc from the visit of μ_{c_1} to the next vertex visit μ_{c_2} , tracing the single edge E_{c_1} . Its interior has no crossing visit, by the preceding paragraph. Its endpoints are not crossing visits either: a generic polygon has distinct vertices and no vertex on a nonincident edge. Consequently this arc contains no visit of a selected crossing of S .

Oriented smoothing changes the traversal only at selected crossing visits (Definition 3.2). Every directed segment of the chosen arc, including its passages into the two vertex visits, is therefore retained with the same direction. More explicitly, subdivide the traversal circle at all polygon-vertex visits and all selected crossing visits. At a selected pair smoothing interchanges the two outgoing continuations; it leaves every other local continuation unchanged. The arc joining c_1 to c_2 has no selected subdivision point, so it remains a directed path in the reconnected graph. Each vertex of that graph has one incoming and one outgoing edge, hence the graph is a disjoint union of directed cycles. The path lies on one such cycle. Thus both original vertices belong to the same subpolygon, say Q , of the decomposition S .

At each of these original vertices, smoothing has changed neither of the two incident edge directions. Both remain actual corners of Q , since their original turn signs in (313) are nonzero. Their turns on Q are the same as on P , so one is left and the other right. Therefore Q is mixed, and S is not uniform (Definition 3.6). This argument also covers $S = \emptyset$, when the one subpolygon is P itself.

Since S was arbitrary, no decomposition of P is uniform. The sum defining C in Definition 3.28 has an empty index set, so

$$C(P) = (-1)^{\ell(P)} \sum_{\substack{S \in \text{Ind}(G_P) \\ S \text{ uniform}}} (-1)^{|S|} \prod_{Q' \text{ subpolygon of } S} c(Q') = 0. \quad (314)$$

No corner coefficient has been evaluated or divided out.

Finally, opposite turns at these consecutive vertices and absence of crossings on their connecting edge depend only on the turn word and the crossing set. Those data are constant on each labelled chamber (Proposition 1.11). Hence the same argument applies to every polygon in that labelled no-loop chamber. The existence of the pair and the argument are preserved by cyclic relabelling,

so they also apply to its polygon chamber. This proves $C(P_{\text{no}}) = 0$ without needing a separate polynomial-invariance argument. $\square$

Proof of Theorem 4.8



Figure 17: Local carrier pictures for soft insertion in the corner state sum, accompanying the retained and smoothed terms in the proof.

Proof of 4.8. Write $M = \mu_j$, $M_\varepsilon = M + \varepsilon q$, $u = \tilde{\lambda}_{j-1}$, $v = \tilde{\lambda}_j$, $D_\varepsilon = v - \varepsilon q$ and $\tau = \tau_j(P)$. Lemma 2.23 supplies one initial generic interval. By Proposition 4.2, C is constant on that chamber, so it suffices to prove the identity on a smaller initial interval. All choices below can hold simultaneously because there are finitely many supports and carriers.

Carrier records and coefficients. Use the marked successor model and inherited orientations of Lemma 3.5. For corresponding carriers with the same retained crossing visits, the cyclic successors and pairings agree. Each carrier crossing germ follows its containing edge of its own ambient polygon in the inherited direction, up to a positive scalar multiple. Lemma 2.23(iii) preserves the determinant sign of each ordered pair of those directions. By Definition 3.8, the same corresponding visit is therefore over, and every crossing sign is positive. These are isomorphic named records, including the correspondence of oriented crossing-free carriers when the visit sets are empty. Lemma 3.20 equates their source polynomials; the common Laurent substitution and Theorem 3.18 identify these with equal H^+ -polynomials. If the corresponding rotations also agree, their crossing counts m_Q , rotations r_Q and $d_Q = 1 - m_Q - |r_Q|$ agree. Definition 3.28 then gives equal corner coefficients $c(Q)$. We now establish each needed carrier correspondence and rotation assertion.

Same-sign sector: $\chi_- = \chi_+ = -\tau$. The inherited Gauss words agree by Lemma 2.23(iv), so the independent supports are precisely the same sets S . In the successor cycles of Lemma 3.5(i), contracting the inserted ordinary vertex M_ε identifies each carrier with its parent carrier. Exactly the carrier through M gains that vertex. Its old turn τ is replaced by two turns τ, τ , by Lemma 2.23(ii). All other original and smoothing corner signs persist by clauses (ii),(iii) and Lemma 3.5(ii). Thus uniformity is equivalent on the two sides of this correspondence.

For rotation, let θ be the old principal turn from u to v , and let α be that from u to q . The same-sign determinant conditions put q in the oriented turn wedge: $0 < \alpha < \theta < \pi$ if $\tau = 1$, and $-\pi < \theta < \alpha < 0$ if $\tau = -1$. Indeed the two strict half-plane inequalities choose precisely the intersection between the rays of u, v inside their principal angular sector. The two new principal turns tend to α and $\theta - \alpha$, whose sum is θ , with no wrap. Every other carrier corner has directions converging to its old regular directions, including smoothing corners involving the return edge and the successor corner at $B = \mu_{j+1}$. The full turn sum therefore converges to its parent turn sum. Lemma 3.5(ii) gives regularity of the nearby carriers, and Lemma 1.20(i) makes their rotations integers. Convergence to the old integer implies equality for all sufficiently small parameters. This applies to every carrier, including geometrically moving carriers that do not pass through M .

Their retained visit records agree by the successor correspondence and Lemma 2.23(iii). The preceding coefficient comparison gives equal products of $c(Q)$. Support sizes do not change, while

$\ell(P_\varepsilon) = \ell(P) + [\tau = 1]$. Substituting in Definition 3.28 gives $C(P_\varepsilon) = (-1)^{[\tau=1]}C(P) = -\tau C(P)$, the required multiplier in this sector.

Mixed sector: $\chi_- \neq \chi_+$. The soft edge has no crossing by Lemma 2.23(i). Every selected reconnection therefore leaves the directed path from M to M_ε intact. These are consecutive corners of one carrier, and their turns $-\chi_-$ and $-\chi_+$ are opposite. No support is uniform, so Definition 3.28 gives $C(P_\varepsilon) = 0$, the required mixed-sector value.

Loop sector: $\chi_- = \chi_+ = \tau$. Let y be the newborn crossing. Its visits are adjacent in the Gauss word by Lemma 2.23(iv), so it interlaces no crossing. The independent supports are exactly S and $S \cup \{y\}$, with $S \in \text{Ind}(G_P)$. For a support omitting y , Lemma 3.5(iii) puts both visits of y on one carrier Q' . Its short arc through M, M_ε has no other crossing visit; hence y remains isolated in the interlacement graph of Q' . If the decomposition is not uniform its term is excluded. If it is uniform, Lemma 3.70(ii) gives $c(Q') = 0$. Thus every support omitting y contributes zero, including those whose corresponding parent support was not uniform.

For $S \cup \{y\}$, process y first, using order independence in Lemma 3.5(i). Let a be its visit on the incoming edge and b its return-edge visit. The clean short arc is $a \rightarrow M \rightarrow M_\varepsilon \rightarrow b$. Swapping the two outgoing successors produces the triangle cycle (b, M, M_ε) and the residual cycle containing a . The triangle smoothing turn follows D_ε into u and has sign $-\tau$; the turns at M, M_ε also have sign $-\tau$. The three noncollinear segments form an embedded triangle, so Lemma 3.70(i) gives its coefficient 1.

All old crossing visits lie on the residual cycle. It replaces the old corner M by the smoothing corner y , whose incoming and outgoing directions are u, D_ε , of turn sign τ by Lemma 2.23(ii). The remaining swaps S act only on that residual cycle and give carriers in bijection with the carriers of S on P . Their original and smoothing corner signs agree; at the distinguished carrier, the corner M is replaced by y of the same sign. Consequently $S \cup \{y\}$ is uniform exactly when S is uniform. No (G1) claim about the intermediate residual polygon is needed: all resulting curves are carriers of the generic P_ε .

As $y \rightarrow M$ and $D_\varepsilon \rightarrow v$, every residual principal turn converges to its old regular counterpart. Integer-valuedness from Lemma 1.20(i) again gives equality of all rotations on a common small interval. Their retained crossing successors and pairings agree under the inherited correspondence, and the corresponding directed germs retain their ordered determinant signs, so the earlier named-record argument gives equal positive-lift polynomials and corner coefficients. This also covers all other corresponding carriers, whose positions may have moved.

Every surviving term thus gains exactly one selected crossing and one triangle coefficient 1. The original left-turn count changes by $\ell(P_\varepsilon) - \ell(P) = -[\tau = 1] + 2[\tau = -1]$. Definition 3.28 gives

$$C(P_\varepsilon) = (-1)^{-[\tau=1]+2[\tau=-1]+1}C(P).$$

For $\tau = 1$ the exponent is zero; for $\tau = -1$ it is three. The multiplier is therefore τ , as required. The three sector multipliers are exactly $(\chi_- + \chi_+)/2$, completing the proof on the unchanged domain. $\square$

Proof of Lemma 4.11

Proof of 4.11. Write $T = \{x_{ef}, x_{eg}, x_{fg}\}$ for the three crossing labels. At the central triple point q , the three edge directions are pairwise nonparallel: otherwise two supporting lines coincide, forcing three original vertices to be collinear, contrary to $Z_{\text{pt}} = \emptyset$. The three direction determinants are therefore nonzero and retain their signs near the centre.

Choose a closed event disc about q whose intersection with the central polygon consists only of these three strand segments, with their endpoints on its boundary. The point q is interior to



Figure 18: A transitive triple configuration in the positive lift. Its local over/under relation is a total order, as in Lemma 4.11(i).

the three edges and is not an original vertex. A fourth edge through q would be remote from the first three, since adjacent edges under (G1) meet only at their common endpoint; it would give another triple in Z_c , a contradiction. The other finitely many compact edges thus have positive distance from q . Choose the disc radius small enough to exclude them and all original vertices, with its boundary transverse to the three strands. After shrinking the parameter interval, the disc still contains exactly these three strand segments, all three crossings of T , and no original vertex or other crossing. Lemma 1.28 gives adjacent pairs of local visits. Lemma 1.31, using the sign-change condition in Definition 1.29(T), reverses each local pair across the wall and preserves every other visit order.

Let Q_- be any subpolygon on one side retaining all three crossings, and fix its decomposition S . Definitions 3.1 and 3.2 give $S \cap T = \emptyset$. Delete the six visits of T from both parent traversal lists. The remaining cyclic lists of all vertices and other crossing visits are identical. In particular all selected pairs have the same cyclic order, so S is independent on the other side as well.

Apply the successor description of Lemma 3.5(i). Every deleted visit is unselected, hence has its ordinary outgoing traversal arrow. Suppressing these visits commutes with all selected successor swaps: following the unchanged arrows until the next remaining mark has the same endpoint on both sides. The resulting permutations on the remaining marks are identical, and their cycles identify the carriers on the two sides. No entire carrier is suppressed: a cycle with only unselected visits follows the original traversal and reaches an original vertex, which has not been deleted. Reinserting each local pair merely subdivides the same oriented arc of the corresponding cycle, with that pair in the reversed order. The corresponding carrier Q_+ therefore retains all three crossings. The disc has no selected smoothing site or original vertex, so both carriers contain all three entire local strand segments with their inherited orientations.

Lemma 3.5(ii),(iii) makes these valid carrier diagrams on both punctured sides. Each local directed germ is a positive multiple of its parent edge direction. By Definition 3.8, its strand from E_i passes over that from E_j exactly when $\det(u_i, u_j) > 0$. Thus either carrier's local tournament is the parent's determinant tournament.

For this one carrier, the exterior diagrams may be identified by planar isotopy. Regard its retained crossings as four-valent vertices, its original and smoothing corners as two-valent vertices, and the six transverse event-boundary intersections as boundary vertices. At a selected crossing this carrier has only one corner, by Lemma 3.5(i); corners belonging to other carriers are not part of this exterior graph. Away from the event disc, the graph of the central carrier — the carrier of S traced on the central polygon $P(0)$, of which the carriers on both sides are small perturbations — has distinct such vertices, noncollapsed arcs and no degenerate incidence: the only parent triple is inside the disc, (G1) excludes vertex incidences, and all other intersections are transverse, with distinct crossing parameters on each edge. Choose disjoint small vertex neighbourhoods, thin strips around the intervening embedded arcs, and collars of the six boundary endpoints. Their cyclic attachment orders persist under a sufficiently small perturbation. Straightening in these neighbourhoods, strips

and collars gives an ambient planar isotopy of the disc exterior, preserving its boundary setwise while moving the six marked endpoints to their counterparts. The nonzero direction determinants at persistent crossings preserve the physical positive over/under choices during this identification. Extending the boundary identification across the disc only identifies coordinates; the two interior triangle drawings are still compared by the local move below.

A tournament on three strands is a total order unless it is a directed cycle. It is cyclic exactly when $\det(u_e, u_f)$, $\det(u_f, u_g)$ and $\det(u_g, u_e)$ have one sign, equivalently when $(\delta_1, \delta_2, \delta_3)$ is $(+, -, +)$ or $(-, +, -)$. In all other cases call the three strands top, middle and bottom according to their total order. The local triangle interchange has the top strand over both others and the bottom strand under both others, so it is the oriented Reidemeister III move. This proves (i) for the corresponding carrier diagrams.

In a cyclic case the determinant signs still persist, so the cyclic over/under pattern is the same on both sides. The three reversed local visit pairs give opposite triangle configurations. This proves (ii), with no cyclic knot-invariance conclusion. $\square$

Proof of Proposition 4.13

Proof of 4.13. We use the actual carrier sum of Lemma 3.31. For a carrier L write $P_L = H_L^+$, $w_L = m_L$, $R_L = |\text{rot}(L)|$, $d_L = 1 - w_L - R_L$, and $\omega_L = [a^{d_L} z^0]P_L$. The term of a support is the product of $\text{wt}(L)\omega_L$ over its carriers. All diagrams below are actual carrier diagrams or their explicitly specified switched or smoothed diagrams. Products of owned block polynomials are justified by Lemma 3.67; equal decorated records give equal polynomials by Lemma 3.20. This does not identify the present definition with an unproved state sum.

Local geometry and (ii).

Lemma 1.28 gives a fixed crossing set and three adjacent pairs of local visits. Lemma 1.31 gives their reversal across this simple wall and fixes all other visit orders. Choose the event disc small enough to contain only the three strand segments: its central point is interior to them, no vertex lies there, and another strand through that point would violate simplicity. The three shrinking segments between local visits contain neither a vertex nor another crossing.

An adjacent transposition of visits of two distinct crossings changes only their mutual interlacement. Indeed it moves one visit past only the other; every other pair keeps its four-point cyclic order. For the two named crossings, exactly one endpoint passes through the boundary of the interval cut by the other chord, so alternation toggles. There are three such transpositions, one for each local pair. Thus the local graph is complemented and every other edge is unchanged. On three vertices the complementary graphs have edge counts 0, 3 or 1, 2; the latter graphs are an edge and a path. This proves (ii).

Rotations and (i).

Every selected crossing that persists gives two corners with nonzero direction determinants near the centre. Polygon corners likewise remain nonzero by (G1). Carriers with fixed successor patterns and no collapsed edge between selected local crossings are therefore regular families through the wall. Their individual directions may change. Their principal turns vary continuously without meeting $\pm\pi$; their integer rotation numbers are constant by Lemma 1.20. Their corner signs and weights are constant as well. For the support in (i), no local crossing is selected. The carrier containing the triple undergoes the transitive oriented RIII change of Lemma 4.11(i). Theorem 3.18 — Reidemeister-III invariance of the local polynomial, from Literature input 2, and its identification with H — therefore preserves its polynomial. It has the same retained crossing labels, hence the same writhe; the preceding regular-family argument preserves its rotation and weight. Every other

carrier has the same decorated record, writhe, rotation and weight. Multiplying their coefficient reads proves (i), also when a factor is zero.

Availability and the outside factor.

For (iii) fix the stated outside independent support Q and write T for the three local crossings. We prove the identity already under the weaker premise that *every selected member of Q* interlaces no member of T . The printed premise of (iii), which quantifies over every outside crossing, implies this premise. We do not change that printed statement.

For any outside crossing y , its two visits avoid the interiors of the three adjacent two-visit clumps. The two arcs cut out by y thus colour the three whole clumps in two colours. A member of T interlaces y exactly when its two clumps have different colours. Three objects have either zero or two bichromatic pairs. Consequently the local mask $I(y) = N(y) \cap T$ has size zero or two. It is fixed across the wall by (ii). More generally,

$$\text{avail}(Q) = T \setminus \bigcup_{y \in Q} I(y)$$

has size three, one or zero, since the union of distinct two-element subsets of a three-element set is the whole set. Only size three is used in the proof that follows.

Smooth Q first. A resulting carrier containing no local visit, including a local smoothing-site visit, survives any subsequent independent local smoothings arc for arc. Conversely undoing those local reconnections on a triangle-disjoint carrier changes no half-edge of that carrier. This is a bijection of the exterior carriers in all rows. Their self-crossing sets also stay fixed. To see this, an undominated exterior crossing adjacent to an available member of T would belong to its connected block and hence to its owner, by Lemma 3.67. That exterior owner would contain local visits, a contradiction. Adding local selections removes only those labels and their neighbours, by (250); it therefore removes no exterior self-crossing.

Erasing the six local visits identifies these exterior successor cycles across the wall. Their decorated records, positive writhes and corner signs agree; the regular-family argument gives equal rotations. Their polynomials and coefficient slots agree as well. Thus all rows have one common exterior factor C_Q , which may be zero. For a side ν write $T_\nu(J)$ for the *complete* term of $Q \cup J$, setting it to zero when that support is absent. We only multiply by C_Q ; we never divide by it.

Each member of Q has mask zero. Its visits lie in one of the three gaps between the local clumps: the three gaps have distinct vectors of sides relative to the three local chords. Smoothing Q therefore acts separately inside each gap. A printed gap below means its boundary-to-boundary successor path after those smoothings; internal closed successor cycles are precisely the exterior carriers just handled. This supplies the same local word calculation for arbitrary Q , not just for $Q = \emptyset$.

Sign classification.

In cyclic strand order put $a = (u_1, u_2)$, $b = (u_1, u_3)$, $c = (u_2, u_3)$, and let s_a, s_b, s_c be their determinant signs. Writing δ for the sign of the three-line determinant, Cramer's rule for the intersection parameters gives the three order signs

$$(q_1, q_2, q_3) = -\delta(s_a s_b, s_a s_c, s_b s_c). \quad (315)$$

For example, if t_{ij} is the parameter on line i of its intersection with line j , the numerator of $t_{12} - t_{13}$ over the common denominator $\det(u_1, u_2) \det(u_1, u_3)$ is the negative three-line determinant. The other two rows follow by expanding $t_{21} - t_{23}$ and $t_{31} - t_{32}$ with the same fixed determinant order. The local graph indicators, read from the three consecutive clumps, are respectively $q_1 = -1$,

$q_2 = +1$, $q_3 = -1$ for ab, ac, bc . They all agree precisely when $s_a = s_c = -s_b$. These are the two extreme sign branches; the other six are generic.

A pair sharing a strand u is called separating when the other two directions lie on opposite sides of u . For pairs ab, ac, bc this requires, respectively, $s_a = -s_b$, $s_a = s_c$, $s_b = -s_c$. In a generic branch exactly one holds. Equation (315) identifies that pair as the independent pair complementary to the path's middle vertex. These statements follow in each of the three sign pairs

(s_a, s_b, s_c)	separating pair	other pairs
$+++ \text{ or } ---$	ac	ab, bc
$++- \text{ or } --+$	bc	ab, ac
$+-- \text{ or } -++$	ab	ac, bc

and do not presume a symmetry of the outside gaps.

Generic orbit: words and common rows.

Choose the ordering of the three directions from their transitive over-order or its reverse that is even relative to their cyclic traversal order. A cyclic relabelling then puts the separating pair at ac and gives $s_a = s_b = s_c = \sigma$, where $\sigma = \pm 1$. No strand orientation is reversed. The words on the path and edge sides are

$$P = abAacBbcC, \quad E = baAcaBcbC. \quad (316)$$

The gaps A, B, C are carried along with this relabelling. The entire support and local-undominated table is

$P : J$	$\emptyset$	a	b	c	ac
$U(Q \cup J) \cap T$	abc	c	$\emptyset$	a	$\emptyset$
$E : J$	$\emptyset$	a	b	c	ab, bc
$U(Q \cup J) \cap T$	abc	b	ac	b	$\emptyset$

where ac is connected in the E, b row. All unlisted supports are absent. The empty row's contact carrier has successor string ABC on both sides. Its local lift has transitive over-order and undergoes ordinary RIII. Its polynomial, writhe, rotation and weight therefore agree exactly as in (i); no bijection of the changing local block partition is required. Hence $T_P(\emptyset) = T_E(\emptyset)$.

For endpoint a the carriers are $BC \mid A$, with residual local word $cBcC$ on P and $bBbC$ on E . Outside survivors after a have mask zero or bc . Thus b, c are twins with respect to those survivors; fixing all outside labels and replacing c by b identifies the two undominated graphs and their owners. This also identifies their decorated records: the relocated first visit changes from the u_2 strand to u_1 , the second remains on u_3 , and $\text{sgn det}(u_2, u_3) = \text{sgn det}(u_1, u_3) = \sigma$ preserves which visit is over. Both crossing signs are positive. The two smoothing turns are $u_1 \rightarrow u_2$ on BC and $u_2 \rightarrow u_1$ on A . The carrier families are regular through the wall, so their rotations, weights, writhes, polynomials and slots agree. Consequently $T_P(a) = T_E(a)$.

For endpoint c the carriers are $AC \mid B$, and the local residual words are $aAaC$ and $bAbC$. Outside survivors have mask zero or ab . Replacing a by b fixes their graph incidences and owners. The first visit stays on u_1 , while the second changes from u_2 to u_3 ; $\text{sgn det}(u_1, u_2) = \text{sgn det}(u_1, u_3)$ preserves the over/under bit. The local turns are $u_2 \rightarrow u_3$ and $u_3 \rightarrow u_2$. The same regular-family and decorated-record argument therefore proves $T_P(c) = T_E(c)$. This verifies each endpoint separately.

Each of the other two independent pair supports is selector-zero. For such a pair the shared strand u is nonseparating, so $\text{sgn det}(u, v) = \text{sgn det}(u, w)$ for the other directions. Its short arc between local visits is intact after smoothing Q . Smoothing the two local crossings puts both ends on one carrier, with turns $v \rightarrow u$ and $u \rightarrow w$, or with v, w interchanged. Those turn signs are opposite. This carrier is mixed whatever happens outside the disc; hence $T_E(ab) = T_E(bc) = 0$.

Generic selected singleton and pair.

The remaining successor rows are

$$P, b : C \mid AB, \quad E, b : C \mid AB, \quad P, ac : C \mid A \mid B.$$

The AB carrier in E, b has two extra positive crossings, with word $aAcaBc$ (spaces are immaterial); its P, b mate has neither. The corner ledger is

$$\begin{array}{c|cc} & C & \text{other carriers} \\ \hline P, b \text{ and } E, b & \sigma & AB : -\sigma \\ P, ac & \sigma, \sigma & A : -\sigma, B : -\sigma \end{array} \quad (317)$$

Every nonlocal corner partitions according to the displayed successor strings. The actual weight definition gives

$$\text{wt}(C_{ac}) = -\sigma \text{wt}(C_b), \quad \text{wt}(A) \text{wt}(B) = \sigma \text{wt}(AB).$$

For left turns an extra corner changes $(-1)^{\text{corner count}}$; for right turns the weight remains one. If inherited signs make one side mixed, they make the matching product mixed as well. Thus the pair's full selector is the negative of the singleton's, including zero, and the two singleton selectors are equal.

The principal-angle formula gives $\text{rot}(C_{ac}) = \text{rot}(C_b)$ and $\text{rot}(A) + \text{rot}(B) = \text{rot}(AB)$. To verify the first equality, the transitive directions lie in a common open half-plane, and the turn from the first to the last equals the sum of the two turns through the middle, all on the same principal branch. The remaining local turns are the negatives of these, giving the second equality. The P, b and E, b rotations agree by their regular families. If the common singleton selector is zero all three terms vanish. Otherwise A, B are uniform of sign $-\sigma$, so Lemma 1.21 gives $R_A + R_B = R_{AB}$.

Let D_H be the actual E, b diagram on AB and D_L its P, b mate. Switching a in D_H makes a, c an opposite-sign oriented RII pair. Its short boundary arcs, including the selected b connector, lie in the clean event disc. No gap arc or spectator is in its lens. Deleting this bigon gives D_L , with its whole outside record. Put $D_0 = \text{smooth}_a D_H$. The skein equation and the two additional positive crossings give

$$\begin{aligned} P_{D_H} &= a^{-2} P_{D_L} + a^{-1} z P_{D_0}, \\ w_H &= w_L + 2, \quad d_H = d_L - 2, \\ \omega_H - \omega_L &= [a^{d_L-1} z^{-1}] P_{D_0}. \end{aligned} \quad (318)$$

The last line is extraction at $a^{d_H} z^0$; its first summand extracts $a^{d_H+2} = a^{d_L}$.

The two components of D_0 have strings A, B . In the full E word their side vectors relative to (b, a, c) are $A = (1, 1, 0)$, $B = (1, 0, 1)$, $C = (0, 0, 0)$. Outside survivors after b therefore have mask zero or ac . Mask-zero self-crossings have exactly the component records of the P, ac carriers. A mask- ac crossing is mixed in D_0 and absent from those pair-row records. The local crossing c is mixed as well. Thus the intrinsic component polynomials are precisely P_A, P_B , also when a component has no self-crossings; this follows from Lemma 3.20 with the actual records just given. Put $f_i = [z^0] P_i$, $D = d_A + d_B$, and let ℓ be the linking number of this actual auxiliary link. All its mixed crossings are positive, so their number is 2ℓ . Removing the one local mixed crossing, which is absent from D_L , gives

$$\begin{aligned} w_L &= w_A + w_B + 2\ell - 1, \\ D &= 2 - (w_A + w_B) - (R_A + R_B) \\ &= 2 - (w_L - 2\ell + 1) - R_{AB} \\ &= d_L + 2\ell. \end{aligned} \quad (319)$$

Lemma 4.5 gives the actual mixed-link row, without requiring that link to be split. Equations (318) and (319) therefore yield

$$\begin{aligned}\omega_H - \omega_L &= [a^{d_L-2+2\ell}]f_A f_B - [a^{d_L+2\ell}]f_A f_B \\ &= [a^{D-2}]f_A f_B - [a^D]f_A f_B = -\omega_A \omega_B.\end{aligned}\quad (320)$$

For the last equality, if either row is zero all these coefficients are zero. Otherwise the actual uniform carriers are subpolygons of a decomposition of the generic side (Lemma 3.5(i)), the domain of Theorem 3.64, so their exponents are bounded below by d_A, d_B . No product exponent is below D , and its coefficient at D is the product of the two coefficients at those bounds. Crossing-free carriers are included in that theorem and have polynomial one. The common C carrier has the same intrinsic record, writhe and rotation in all three rows; its coefficient read is common. Multiplying (320) by it, the exterior factor and the singleton selector gives

$$T_E(b) = T_P(b) + T_P(ac). \quad (321)$$

Together with the common and zero rows, this proves equality of the two full-availability sums in the generic orbit.

Extreme orbit: words and owners.

Relabel the crossings $x = (u_1, u_2)$, $y = (u_1, u_3)$, $z = (u_2, u_3)$. Here z as a crossing label is distinguished by context from the HOMFLY variable. The complete and empty local graphs have words

$$H = xyAzxByzC, \quad L = yxAxzBzyC. \quad (322)$$

The sign classification gives $s_x = s_z = -s_y = \sigma$. On L every local subset is independent; on H only the empty set and the three singletons are. The gap side vectors relative to (x, y, z) are

$$A = (1, 1, 0), \quad B = (0, 1, 1), \quad C = (0, 0, 0). \quad (323)$$

Thus crossings joining AB, AC, BC have masks xz, xy, yz , respectively. Mask-zero crossings have both visits in one gap. The legitimate full support $Q \cup \{x, y, z\}$ on L has the three outer carriers with strings A, B, C and the central triangle. Write $P_G, f_G = [z^0]P_G, w_G$ for the actual outer data, $G = A, B, C$. These symbols always refer to those actual carriers. A crossing-free one has $P_G = f_G = 1$ and $w_G = 0$.

The local smoothing turns on the outer carriers are $u_2 \rightarrow u_1, u_3 \rightarrow u_2, u_1 \rightarrow u_3$ on A, B, C . Each has sign $s_o = -\sigma$. The opposite three turns make the central triangle uniform of sign σ . It is the small triangle inside the event disc, has three corners and no crossings, and therefore has $R = 1$, slot zero and coefficient one.

Extreme pair rows vanish.

In a local pair row on L , the remaining local crossing r is undominated. Any outside neighbour of r has a two-letter mask and is consequently adjacent to one of the selected pair, so is dominated. Thus r is an isolated self-crossing on its owner. If that owner is mixed its weight is zero; if it is uniform Lemma 3.68 makes its corner coefficient zero. Each pair term on L therefore vanishes. The corresponding pairs on H are absent.

Extreme singleton transport.

Here are the three successor calculations separately. The H words retain the visits of dominated local crossings only to identify the arcs; deleting them gives the common and affected strings.

j	common	affected	H strings	L affected word	q	r
x	A	BC	$yAz \mid ByzC$	$zBzyCy$	y	z
y	C	AB	$Azx B \mid zCx$	$xAxzBz$	x	z
z	B	CA	$xB y \mid CxyA$	$yCyxAx$	y	x

(324)

The affected carrier has one selected-local corner of sign σ ; the common carrier has its mate of sign $-\sigma$. This follows row by row from $u_1 \rightarrow u_2$, $u_3 \rightarrow u_1$, $u_2 \rightarrow u_3$ on the affected carriers. Their corner lists and weights agree across the wall. The two carriers in each row form regular families through it: the unselected local crossings create no corner and the only selected local turn has nonzero determinant. Their rotations agree.

After deleting the displayed local visits, the common records are identical. An outside survivor after j has mask zero or the complementary pair mask, by the parity argument. The latter has one visit in each affected gap; all these crossings are retained in the clean affected diagram on H and in its L mate. Call these actual diagrams D_0, D_+ , and their common absolute rotation R . The only additional self-crossings of D_+ are q, r in the table. Switching q makes an opposite-sign oriented RII pair. The two short bigon edges consist of adjacent local strand subarcs and the already selected j connector; they lie in the clean event disc, without any gap arc or outside strand in its lens. Deleting the pair gives exactly D_0 with its outside record. Thus, putting $d_0 = 1 - w_0 - R$ and $D_* = \text{smooth}_q D_+$, the skein equation gives

$$\begin{aligned} w_+ &= w_0 + 2, & d_+ &= d_0 - 2, \\ P_{D_+} &= a^{-2}P_{D_0} + a^{-1}zP_{D_*}, \\ \omega_+ - \omega_0 &= [a^{d_0-1}z^{-1}]P_{D_*}. \end{aligned} \tag{325}$$

The two components of D_* , for the three rows of (324), are respectively

$$C \mid zBz, \quad A \mid zBz, \quad C \mid xAx. \tag{326}$$

Every complementary-mask retained crossing is mixed; every mask-zero crossing lies in the indicated gap with exactly its full-support outer record. The extra self-crossing r on the second component is isolated: any outside neighbour also meets the selected q and is dominated in the legitimate support $Q \cup \{j, q\}$. Lemma 3.67 and the single-crossing polynomial one therefore give component rows (f_C, f_B) , (f_A, f_B) , (f_C, f_A) in the three cases. Their component writhes are $(w_C, w_B + 1)$, $(w_A, w_B + 1)$, $(w_C, w_A + 1)$: the singleton polynomial is one, but its writhe has not been discarded.

If the matched singleton selector is zero, both complete terms are zero. Otherwise the affected parent is uniform of sign σ . The two additional smoothings q, r split it into the two clean outer carriers in (326) and the central triangle. For $j = x$, $q = y$ puts the $-\sigma$ corner on C and $r = z$ puts one on B . For $j = y$, $q = x$ puts one on A and $r = z$ puts one on B . For $j = z$, $q = y$ puts one on C and $r = x$ puts one on A . Each clean outer carrier has exactly one dissent and at least two inherited corners, since Lemma 3.5 gives at least three corners on every carrier. Each successive smoothing introduces two opposite principal turns and leaves the other turns unchanged. Signed rotation is additive by Lemma 1.20. The one-dissent rotation result, with reversal if necessary, puts both outer rotations on the parent's sign ray; the central triangle has that same sign. Consequently, for the two clean outer absolute rotations,

$$R_1 + R_2 + 1 = R. \tag{327}$$

Let ℓ be the linking number in D_* and let w_1, w_2 be the *clean full-support* outer writhes, without the extra r . The clean diagram D_0 contains the two outer self-crossing sets and exactly 2ℓ positive crossings that become mixed in D_* . Thus

$$\begin{aligned} w_0 &= w_1 + w_2 + 2\ell, \\ D := d_1 + d_2 &= 2 - (w_1 + w_2) - (R_1 + R_2) \\ &= 2 - (w_0 - 2\ell) - (R - 1) \\ &= d_0 + 2 + 2\ell. \end{aligned} \tag{328}$$

Lemma 4.5 gives $[z^{-1}]P_{D_*} = a^{-2\ell}(a - a^{-1})f_1f_2$ with the intrinsic component polynomials just identified. The floor theorem (Theorem 3.64) applies to the actual one-dissent outer carriers, subpolygons of a decomposition of the generic side (Lemma 3.5(i)). If a row is zero the following support assertion is vacuous. Otherwise its exponents are at least d_i , so the product has none below D . Extracting the two terms in (325) gives

$$\begin{aligned}\omega_+ - \omega_0 &= [a^{d_0-2+2\ell}]f_1f_2 - [a^{d_0+2\ell}]f_1f_2 \\ &= [a^{D-4}]f_1f_2 - [a^{D-2}]f_1f_2 = 0.\end{aligned}\tag{329}$$

Common-carrier and exterior reads agree, as do every weight and rotation. Multiplication now proves each of $T_H(x) = T_L(x)$, $T_H(y) = T_L(y)$, $T_H(z) = T_L(z)$.

Extreme empty/full couple: two switches.

The empty local support has one contact carrier on each side, with diagrams D_H, D_L , common positive writhe w , absolute rotation R , weight W , and slot $d = 1 - w - R$. The common writhe counts the same retained labels. The unsmoothed local crossings are not corners, so the regular-family argument supplies the common rotation and weight. Its row difference is $C_QW(\omega_H - \omega_L)$.

The original local over-order is cyclic. When $\sigma = +1$ it is u_1 over u_2 , u_3 over u_1 , u_2 over u_3 ; switching x makes the total order $u_2 > u_3 > u_1$. When $\sigma = -1$ all arrows reverse and the switched order is $u_1 > u_3 > u_2$. The two switched diagrams are therefore related by ordinary RIII with the same outside record. Their HOMFLY values agree. Subtract the two skein equations at x ; their switched terms cancel and give

$$P_{D_H} - P_{D_L} = a^{-1}z(P_{D_H^x} - P_{D_L^x}),\tag{330}$$

where the superscript x denotes oriented smoothing, not switching. After that smoothing the successor rows are

$$H : yAz \mid ByzC, \quad L : A \mid zBzyCy.$$

Switching y in each leaves y negative and z positive. Their short local arcs have the orientation-compatible pairing inherited from the x connector, and bound an empty RII bigon inside the event disc. Deleting it on both sides leaves the same boundary pairing $A \mid BC$ and the same outside record. The two switched link polynomials agree, by RII invariance and Lemma 3.20. Subtracting these two skeins at y and substituting into (330) yields

$$P_{D_H} - P_{D_L} = a^{-2}z^2(P_{D_H^{xy}} - P_{D_L^{xy}}).\tag{331}$$

Only matched switched diagrams have been identified: neither is identified with the opposite *unswitched* diagram.

Smoothing x on either side splits the contact knot. On H , y is then mixed, so its smoothing rejoins the two components; D_H^{xy} is a knot. On L , y is a self-crossing and its smoothing splits again; $J = D_L^{xy}$ has three components

$$A, \quad C, \quad zBz.$$

Knot parity in Theorem 3.18 gives $[z^{-2}]P_{D_H^{xy}} = 0$. In J , the mask-zero crossings give the full-state outer records, while every two-letter-mask crossing is mixed. The remaining z is isolated on its component, since an outside neighbour also meets selected x or y . Its polynomial factor is one and its writhe is one. The intrinsic component rows are therefore f_A, f_C, f_B , by Lemma 3.67. Write $K = f_A f_B f_C$ and Λ for the sum of the three pairwise linking numbers of J .

For clarity the lowest row needed here follows from the skein relations themselves. A two-component link has no z^{-3} row: switch mixed crossings until it is split. Each smoothing is a knot,

whose z^{-4} row is zero; the split endpoint is $(a - a^{-1})z^{-1}$ times two knot polynomials and also has no z^{-3} row. The skein relation propagates that zero back. For three components, switching a positive mixed crossing to a negative one changes the total linking number by minus one. Its smoothing has two components and contributes no z^{-3} row. Thus at row z^{-2} the positive value is a^{-2} times the negative value; the reverse switch gives the inverse factor. Switch until the components have a total over-order, and translate their successive height levels apart to obtain the split endpoint. The intrinsic component diagrams are unchanged. The split polynomial is $((a - a^{-1})/z)^2 P_A P_B P_C$, by the split product proof of Lemma 4.5, applied successively (or Theorem 3.24 to its three clean one-component blocks). The net linking change gives

$$[z^{-2}]P_J = a^{-2\Lambda}(a - a^{-1})^2 K. \quad (332)$$

Extracting $a^d z^0$ in (331) now gives

$$\omega_H - \omega_L = -[a^{d+2+2\Lambda}](a - a^{-1})^2 K. \quad (333)$$

Extreme empty/full couple: selectors and floor.

All inherited corners of the empty contact carrier partition among the three outer carriers. Each outer has its one local corner of sign $s_o = -\sigma$, while the central triangle has three of the opposite sign. Let W_{full} be the product of these four carrier weights. If the parent is mixed, $W = 0$. A nonzero full selector would force all outer carriers to be uniform with sign s_o , hence all inherited parent corners to have that sign, a contradiction. Both sides of the proposed empty/full identity are then zero.

Assume the parent uniform and put $\chi = 1$ if its sign is s_o , $\chi = 0$ if its sign is $-s_o$. If $\chi = 1$, all outer carriers are uniform and $W_{\text{full}} = -W$: when $s_o = -1$ the three outers are all right and the left central triangle contributes $(-1)^3 = -1$; when $s_o = +1$ the outers contain three more left corners than the parent and the right central weight is one. If $\chi = 0$, each outer has exactly one dissent and at least two inherited corners, so it is mixed and $W_{\text{full}} = 0$.

Every smoothing creates two opposite principal turns. Summing all turns across the three smoothings therefore gives

$$\text{rot}(A) + \text{rot}(B) + \text{rot}(C) + \text{rot}(\text{central}) = \text{rot}(\text{parent}).$$

In the live branch $\chi = 1$ the outer rotations have the parent's sign and the central triangle has the opposite sign and absolute rotation one. In the dead branch $\chi = 0$ the one-dissent rotation lemma puts the outer rotations on the parent's sign ray, as does the central triangle. Hence

$$R_A + R_B + R_C = \begin{cases} R + 1, & \chi = 1, \\ R - 1, & \chi = 0. \end{cases} \quad (334)$$

The empty parent retains the three positive local crossings, the outer self-crossings, and all positive mixed crossings of J . There are 2Λ of the latter. Consequently

$$\begin{aligned} w &= 3 + w_A + w_B + w_C + 2\Lambda, \\ D := d_A + d_B + d_C &= 3 - (w_A + w_B + w_C) - (R_A + R_B + R_C) \\ &= 6 - w + 2\Lambda - (R_A + R_B + R_C) \\ &= \begin{cases} d + 4 + 2\Lambda, & \chi = 1, \\ d + 6 + 2\Lambda, & \chi = 0. \end{cases} \end{aligned} \quad (335)$$

Each actual outer carrier is uniform or exactly one-dissent. Each is a subpolygon of a decomposition of the generic side (Lemma 3.5(i)), the domain of Theorem 3.64; with reversal for negative inherited signs, that theorem bounds the support of f_G below by d_G . If some $f_G = 0$ then $K = 0$ and all the following coefficients vanish, without any degree assigned to the zero polynomial. Otherwise no exponent of K is below D and $[a^D]K = \omega_A \omega_B \omega_C$.

Expand $(a - a^{-1})^2 = a^2 - 2 + a^{-2}$ in (333). In the live branch its three sampled degrees of K are $D - 4, D - 2, D$; only the last can contribute. In the dead branch they are $D - 6, D - 4, D - 2$, all below the floor. Thus

$$\omega_H - \omega_L = \begin{cases} -\omega_A \omega_B \omega_C, & \chi = 1, \\ 0, & \chi = 0. \end{cases} \quad (336)$$

The central coefficient is one. In the live branch multiplication by $C_Q W$ and $W_{\text{full}} = -W$ gives

$$\begin{aligned} T_H(\emptyset) - T_L(\emptyset) &= -C_Q W \omega_A \omega_B \omega_C \\ &= C_Q W_{\text{full}} \omega_A \omega_B \omega_C = T_L(xyz). \end{aligned} \quad (337)$$

The mixed-parent and uniform-dead branches were both zero above. Together with the three singleton transports and the three pair zeros, this proves equality of the full-availability sums in the extreme orbit and completes (iii).

This proof uses neither the triple-wall theorem nor the aggregate vertex-edge conclusion. Each skein triple, component, owner, rotation and coefficient slot has been specified for this triple wall. No selector, exterior factor or coefficient has been cancelled. The availability-one and availability-zero fibres of a general wall are not proved here; no general triple-wall conclusion follows. $\square$

Proofs for Section 5

Proof of Lemma 5.2

Proof of 5.2. First, every common positive divisor of r and N divides $N - 2r = 1$, so $\gcd(r, N) = 1$. For integer indices j, k , if $rj \equiv rk \pmod{N}$ then N divides $r(j - k)$. Multiplying that divisibility by 2 and using $2r = N - 1$ shows that N divides $j - k$. Thus multiplication by r permutes the residue classes modulo N , so the displayed star vertices are distinct. Moreover R^k is the identity precisely when kr/N is an integer. The same divisibility argument makes this equivalent to $N \mid k$; the least positive such k is N , proving that R has order N .

(i) $(\sigma K_r)_t = \mu_{t+1} = u_{rt} = R(u_{r(t-1)}) = R(\mu_t)$. Hence $\tilde{\lambda}_t = R(\tilde{\lambda}_{t-1})$ for every t , so each principal turn is the angle of R reduced to $(-\pi, \pi)$, which is $2\pi r/N$ itself because $2r < N$; the N turns sum to $2\pi r$. (ii) E_t is the chord from u_a to u_{a+r} with $a = r(t-1)$. The counterclockwise arc from u_a to u_{a+r} has angle $2\pi r/N < \pi$, so the midpoint of the chord is at distance $\cos(\pi r/N)$ from the origin and the chord line is tangent there to the circle of that radius; a chord directed along a counterclockwise arc of angle less than π has the centre on its left. (iii) (G1): the residue permutation proved above gives N distinct vertices on the unit circle, and no three points of a circle are collinear. (G2): the N edge lines are pairwise distinct tangent lines of the circle of radius ρ (their tangency points are the N distinct images of one midpoint under the powers of R , whose order was proved to be N). Through a point outside a circle pass exactly two tangent lines of the circle, so no three of the edge lines are concurrent, and a fortiori no three edge interiors are. The statements for K_{-r} are Lemma 1.24(ii)–(iv). (iv) The four determinants are $\det((2, 2), (0, 2)) = 4$, $\det((2, 2), (2, 0)) = -4$, $\det((0, 2), (2, 0)) = -4$, $\det((-2, 0), (0, -2)) = 4$; the turns are $\tau_1 = \chi_{412} = \chi_{124}$, $\tau_2 = \chi_{123}$, $\tau_3 = \chi_{234}$, $\tau_4 = \chi_{341} = \chi_{134}$ (Lemma 1.4(i)). The edges $E_1 = [(0, 0), (2, 2)]$ and $E_3 = [(0, 2), (2, 0)]$ meet at $(1, 1)$, while E_2 and E_4 are disjoint parallel segments; there is one crossing point, so (G2) holds. The principal turns are $-\frac{3\pi}{4}, \frac{3\pi}{4}, \frac{3\pi}{4}, -\frac{3\pi}{4}$, which sum to 0. $\square$

Proof of Lemma 5.3

Proof of 5.3. For finitely many planar unit vectors, lying in no closed semicircle is equivalent to the origin being in the interior of their convex hull. Indeed a supporting line through the origin for a boundary point of that convex polygon, or a separating line if the origin is outside it, puts all vectors in one closed half-plane through the origin. Conversely, a closed half-plane through the origin cannot contain a neighbourhood of the origin. This is also valid for a segment or a point as convex hull, since either is contained in a closed half-plane through the origin whenever the origin is not an interior point in the plane.

Fix a parameter and put $b = \sum_i u_i$. Interiority gives an $\epsilon > 0$ such that $-\epsilon b$ belongs to the convex hull. Write $-\epsilon b = \sum_i a_i u_i$ with $a_i \geq 0$. Then

$$\sum_i (a_i + \epsilon) u_i = -\epsilon b + \epsilon \sum_i u_i = 0, \quad a_i + \epsilon > 0. \quad (338)$$

Thus there is a strictly positive closing vector at each parameter.

Fix such a vector l at t_0 . Choose p, q with $\det(u_p(t_0), u_q(t_0)) \neq 0$, possible because the vectors are not all parallel. Let $E(t) = \sum_i l_i u_i(t)$ and $D(t) = \det(u_p(t), u_q(t))$. Correct just two coordinates by

$$\delta_p(t) = -\frac{\det(E(t), u_q(t))}{D(t)}, \quad \delta_q(t) = -\frac{\det(u_p(t), E(t))}{D(t)}. \quad (339)$$

The determinant expansion of a vector in the basis (u_p, u_q) gives $E + \delta_p u_p + \delta_q u_q = 0$. The correction is continuous and vanishes at t_0 , so all corrected lengths remain positive on a neighbourhood of t_0 . This constructs local continuous positive closing vectors.

For a one-point parameter interval the pointwise choice already suffices. Otherwise choose a finite cover by proper smaller relative intervals, with nonempty complements, whose closures lie in these neighbourhoods. Continuous nonnegative weights subordinate to the smaller intervals, summing to one, can be obtained by taking distances to their complements and normalizing their sum. Multiply each weight by the local closing vector on its neighbourhood and sum. The products extend continuously by zero outside the smaller intervals because the local vectors are bounded on their closures. The result is a continuous positive closing vector: closure is linear and every positively weighted local vector has all entries positive. Finally, for fixed directions the set of strictly positive closing lengths is convex. Straight segments in that set join any prescribed endpoint lengths to the selected ones. $\square$

Proof of Theorem 5.5

Proof of 5.5. Use tail directions $d_i = \mu_{i+1} - \mu_i$, unit directions $u_i = d_i/|d_i|$, and principal turns $\vartheta_i \in (-\pi, \pi)$ from u_{i-1} to u_i . Multiplication of $u_i = e^{i\vartheta_i} u_{i-1}$ around the cycle shows that $\sum_i \vartheta_i$ is an integer multiple of 2π . Principal turns vary continuously on the regular locus, so their normalized sum is constant on every path there. The symbols τ_i remain the signs of these real principal turns, not the angles themselves.

Reconstructing a polygon from directions. A direction path with no antiparallel successive pair and with no closed semicircle containment can be closed using Lemma 5.3. Starting at any continuous base point, successively add the positive vectors $l_i(t)u_i(t)$; the final sum is zero, so this gives a continuous closed labelled n -tuple. Every edge is nonzero, and every successive pair has the prescribed principal turn in $(-\pi, \pi)$. Consequently the tuple stays in $\mathcal{R}_n$. Interpolate the base point between those of P, P' and use the endpoint length segments of that lemma to recover the exact endpoint tuples. No vertex is inserted or deleted.

Nonzero rotation. Choose real arguments β, β' for the first unit directions and interpolate these arguments and all principal turns linearly. The interpolated turns lie strictly between $-\pi$ and π and have sum $2\pi r$. The lifted edge arguments, starting with $\theta_0(t) = \beta(t)$, are

$$\theta_k(t) = \beta(t) + \sum_{m=1}^k \vartheta_{m+1}(t), \quad 0 \leq k \leq n, \quad \theta_n(t) - \theta_0(t) = 2\pi r, \quad (340)$$

where turn indices are cyclic. The first n entries give the unit directions; the last repeats the first direction because r is an integer. If these directions lay in a closed semicircle, Lemma 5.4 would put the whole lifted sequence in one interval of length π , implying $|2\pi r| \leq \pi$. This contradicts $r \neq 0$. The reconstruction above therefore joins the exact labelled endpoints, with no cyclic shift.

The zero-rotation angle domain. For $r = 0$, the lifted angles close as real numbers and form a cyclic vector $\theta = (\theta_1, \dots, \theta_n)$ satisfying

$$|\theta_{i+1} - \theta_i| < \pi \quad \text{cyclically.} \quad (341)$$

The directions of any regular closed polygon lie in no closed semicircle. For if a normal $v \neq 0$ had $\langle v, d_i \rangle \geq 0$ for every edge, closure would force every scalar product to be zero. All directions would then be on the two rays of one line. Antiparallel adjacent directions are forbidden, so every edge would point along the same ray. Their positive length sum could not be zero. By Lemma 5.4, the cyclic angle vector of a zero-rotation polygon consequently has $\max_i \theta_i - \min_i \theta_i > \pi$.

For an ordered pair of indices with forward cyclic distance $d \in \{2, \dots, n-2\}$, define

$$\mathcal{A}(h, d) = \{\theta : (341) \text{ holds, } \theta_h - \theta_{h+d} > \pi\}. \quad (342)$$

These open convex charts cover all zero-rotation endpoints: take a maximum and a minimum; they cannot be cyclically adjacent by (341). Conversely every chart point has directions lying in no closed semicircle, by Lemma 5.4. Every continuous angle path inside the union can thus be reconstructed as above.

Connected chart overlaps for $n \geq 5$. For $2 \leq d \leq n-3$, both of the following chart pairs intersect:

$$\mathcal{A}(h, d), \mathcal{A}(h, d+1); \quad \mathcal{A}(h, d+1), \mathcal{A}(h+1, d). \quad (343)$$

Here is an explicit witness, reading indices $s = 0, \dots, n-1$ from h . Set the high value to $3\pi/2$ and the low value to zero. For the first pair, put the high value at $s = 0$, descend linearly to zero at $s = d$, keep zero through $s = d+1$, and rise linearly back to the high value at $s = n$. For the second pair, keep the high value at $s = 0, 1$, descend linearly to zero at $s = d+1$, and rise linearly back to the high value at $s = n$. Every ramp has at least two steps, so each step has magnitude at most $3\pi/4 < \pi$. Each specified high-minus-low difference is $3\pi/2 > \pi$, proving both overlaps. The cyclic last step is one of these ramp steps.

The first moves reduce any distance to two. At distance two, the two moves $\mathcal{A}(h, 2) \rightarrow \mathcal{A}(h, 3) \rightarrow \mathcal{A}(h+1, 2)$ advance the high index by one, and are valid even for $n = 5$. Hence their overlap graph is connected. Join endpoint angle vectors to consecutive overlap witnesses by straight segments in their convex charts. The resulting finite path of feasible angle vectors reconstructs a regular path from P to P' , again with no relabelling.

The exceptional four-vertex fibre. For $n = 4$ the charts are just $\mathcal{A}(h, 2)$ for four indices h . They are nonempty: values $3\pi/2, 3\pi/4, 0, 3\pi/4$, beginning at h , satisfy all their inequalities. If $\theta_h - \theta_{h+2} > \pi$, the two intermediate entries satisfy

$$\theta_h > \theta_{h+1} > \theta_{h+2}, \quad \theta_h > \theta_{h+3} > \theta_{h+2}. \quad (344)$$

For example $\theta_{h+1} > \theta_h - \pi > \theta_{h+2}$, and $\theta_{h+1} < \theta_{h+2} + \pi < \theta_h$; the other case is identical. Thus the turn signs, beginning at vertex h , are $(+, -, -, +)$. The four charts force the four distinct cyclic rotations of this word. They are therefore disjoint, and every rotation-zero four-tuple belongs to one of them. In particular no such tuple has a zero turn. Along a regular zero-rotation path all four nonzero turn signs remain fixed by continuity, so different words cannot be joined. Two tuples with the same word belong to the same convex angle chart; their real argument choices can differ by a common multiple of 2π , which changes neither the chart inequalities nor the reconstruction. Interpolation inside that chart and the positive-length construction join them. There are exactly four components.

For $K_0 = ((0, 0), (2, 2), (0, 2), (2, 0))$ the corner determinants are $(-4, 4, 4, -4)$. Its principal turns are the corresponding signs times $3\pi/4$, so its rotation is zero. Cyclic shifts rotate this sign word and exhibit all four components. Shifting P' to match the word of P gives the asserted path in the exceptional case.

Finally a regular triangle is noncollinear: three collinear nonzero closing edges must reverse direction at a vertex. All three turn determinants of a noncollinear triangle equal its oriented area determinant, so the three principal turns have one sign. Their sum is a nonzero multiple of 2π of absolute value less than 3π ; its rotation is $+1$ or -1 . There is no zero-rotation case at $n = 3$. These cases exhaust admissible arities. Passing any constructed path to the cyclic quotient identifies its shifted target with the original target orbit, proving the orbit statement. $\square$

Proof of Lemma 5.7

Proof of 5.7. Put $u = \tilde{\lambda}_{j-1}$, $v = \tilde{\lambda}_j$, and $\theta = \text{Arg}(v/u) \in (-\pi, \pi)$, using the complex identification of the oriented plane. Let $\alpha = \text{Arg}(q/u)$ and $\beta_0 = \text{Arg}(v/q)$ be principal turns. All three exist because admissibility makes the relevant pairs nonparallel. At the new point the actual turn is $\beta_\varepsilon = \text{Arg}((v - \varepsilon q)/q)$, which converges to β_0 . The turn at the successor vertex also changes, but converges to its old value because its incoming direction $v - \varepsilon q$ tends to v and the old corner is regular. All other original turns except the replaced turn at j are unchanged. By Lemmas 2.23 and 1.20(ii), the rotation of P_ε is one fixed integer on a sufficiently small positive interval. Taking the limit of the full sum of turns therefore gives the exact integer identity

$$2\pi(\text{rot}(P_\varepsilon) - \text{rot}(P)) = \alpha + \beta_0 - \theta. \quad (345)$$

The left side is constant on that interval, not an assertion that the successor turn is unchanged at finite ε .

The three direction ratios multiply as $(q/u)(v/q) = v/u$, so $\alpha + \beta_0 - \theta$ is a multiple of 2π . Moreover $\text{sgn } \alpha = -\chi_-$ and $\text{sgn } \beta_0 = -\chi_+$. In the same-sign sector both have the sign $\tau_j = \text{sgn } \theta$. Their sum and θ lie in the same open interval of length 2π , so the multiple is zero. In the mixed sector α, β_0 have opposite signs, making their sum lie in $(-\pi, \pi)$; the difference from θ is again strictly between -2π and 2π , so is zero. In the loop sector both signs are $-\tau_j$. When $\tau_j = 1$, the difference lies in $(-3\pi, 0)$ and must be -2π ; when $\tau_j = -1$ it lies in $(0, 3\pi)$ and must be 2π . Equation (345) gives exactly the stated rotations.

For the cone description, set the angle of u equal to zero and that of v equal to θ . The determinant signs change only at the four rays $u, v, -u, -v$. The same-sign inequalities put q strictly inside the turn wedge from u to v of angular width $|\theta| < \pi$; the loop inequalities put it inside the opposite wedge. The remaining two open angular intervals are exactly the mixed sector. This proves the cone statements as well. $\square$

Proof of Proposition 5.9

Proof of 5.9. Exhaustion. If $(n-1, r)$ is admissible, case (Z) applies. Otherwise $n-1 \geq 3$, so Definition 1.25 says that either $2|r| \geq n-1$ or $(n-1, r) = (3, 0)$. In the first case, $2|r| < n$ and integrality give $2|r| = n-1$. Thus $n = 2|r| + 1$, and $n \geq 4$ forces $|r| \geq 2$, giving (L). The exceptional predecessor gives exactly $(n, r) = (4, 0)$, case (L₀). These alternatives are disjoint: in case (L) the predecessor $(n-1, r) = (2|r|, r)$ has $2|r| = n-1$ and in case (L₀) it is $(3, 0)$, so in neither is it admissible (Definition 1.25), which excludes (Z); and (L) has $|r| \geq 2$ while (L₀) has $r = 0$. In particular $(4, 1)$ and $(4, -1)$ have admissible triangular predecessors and belong to (Z), not (L).

Sectors. By Lemma 5.7 each sector is a nonempty open cone of directions; admissibility removes finitely many lines ($\det(q, \mu_k - \mu_j) = 0$, $k \neq j$), so each sector contains admissible vectors. For such q and $0 < \varepsilon < \varepsilon_0$, as Definition 5.8 requires, Lemma 2.23 makes P_ε generic and Lemma 5.7 gives its rotation.

(Z) A generic $(n-1)$ -gon of rotation r exists by Lemma 1.26; in the mixed sector the rotation is preserved.

(L) Let $r \geq 2$. The star K_{r-1} is generic with all turns left and rotation $r-1$ (Lemma 5.2; for $r=2$ it is the counterclockwise triangle). Subdivide one edge at its midpoint, obtaining a $2r$ -gon in $\mathcal{R}_{2r}$ with one zero turn and rotation $r-1$ (Lemma 1.20(iii)). Write the subdivided edge as $B-A=w$ and its moved midpoint as $M=(A+B)/2+\delta v$, where $\delta > 0$ and $\det(w, v) > 0$: the displacement is to the left of the directed edge. The new corner determinant is

$$\det(M-A, B-M) = -\delta \det(w, v) < 0. \quad (346)$$

Thus its turn is right. The two neighbouring turn determinants are strictly positive at $\delta = 0$ and remain positive for a short move. All edge lengths remain positive; the initially positive-flat new corner and all other regular corners stay in the open regular locus. Rotation therefore remains $r - 1$ by Lemma 1.20(ii). The generic locus is open (Proposition 1.11) and dense (the density clause of Lemma 1.26), so a small perturbation gives a generic $2r$ -gon P of rotation $r - 1$ with a vertex j of turn $\tau_j = -1 = -\text{sgn}(r)$. A loop-sector insertion at j gives a generic $(2r + 1)$ -gon of rotation $r - 1 - \tau_j = r$. For $r \leq -2$ apply the reflection $(x, y) \mapsto (x, -y)$ to the construction for $-r$: it negates every chirotope and every rotation number, maps generic polygons to generic polygons, and maps a loop-sector insertion to a loop-sector insertion (the sectors are defined by sign patterns relative to τ_j , which all flip together).

(L₀) Insert a loop at any vertex of the counterclockwise triangle: the rotation becomes $1 - 1 = 0$, and the result is generic with four vertices. $\square$

Proofs for Section 6

Proof of Theorem 6.2

Proof of 6.2. Strong induction on n . The case $n = 3$ follows from Lemma 6.1(i).

Let $n \geq 4$, let P be generic of rotation r , and let $g \neq h$ be two edge labels. For a generic labelled tuple Q with n vertices put $\Delta_{gh}(Q) = A_g(Q) - A_h(Q)$. Choose a target Z in the fibre (n, r) as follows.

- If (n, r) is not minimal: Z is a zero anchor whose insertion vertex j is chosen with $j \notin \{g, h\}$ (Proposition 5.9; there are $n-1 \geq 3$ choices of j). Its soft edge is $E_j(Z)$. Choose its parameter below the geometric bound of Definition 5.8 and below the two bounds in Theorem 2.24 for the roots g and h . All three bounds are positive, so such a choice exists.
- If (n, r) is minimal with $|r| \geq 2$: $Z = K_r$.
- If $(n, r) = (4, 0)$: $Z = K_0$.

Lemma 5.6 gives k ($k = 0$ in the first two cases) and a path from P to $\sigma^k Z$ with finitely many simple walls. Along the path Δ_{gh} is constant: on generic segments by Proposition 2.7; at (T), (E), (C) by Theorem 2.15, which gives zero jump for each root separately; at a flat wall at i the jump of Δ_{gh} is $A_{D_i(g)}(Q) - A_{D_i(h)}(Q)$ with Q the generic deletion (Theorem 2.12), which vanishes by induction; at a vertex-edge wall the jump is $s(A_{h_1}(\lambda_1)A_{h_2}(\lambda_2) - A_{h'_1}(\lambda_1)A_{h'_2}(\lambda_2))$ with $(h_1, h_2) = H(g)$, $(h'_1, h'_2) = H(h)$ and generic halves (Theorem 2.14), which vanishes by induction applied to each half. Hence $\Delta_{gh}(P) = \Delta_{gh}(\sigma^k Z)$.

At the target: in the non-minimal case $k = 0$ and neither E_g nor E_h is the soft edge of Z , so Theorem 2.24 in the mixed sector gives $A_g(Z) = A_h(Z) = 0$. For $Z = K_r$ with $r \geq 2$: by Lemma 5.2(i) and Proposition 2.26(i), $A_{g+1}(K_r) = A_g(\sigma K_r) = A_g(R(K_r)) = A_g(K_r)$, the last step because every chirotope, hence every gate, is invariant under a rotation of the plane; so all roots of K_r have one value. For $r \leq -2$, $K_r = \overline{K_{|r|}}$ by Definition 5.1, and Proposition 2.26(ii) gives $A_h(K_r) = (-1)^N A_{1-h}(K_{|r|}) = -A_{1-h}(K_{|r|})$ with $N = 2|r| + 1$ odd; since $h \mapsto 1 - h$ is a bijection of $\mathbb{Z}/N$, all roots of K_r have one value as well. Cyclic covariance, Proposition 2.26(i), covers every shift. For $Z = K_0$, Lemma 6.1(ii) gives -1 at every root of every shift. In every target case $\Delta_{gh}(\sigma^k Z) = 0$. $\square$

Proof of Corollary 6.12

Proof of 6.12. Theorem 6.11 identifies C with A on every generic polygon, and Definition 3.28 defines C on generic polygons only; so before $C = A$ is substituted in an identity of Corollary 6.3, every argument of that identity must be shown generic. The two sides of a wall germ are generic by Definition 1.27. At a simple flat wall the deletion $P(0) \setminus j$ is generic by Lemma 1.34(i), and at a simple vertex-edge wall both halves are generic by Lemma 1.34(ii). The soft family P_ε is generic for all sufficiently small $\varepsilon > 0$ in every sector by Lemma 2.23. Reversal permutes the vertex triples and preserves the edge segments, so it preserves (G1) and (G2). The two triangles $K_{\pm 1}$ are generic by Lemma 5.2(iii).

For the cusp law let the simple cusp wall be at j and write $A = \mu_{j-1}(0)$, $M = \mu_j(0)$ and $B = \mu_{j+1}(0)$. The deletion $Q = P(0) \setminus j$ has the vertices $\mu_i(0)$, $i \neq j$, and its edges are the parent edges $E_i(0)$, $i \notin \{j-1, j\}$, together with the fused edge $[A, B]$; it satisfies (G1) by the hypothesis of the cusp law. For (G2): by Definition 1.29(K) the three points are collinear with M outside the closed segment $[A, B]$, so $[A, B]$ is contained in the longer of $[A, M] = E_{j-1}(0)$ and $[M, B] = E_j(0)$,

and its relative interior in the relative interior of that edge. Suppose three distinct edges of Q had a common point in their relative interiors. Two adjacent edges of Q are not collinear, by (G1) for Q , and so meet only at their common vertex; hence the three edges are pairwise remote in Q . Send the fused edge, if it is among them, to the longer parent edge containing it, and every other edge to itself. This map is injective, since neither $E_{j-1}(0)$ nor $E_j(0)$ is an edge of Q , and the common point lies in the relative interiors of the three image edges. The image edges are pairwise remote in $P(0)$: two surviving edges are adjacent in $P(0)$ exactly when they share a vertex, which is a vertex of Q , hence exactly when they are adjacent in Q ; and the parent neighbours of the longer cusp edge are the other cusp edge, which is not an edge of Q , and one of the two edges of Q adjacent to the fused edge, both excluded because the three edges are pairwise remote in Q . This is a point in the relative interiors of three pairwise remote edges of the centre, contrary to $Z_c = \emptyset$ in Definition 1.29(K). Thus Q is generic; the argument does not require the cusp to be empty.

Now apply each identity of Corollary 6.3 to its stated arguments and substitute $C = A$ at every one of them, including the deletions and halves just checked. For a chamber-side identity each sufficiently small punctured side lies in one generic chamber, on which $C = A$ holds pointwise, so the constant side values of C and of A agree; no value at the wall centre and no limit is used. This gives every identity for C on the domain stated in Corollary 6.3, including the displayed cusp law. $\square$

Proofs for Section 7

Proof of Proposition 7.1

Proof of 7.1. First let $r > 0$, $N = 2r + 1$, and denote the traversal vertices by $v_t = \mu_{t+1} = u_{rt}$, with t used only as a zero-based position in this local calculation. For $0 < \ell < N$ let $\text{res}(\ell)$ be the least residue of $r\ell$ modulo N . Rotating the starting circle point to residue zero shows, for $0 < p < \ell < N$, that

$$\chi(v_t, v_{t+p}, v_{t+\ell}) = 1 \iff \text{res}(p) < \text{res}(\ell). \quad (347)$$

This is the orientation of a triple of distinct points in their positive circular order. Since $2r = N - 1$, integer reduction gives

$$\text{res}(2a + 1) = r - a \quad (0 \leq a < r), \quad \text{res}(2a) = N - a \quad (1 \leq a \leq r). \quad (348)$$

In particular the residues of odd lengths are at most r and decrease with length, whereas those of positive even lengths are at least $r + 1$.

At a tree vertex let its positive child lengths be $\ell_1, \dots, \ell_J$, with total $\ell < N$ and partial sums s_k . When all child lengths are odd, the sum of two consecutive child lengths is even. Applying (347) to their endpoints gives $d_k = -1$ at every cut. The far sign satisfies

$$h_k = 1 \iff \text{res}(s_k) > \text{res}(\ell). \quad (349)$$

If ℓ is odd, J is odd. For odd k , the two residues in (349) are odd-length residues and $s_k < \ell$, so the inequality holds. For even k , its left residue is at least $r + 1$ and its right residue at most r , so it again holds. Thus every ordinary gate passes; there are $J - 1$ even many factors $d_k = -1$, and the ordinary weight before its extra vertex minus is 1. If ℓ is even, s_1 is odd and its residue is at most r , while the total residue is at least $r + 1$. Hence $h_1 = -1 = d_1$ and the ordinary gate fails.

These facts prove inductively that all even ordinary branch sums vanish. In fact the claim holds at length two, and at greater length any nonzero product of proper child sums has all child lengths odd by induction; the preceding even-total gate then kills it. For odd total length, only odd child lengths remain, with odd arity $J \geq 3$, and the ordinary top sign is -1 . Let $G(x) = \sum_{\ell \geq 1} b_\ell x^\ell$, read coefficientwise only through the largest required length $2r$. The recursion of Definition 2.4 therefore gives

$$G = x - \sum_{j \geq 1} G^{2j+1} = x - \frac{G^3}{1 - G^2}. \quad (350)$$

The denominator is invertible as a formal series because $G(0) = 0$. Multiply by $1 - G^2$ to obtain $G - G^3 = x - xG^2 - G^3$; cancel the cubic terms to obtain $G = x - xG^2$. Writing $G = xF$ gives

$$F = 1 - x^2 F^2. \quad (351)$$

Its constant coefficient is 1, and successive coefficient comparison gives $b_{2a+1} = (-1)^a \text{Cat}_a$, with $b_{2a} = 0$. Here the Catalan recurrence follows directly from $C(t) = 1 + tC(t)^2$: the solution with constant term one is $(1 - \sqrt{1 - 4t})/(2t)$, whose binomial expansion has coefficient $\frac{1}{a+1} \binom{2a}{a}$. Substituting $t = -x^2$ verifies the stated closed form rather than assuming a normalization for b .

At the barred root the total length is $2r$ and its residue is $r + 1$. Only odd child lengths can contribute. Two odd children give one passing barred gate of value -1 . Four or more odd children fail at the second gate: its partial length is a proper positive even length, so its residue is greater

than $r + 1$, giving $h_2 = 1$ while $d_2 = -1$. The unary barred root contributes the already zero branch b_{2r} . Consequently

$$A_g(K_r) = - \sum_{a+b=r-1} b_{2a+1} b_{2b+1}, \quad (352)$$

$$= (-1)^r \sum_{a+b=r-1} \text{Cat}_a \text{Cat}_b = (-1)^r \text{Cat}_r. \quad (353)$$

The start t was arbitrary, so this is a separate calculation for every physical root, not an appeal to general root independence. Since N is odd, Proposition 2.26(ii) negates the value under physical reversal, with the physical root map $g \mapsto 1 - g$. This proves the negative-rank case of the stated formula as well. The computation includes $r = 1$: the sole barred binary gate is -1 .

For K_0 , Lemma 6.1(ii) supplies the rootwise value. Its equation (58) is the direct four-vertex tree recursion, and table (59) computes its value -1 at each of the four physical SM roots using exactly the coordinates of Definition 5.1. Those four explicit determinant rows, rather than the general root-independence conclusion, supply this rank-zero computation. Cyclic covariance in Proposition 2.26(i) gives every shifted tuple and root. This completes all cases of the stated formula. $\square$

Proof of Proposition 7.3

Proof of 7.3. Take $r > 0$ and put $N = 2r + 1$. At a selected crossing with oriented directions v, w , the two smoothing corners have determinants $\det(v, w)$ and $\det(w, v)$, which are nonzero and opposite. If no subpolygon of S were mixed, the carrier containing a right smoothing corner would turn right everywhere. It could contain no original vertex, since all original vertices turn left. Its segments are forward subsegments of star edges, each having the origin strictly on its left.

Such a closed all-right polygon is impossible. Translate the common left point to the origin and number the carrier vertices $p_0, \dots, p_{q-1}$ cyclically and choose successive polar lifts α_i with $\alpha_{i+1} - \alpha_i \in (0, \pi)$. Indeed strict leftness is $\det(p_i, p_{i+1}) > 0$, so precisely this positive principal polar increment exists on each edge. Its direction lift β_i lies in $(\alpha_{i+1}, \alpha_i + \pi)$. Consecutive such lifts differ by a number in $(-\pi, \pi)$: both lie above α_i and below $\alpha_i + \pi$ at their shared corner after using the neighbouring intervals. Thus a right turn forces $\beta_i < \beta_{i-1}$. Extend the list periodically. Closure gives $\alpha_q = \alpha_0 + 2\pi w$ with a positive integer w , hence $\beta_q = \beta_0 + 2\pi w > \beta_0$, contradicting strict decrease around the closed carrier. The left/right-reversed argument applies to K_{-r} . Thus every nonempty support dies before coefficient extraction.

For completeness, every pair of nonadjacent physical star edges crosses. Here is the endpoint count without a second division by two. Rotate one edge to the chord with circular endpoints $0, r$. A second edge has endpoints $b, b + r$ modulo N , with $b \notin \{0, r, r + 1\}$ exactly when it is nonadjacent. If $1 \leq b \leq r - 1$, its first endpoint is strictly in $(0, r)$ and its second lies in (r, N) . If $r + 2 \leq b \leq 2r$, its first endpoint lies in (r, N) and its second reduces to $b - r - 1 \in (0, r)$. These are all nonadjacent cases. Two chords of a circle cross in their relative interiors precisely when their endpoints alternate: the line of one chord separates the two arcs of its complement, so the other chord crosses it precisely when its endpoints are in the different arcs. Thus each edge crosses exactly $N - 3$ other edges. Counting unordered pairs once gives

$$m = \frac{N(N - 3)}{2} = \frac{(2r + 1)(2r - 2)}{2} = N(r - 1). \quad (354)$$

Each crossing is counted twice in $N(N - 3)$, not in $N(r - 1)$. For $r = 1$ the count is zero and there is no nonempty support. The empty support always has the single uniform carrier K_r . Its rotation

and all original turn signs are Lemma 5.2(i). For negative r , reversal preserves the crossing count and negates all turns and the rotation, giving the stated absolute-value formula. $\square$

Proof of Lemma 7.4

Proof of 7.4. Round the star's corners in disjoint small neighbourhoods missing all crossings; this adds or deletes no crossing visit. We construct a finite positive braid word with that full decorated record, and convert it to the displayed word by operations preserving the entire record.

Polar cutting and an actual closure. By Lemma 5.2(ii), the origin is strictly left of every forward star edge. Along the straight segment $p(t)$ this gives

$$\frac{d\theta}{dt} = \frac{\det(p(t), p'(t))}{|p(t)|^2} > 0. \quad (355)$$

Each chord advances the polar argument by $2\pi r/N < \pi$, so the whole traversal advances it by $2\pi r$. Choose a ray missing all crossings and vertices. It has exactly r transverse preimages: a continuous strictly increasing angular lift covers each residue modulo 2π exactly r times in one traversal. Cutting at these preimages gives r strands, each strictly increasing in angular coordinate through one turn. At a vertex the angular coordinate remains strictly increasing on the two one-sided segments.

On the cut annulus use page coordinates $(x, y) = (\theta, -\rho)$. Their Jacobian relative to the oriented plane is $1/\rho > 0$. They preserve every decorated crossing sign. The radial positions of the r strands have a fixed order between crossing angles. At a crossing the two equal positions are adjacent, since a third at that radius would be a forbidden triple point. Two crossings at the same angle use disjoint pairs of positions, for the same reason. Read such crossings in either order; no physical strand visits both, so their order changes no strand's named crossing record.

Replace the intervals between crossings by parallel ports and each crossing by the standard two-strand braid gadget, with port indices in increasing radial order. A clockwise rotation of the usual upward braid picture gives exactly this order in $(\theta, -\rho)$. Every positive determinant crossing is therefore a positive Artin letter σ_i . This constructs a finite positive word w on r strands. Reconnect equal port positions at the cut. The ordinary plane closure can be drawn with crossing-free exterior arcs: their endpoint pairs are nested around the rectangle; choose successively wider exterior U-shaped arcs. Their interiors avoid the rectangle and each other. These arcs add no crossing visits.

Following any strand through the gadgets and the closing ports now gives exactly its original oriented cyclic successor sequence. Every crossing name, its two paired visits, over/under choice and sign is retained. Thus the constructed actual closure and the actual positive star diagram have identical full decorated records. They are one component. Every component circle is therefore included in the record bijection; no exterior-face prescription is imposed.

The repeated sector word. Let F_i be the traversal edge from u_{ir} to $u_{(i+1)r}$, so $F_i = E_{i+1}$ in the SM convention. Rotation through $2\pi/N$ sends F_i to F_{i-2} , since $-2r \equiv 1 \pmod{N}$. Its action on unordered crossing pairs is free. A stabilizing rotation either fixes both indices, hence is the identity, or exchanges them. In the latter case twice its index shift is zero modulo odd N , so again the shift is zero and no nontrivial exchange occurs. There are $N(r-1)$ crossings by (354), and each orbit has N elements; hence there are $r-1$ orbits.

Choose the initial sector boundary away from all crossing and vertex angles. Each open sector of width $2\pi/N$ contains one crossing from each orbit. Rotation identifies its radial port ordering with that of the next sector. Ordering any simultaneous disjoint crossings in the same way under

these rotations makes the full word literally

$$w = c^N, \quad |c| = r - 1, \quad (356)$$

with c positive. If c omitted σ_k for some $1 \leq k < r$, its port permutation would preserve the two nonempty sets $\{1, \dots, k\}$ and $\{k + 1, \dots, r\}$. The same would hold for c^N . The cycles of the closure's directed port network are exactly the cycles of this permutation, so its closure would have at least two components. This contradicts the single star traversal. Thus c contains each of the $r - 1$ simple generators exactly once.

A finite conversion to the ascending word. For such a word, orient the path on indices $1, \dots, r - 1$ by putting an arrow from the earlier to the later of each neighbouring pair. Any two orders with the same path orientation differ by commutations of nonneighbouring letters. Here is a finite check: take the first letter of the target order and move it left in the current order. Every letter it passes must be incomparable in the partial order generated by those arrows; otherwise one of the two orders would violate that order. In particular their indices are nonneighbouring, so each interchange is a permitted disjoint commutation. Repeat on the remaining suffix. Each step fixes one more initial position.

A source of the path orientation can in this manner be put first. Moving the first letter to the end flips just its incident arrows. Likewise a sink can be put last and moved to the beginning. These source/sink flips connect all orientations of a finite tree. To prove this, remove a leaf and use induction on the smaller tree. Lift each source or sink flip there: if its vertex is the leaf's neighbour and the leaf edge obstructs the flip, first flip the leaf, which is always a source or sink. Then the required flip is allowed in the full tree. After completing the smaller tree's sequence, one last leaf flip sets its edge as desired. The one-vertex case is immediate. This proves the induction and supplies a finite sequence on the path. Consequently disjoint commutations and cyclic first/last moves convert c to $\sigma_1 \cdots \sigma_{r-1}$.

We verify those operations on the entire N th power as actual closed diagrams. Commuting two nonneighbouring letters uses disjoint physical strands: no strand visits both gadgets, and their disjoint port transpositions commute. The whole directed network, including every cyclic component and every decorated visit, is unchanged. Apply this commutation in each of the N copies. If $c = au$, moving its first letter to the end replaces $(au)^N$ by $(ua)^N$. The latter is literally the same cyclic sequence of gadgets cut just after its first a . Closing equal port positions gives the same directed network; numbering at the new cut means actual positions, not a preassigned physical strand label. Last-to-first moves are the inverse operation. All intermediate diagrams therefore have the same full decorated record. Composing these finitely many bijections gives the required bijection with the final standard closure.

The final closure is exactly $(\sigma_1 \cdots \sigma_{r-1})^N$ with $N = 2r + 1$, the specified $T(r, 2r + 1)$. The gcd is one since any common divisor of r and $2r + 1$ divides 1. This agrees with the single-component record already proved. Lemma 3.20 now gives equality of the source values for the two actual diagrams, since all cyclic successors, paired visits, signs, over/under bits and component circles have been matched. Apply the common substitution of Theorem 3.18 to obtain (63). No ambient equivalence theorem or torus coefficient formula is used for this scalar equality. $\square$

Proof of Lemma 7.8

Proof of 7.8. In $\sigma_i \sigma_i^{-1}$ or $\sigma_i^{-1} \sigma_i$, one physical strand is over at both crossings and the other is under at both. The signs are opposite. A small rectangle about the two adjacent positions $i, i + 1$ excludes all other strands and has empty bigon interior. Ordinary Reidemeister II removes the pair

relative to its four collars, changing no other crossing and giving the closure of the word with the pair deleted. Local invariance proves equality. This does not cancel a two-positive-crossing bigon.

For $|i - j| \geq 2$, the gadgets σ_i^ε and σ_j^η , with $\varepsilon, \eta \in \{1, -1\}$, involve disjoint physical strands. No strand visits both gadgets, so their interchange changes no strand's visit order. Their position transpositions commute because their supports are disjoint. The closed directed port networks therefore have identical oriented cycles under the crossing-name bijection, with every over/under bit and sign preserved. Lemma 3.20 gives equality for any exterior word and any number of closure components.

For $\sigma_i\sigma_{i+1}\sigma_i$ and $\sigma_{i+1}\sigma_i\sigma_{i+1}$, call the three incoming physical strands A, B, C . In the first word, A crosses above B and then above C , while B crosses above C . In the second, B crosses above C first, then A above C , then A above B . Both have the strict height order $A > B > C$ and the same outgoing permutation. They are the two sides of ordinary oriented Reidemeister III, with no extra strand in the local three-position rectangle. The local move relative to the six collars proves equality. Switching all three crossings gives the negative version with height order $C > B > A$.

Each defining relation can be substituted a finite number of times inside an arbitrary word, and inverse insertion and cancellation have just been proved. This gives the claim for the Artin quotient. It is not a representation theorem for arbitrary geometric braids, a classification of isotopic closed braids, or the polygon state-sum R3 assumption. $\square$

Proof of Lemma 7.9

Proof of 7.9. Form the directed port network on the cyclic sequence of crossing gadgets. The network of $D_n(uv)$ cuts this sequence just before u ; cutting just before v gives the network of $D_n(vu)$. Each gluing uses equal positions at consecutive gadget boundaries. Thus the new port numbering is the actual left-to-right position numbering at the new cut, not a fixed label assigned to a physical strand before its passage through u .

Keep every crossing gadget and both visits. The cyclic network is unchanged by the choice of cut. Following an output reaches the same next crossing input in both descriptions, and a whole cycle without crossings is retained. The induced map preserves every oriented component cycle, crossing pairing, over/under bit and sign. The actual plane closures have these records because their exterior arcs add no visits. Lemma 3.20 proves (65).

This includes an empty block, empty total word, and disconnected closure. A newly assigned component order need not be preserved: that order is auxiliary for the source value. No physical root of an original polygon is relabeled by this diagram equality.

For vwv^{-1} , slide v from the beginning to the end to obtain $wv^{-1}v$. Here v^{-1} is the reversed letter sequence with every exponent negated. Remove the finitely many adjacent inverse pairs in $v^{-1}v$ by Lemma 7.8. This is a finite scalar derivation on the supplied words, not Markov completeness. A cut preserves every exponent, and the inverse pairs cancel their sums, so conjugation preserves the exponent sum as well. $\square$

Proof of Lemma 7.14

Proof of 7.14. On the formal orthonormal basis of partitions let U add a box and D remove one. At any fixed degree these incidence maps are finite and transpose to each other. For distinct equal-size diagrams, a common one-box successor exists exactly when their symmetric difference is two boxes; it is their union, and their common predecessor is their intersection. Otherwise neither exists. For a single diagram its equal-row blocks give one more addable than removable box. Thus, entry by

entry including the empty diagram,

$$DU - UD = I. \quad (357)$$

Let v_0 be the empty diagram. A saturated chain labels each new box by its insertion step and is exactly a standard tableau, so the coefficients of $U^n v_0$ are f_λ . Commuting D through the n factors, and using $Dv_0 = 0$, gives $DU^n v_0 = nU^{n-1}v_0$. Consequently

$$\sum_{|\lambda|=n} f_\lambda^2 = \langle U^{n-1}v_0, DU^n v_0 \rangle, \quad (358)$$

$$= n \sum_{|\mu|=n-1} f_\mu^2 = n!, \quad (359)$$

where the initial sum is 1. No hook-length dimension formula was used.

For $|d| > 1$, both $\mathbf{a}_d$ and $\mathbf{a}_{-d}$ are positive for $q > 0$. Also $\mathbf{a}_1 = 1$, $\mathbf{a}_{-1} = 0$, and $\mathbf{a}_d + \mathbf{a}_{-d} = 1$. Cancelling $1 - q$ in numerator and denominator gives

$$\mathbf{a}_d(1) = \frac{d+1}{2d}. \quad (360)$$

Thus the chosen square root in (71) is smooth when $|d| > 1$, and is identically zero for $|d| = 1$. Forbidden swaps are precisely adjacent boxes in one row or column. All generator matrices are therefore real symmetric and continuous throughout $q > 0$. The relation verification in the selected source has no hidden zero distance: boxes on the same diagonal force two intermediate boxes, hence their entries differ by at least three. Its nontrivial scalar identity can also be checked directly. If $\ell = k + m$ and $k, m, \ell \neq 0$, put $X = q^k, Y = q^m$. Then

$$1 - \mathbf{a}_k - \mathbf{a}_m = -\frac{(1-q)(1-XY)}{(1+q)(1-X)(1-Y)}, \quad (361)$$

$$\mathbf{a}_k \mathbf{a}_m + \mathbf{a}_\ell(1 - \mathbf{a}_k - \mathbf{a}_m) = \frac{q}{(1+q)^2}. \quad (362)$$

For the second line the common numerator is $(1-qX)(1-qY) - (1-q)(1-qXY) = q(1-X)(1-Y)$. The source's three-entry expansion uses this identity for the off-diagonal terms and its product with $\mathbf{a}_k - \mathbf{a}_m$ for the diagonal difference. Together with idempotence and remote commutation it gives the braid relation after $g_i = q - (1+q)e_i^W$. Continuity extends it to $q = 1$.

We spell out the finite linear algebra needed for irreducibility. A finite-dimensional unital $*$ -algebra of matrices has orthogonal invariant complements. On an irreducible module choose a nonzero projection p in the algebra of minimal rank. If $p\mathcal{A}p$ is not scalar, a nonscalar self-adjoint element has a proper nonzero spectral projection, contrary to minimality. If p had orthogonal nonzero vectors v, w in its range, then for all $a, b \in \mathcal{A}$

$$\langle av, bw \rangle = \langle pb^*apv, w \rangle = 0. \quad (363)$$

Both cyclic subspaces would be the whole irreducible module, a contradiction. Thus p has rank one, and apb^* span all matrix units. The irreducible image is the full matrix algebra.

Induct now on n . A nonzero invariant subspace of V_λ contains one restriction block: the inductive central projections in H_{n-1} extract it, and its smaller module is irreducible. Any two removable corners in rows $a < b$ have content difference $\lambda_a - \lambda_b + b - a \geq 2$. Remove both corners, fill the remaining diagram by a standard tableau, and put $n-1, n$ in their two possible orders. The off-diagonal entry of e_{n-1}^W is nonzero. Projecting to the other restriction block extracts a nonzero

vector and hence the whole block. Every block is reached. With one removable corner, restriction alone is irreducible.

Different shapes of size $n \geq 3$ have different predecessor sets by the selected finite diagram fact, so their multiplicity-free restrictions are inequivalent. At $n = 2$ their generator values $q, -1$ are distinct. To see that the simultaneous image is the product of the matrix blocks, the kernels of inequivalent irreducible quotients are distinct maximal two-sided ideals: equal kernels would give equivalent modules for the same full matrix algebra. They are pairwise comaximal; the finite Chinese-remainder decomposition, or equivalently the orthogonal central-ideal decomposition of the finite $*$ -algebra, gives the product. Its dimension is the just-proved $\sum f_\lambda^2 = n!$. The upper bound of Lemma 7.13 therefore forces faithfulness and (75). At $q = 1$ the generator map to adjacent transpositions surjects onto $\mathbb{C}[S_n]$. That algebra has dimension $n!$, so the surjection is an isomorphism. This also establishes the asserted label convention. $\square$

Proof of Lemma 7.15

Proof of 7.15. Pad μ by zeros to k rows, with k strictly larger than its number of positive rows, and set $l_i = \mu_i + k - i$. These are strictly decreasing nonnegative integers. Row i has precisely the hook multiset

$$\{1, \dots, l_i\} \setminus \{l_i - l_j : i < j \leq k\}. \quad (364)$$

Indeed the hook in column c is $l_i - m_c$, where $m_c = c + k - \mu'_c - 1$ increases strictly. Before m_1 the missing integers are the l_j of zero rows. Between m_c and m_{c+1} the missing integers are $c + k - \mu'_c, \dots, c + k - \mu'_{c+1} - 1$, exactly the l_j of rows of length c . After the last m_c they are the l_j of the later rows having length μ_i . These three cases partition $j > i$. For a zero row the exclusions exhaust the whole interval, as required. Thus multiplication over rows gives

$$\prod_{(i,j) \in \mu} (1 - q^{h_{\mu}(i,j)}) = \frac{\prod_{i=1}^k \prod_{s=1}^{l_i} (1 - q^s)}{\prod_{i < j} (1 - q^{l_i - l_j})}. \quad (365)$$

The uncanceled boundary rational function is

$$F(t) = \frac{\prod_{i=1}^k (t - q^{\mu_i - i})}{(t - q^{-k}) \prod_{i=1}^k (t - q^{\mu_i + 1 - i})}. \quad (366)$$

An equal numerator and denominator exponent can occur only for adjacent equal rows. In a maximal block of equal positive rows $a, \dots, b$ of length v , cancellation leaves numerator content $v - b$ and denominator content $v + 1 - a$. The last zero block cancels $t - q^{-k}$ and leaves the new-row addable content. Hence

$$F(t) = \frac{\prod_{d \in \mathcal{B}} (t - y_d)}{\prod_{c \in \mathcal{A}} (t - x_c)}. \quad (367)$$

This also proves that the contents within each set are distinct, the sets are disjoint, and $|\mathcal{A}| = |\mathcal{B}| + 1$. The sums of the exponents in the numerator and denominator of (366) agree: $\sum_i (\mu_i - i) = -k + \sum_i (\mu_i + 1 - i)$. Cancellation preserves this equality. Consequently $\prod_c x_c = \prod_d y_d$ and $F(0) = -1$.

All poles are simple. Multiplying by the denominator and evaluating at each pole proves the finite partial-fraction identity

$$F(t) = \sum_{c \in \mathcal{A}} \frac{x_c D_c}{t - x_c}. \quad (368)$$

There is no polynomial term, since the degrees differ by one. Evaluation at zero and comparison of t^{-1} coefficients at infinity give, respectively,

$$\sum_{c \in \mathcal{A}} D_c = 1, \quad \sum_{c \in \mathcal{A}} x_c D_c = 1. \quad (369)$$

If $u = q^{d_0}$ for a removable content d_0 , then $F(u) = 0$ and u is not a pole. Thus

$$\sum_{c \in \mathcal{A}} \frac{x_c D_c}{x_c - u} = 0, \quad (370)$$

$$\sum_{c \in \mathcal{A}} \frac{D_c}{x_c - u} = -\frac{1}{u}. \quad (371)$$

For the second equation substitute $x_c = (x_c - u) + u$ in the first, then use $\sum D_c = 1$ and divide by the nonzero u . These are finite algebraic identities, not an analytic residue argument.

Suppose the added box is in row i , with content $c = \mu_i + 1 - i$. Only l_i changes, to $l_i + 1$, in (365). The numerator of the weight gains $\mathbf{a} - \mathbf{b}q^c$, and its leading power gains q^{i-1} . Taking the old and new products in the field gives

$$\begin{aligned} \frac{W_{\mu+c}}{W_\mu} &= \frac{q^{i-1}(\mathbf{a} - \mathbf{b}q^c)}{1 - q^{l_i+1}} \prod_{j < i} \frac{1 - q^{l_j - l_i - 1}}{1 - q^{l_j - l_i}} \\ &\quad \cdot \prod_{j > i} \frac{1 - q^{l_i + 1 - l_j}}{1 - q^{l_i - l_j}}. \end{aligned} \quad (372)$$

The pole $x_c = q^{l_i+1-k}$ is not cancelled in (366). Its residue divided by x_c is

$$D_c = x_c^{-1} \frac{1 - q^{-1}}{1 - q^{-l_i-1}} \prod_{j \neq i} \frac{1 - q^{l_j - l_i - 1}}{1 - q^{l_j - l_i}}. \quad (373)$$

The prefactor is $q^{k-1}(1 - q)/(1 - q^{l_i+1})$. Each of the $k - i$ factors with $j > i$ becomes $q^{-1}(1 - q^{l_i+1-l_j})/(1 - q^{l_i-l_j})$ upon converting its two negative exponents. The total power is $q^{k-1-(k-i)} = q^{i-1}$. Comparison with (372) proves (79). Its product form remains valid when a later specialization makes W_μ zero; no specialized division by W_μ is used. $\square$

Proof of Theorem 7.17

Proof of 7.17. Matrix traces are cyclic. Grouping the restriction blocks (72) by μ makes their total coefficient $\sum_c W_{\mu+c} = W_\mu$, proving compatibility. The one-box weight is $(\mathbf{a} - \mathbf{b})/(1 - q) = 1$; compatibility gives normalization at every positive size.

For the Markov equation fix a tableau shape of size $n + 1$. Every $X \in H_n$ preserves the position of $n + 1$, and hence each restriction block. In a term $X_{T,S}(e_n^W)_{S,T}$ of the matrix trace, an off-diagonal nonzero entry of e_n^W must exchange $n, n + 1$. It changes the position of $n + 1$, so $X_{T,S} = 0$. A forbidden exchange already has zero coefficient. Only the diagonal entry $\mathbf{a}_{c_T(n) - c_T(n+1)}$ contributes, even when X is an arbitrary non-diagonal matrix inside each block. Fix a smaller tableau of shape μ ; its largest entry occupies a removable box of content d_0 . Its diagonal X entry is the same for every extension. Summing those extensions multiplies it by $\sum_c \mathbf{a}_{d_0-c} W_{\mu+c} = \eta W_\mu$. Summing over smaller tableaux and shapes proves (87).

For uniqueness, two such traces agree on H_1 . If they agree on H_n , cyclicity and the Markov condition force their value on $Xe_n^W Y$ to be $\eta\tau_n(YX)$. They already agree on H_n , and (73) therefore makes them agree on H_{n+1} . This also shows descent from the square-root matrix field to the rational base field: all values are obtained inductively from base-field coefficients and η .

Normalization and $g_n = q - (1+q)e_n^W$ prove the generator equation and (88). Set $w = 1 - q + z_J$. Multiplying each box numerator by its row power gives

$$q^{n(\lambda)} \prod_{(i,j) \in \lambda} (w - z_J q^{j-i}) = \prod_{(i,j) \in \lambda} (q^{i-1} w - q^{j-1} z_J). \quad (374)$$

Row i contributes $q^{(i-1)\lambda_i}$, whose product is exactly $q^{n(\lambda)}$. The right-hand factors are Jones's Figure 5.3; the denominator is his unnumbered hook diagram immediately before (5.4), printed p. 346 [15]. Thus the proved product is Jones's (5.4), and (86) is his (5.5). This Wenzl–Jones dictionary uses the same q , the same generators and the same partition label, but complementary projectors.

For clarity, substituting (82) of Lemma 7.16 into an individual box numerator gives

$$\mathfrak{a} - \mathfrak{b}q^c = 1 - q^{c+1} - (1+q)(1-q^c)\eta. \quad (375)$$

This is the internally proved normalized formula, not Wenzl's misprinted (3.6). No Schur or Pieri identity was used to establish it. $\square$

Proof of Lemma 7.18

Proof of 7.18. For a finite group, average an inner product to make its representations unitary. Averaging the action on $\text{Hom}(V, W)$ projects onto the intertwining maps; its trace is the inner product of the two characters. Schur's lemma and orthogonal invariant complements show that an actual character of squared norm one is irreducible. These facts are finite-dimensional linear algebra, not a character-table input.

Let V be the coordinate-sum-zero subspace of $\mathbb{C}^n$ with its permutation action, and let χ_k be the character of $\bigwedge^k V$, $0 \leq k < n$. If a permutation has m_ℓ cycles of length ℓ , then

$$\sum_{k=0}^{n-1} \chi_k(g) t^k = \det(I + tg|_V) = \frac{\prod_{\ell \geq 1} (1 - (-t)^\ell)^{m_\ell}}{1 + t}. \quad (376)$$

An ℓ -cycle has the ℓ th roots of unity as eigenvalues, giving its numerator factor; removing the constant line divides by $1 + t$. The number of permutations of that cycle type is $n! / \prod_\ell (\ell^{m_\ell} m_\ell!)$: order the symbols, then divide by rotations within cycles and permutations of equal-length cycles. Therefore, for $n \geq 1$, the character inner products obey

$$\sum_{k,j} \langle \chi_k, \chi_j \rangle t^k s^j = [u^n] \frac{(1+tu)(1+su)}{(1-u)(1-tsu)(1+t)(1+s)}. \quad (377)$$

To obtain this expression, sum over the nonnegative integers m_ℓ using their just-derived cycle count. Before dividing by $(1+t)(1+s)$, the generating series is

$$\exp \left(\sum_{\ell \geq 1} \frac{u^\ell (1 - (-t)^\ell) (1 - (-s)^\ell)}{\ell} \right) = \frac{(1+tu)(1+su)}{(1-u)(1-tsu)}. \quad (378)$$

Expand the four summands in the exponent and use the formal identity $\sum_{\ell \geq 1} v^\ell / \ell = -\log(1 - v)$. Each coefficient of u is finite; no analytic convergence or infinite probabilistic averaging is used. The characters are real integers by (376), so their product is the required character inner-product integrand.

Put $S_m = \sum_{h=0}^m (ts)^h$ for $m \geq 0$ and $S_{-1} = 0$. The numerator coefficient on the right of (377) is $S_n + (t+s)S_{n-1} + tsS_{n-2}$. For $n = 1$ its last term is zero; for every $n \geq 1$ it equals $(1+t)(1+s)S_{n-1}$, since the difference is $(ts)^n - ts(ts)^{n-1} = 0$. Consequently

$$\langle \chi_k, \chi_j \rangle = \delta_{kj}. \quad (379)$$

The n exterior powers are distinct irreducible modules.

Their labels are hooks, not merely modules of matching dimensions. On restriction to S_{n-1} , the sum-zero module is the smaller sum-zero module plus a trivial line. Its k th exterior power is therefore the sum of the smaller k th and $(k-1)$ st powers, omitting out-of-range summands. Induction gives precisely the removable-box predecessors of $(n-k, 1^k)$. The predecessor-set distinction in Literature input 6, the complete family in Lemma 7.14, and the trivial/sign base at $n = 2$ identify the label uniquely. The endpoint exterior powers are the trivial and sign modules.

For an n -cycle c , (376) becomes

$$\frac{1 - (-t)^n}{1 + t} = \sum_{k=0}^{n-1} (-t)^k. \quad (380)$$

Thus all hooks have the claimed value. To exclude nonhooks, the matrix-block decomposition of $\mathbb{C}[S_n]$ makes its irreducible characters a basis of class functions: central class sums and scalar block identities are two bases of the center. The indicator of the conjugacy class of c has squared norm $1/n$, since the centralizer of c consists of its n powers. Its coefficient at an irreducible character is $\chi_\lambda(c)/n$. Parseval gives

$$\sum_{\lambda} |\chi_\lambda(c)|^2 = n. \quad (381)$$

The n hooks already contribute n , leaving every nonhook zero.

A hook tableau is specified by the k entries from $\{2, \dots, n\}$ placed below its corner. Their column order and the remaining row order are forced, proving its binomial dimension. In such a tableau, 2 lies either immediately right of 1 or immediately below it. The first case has $e_1^W = 0$, hence $e_1^J = 1$; the second has $e_1^J = 0$. The first case permits exactly $\binom{n-2}{k}$ choices of the entries below the corner. This proves all of (89), including the complementary-projector convention. $\square$

Proof of Lemma 7.19

Proof of 7.19. The generators are invertible by the explicit formula $g_i^{-1} = q^{-1}(g_i - (q-1))$. The triangular word has both forms

$$\Delta_n = \Delta_{n-1}(g_{n-1} \cdots g_1) = \delta_n \Delta_{n-1}. \quad (382)$$

The first is its definition. For the second, insert the induction formula for Δ_{n-1} and commute g_{n-1} past Δ_{n-2} . Reversing all word orders in these two forms proves by induction that word reversal fixes Δ_n . For $i \geq 2$, the descending suffix satisfies $(g_{n-1} \cdots g_1)g_i = g_{i-1}(g_{n-1} \cdots g_1)$ by (74); induction then gives $\Delta_n g_i = g_{n-i} \Delta_n$. For $i = 1$ use the second form in (382), induction in Δ_{n-1} , and $\delta_n g_{n-2} = g_{n-1} \delta_n$, the ascending version of that same word calculation. The $n = 2$ case is direct.

Apply this conjugation to the word Δ_n itself, then cancel the invertible right Δ_n . Index reversal $g_i \mapsto g_{n-i}$ therefore fixes Δ_n as well. Combining index reversal and word reversal with (382) gives

$$\Delta_n = \text{sh}_1(\Delta_{n-1})\delta_n, \quad (383)$$

where sh_r raises every generator index by r . For $r = n - k$ this yields the ladder identity

$$\text{sh}_r(\Delta_k)\delta_n = \delta_r \text{sh}_{r-1}(\Delta_{k+1}), \quad 1 \leq k < n. \quad (384)$$

Indeed the prefix $g_1 \cdots g_{r-1}$ of δ_n commutes past the shifted triangle. The remaining equality is precisely (383) on generators $g_r, \dots, g_{n-1}$. Induction using (384) now proves

$$\delta_n^k = (\delta_n \delta_{n-1} \cdots \delta_{n-k+1}) \text{sh}_{n-k}(\Delta_k), \quad 1 \leq k \leq n. \quad (385)$$

The base $k = 1$ has $\Delta_1 = 1$. Multiplication on the right by δ_n and (384) give the next case. At $k = n$, iteration of (382) makes the first factor Δ_n and the second is also Δ_n . Hence $\delta_n^n = \Delta_n^2$. Applying the conjugation formula twice shows that this square commutes with every generator. Only braid-word relations, not equivalence or classification of closed braids, were used. $\square$

Proof of Lemma 7.20

Proof of 7.20. In an irreducible V_λ of dimension d_λ , centrality of $\Delta_n^2 = \delta_n^n$ (Lemma 7.19) makes it scalar. The generators are conjugate: the braid word $g_i g_{i+1}$ conjugates g_i to g_{i+1} . Thus the common rank $\rho_\lambda = \text{rank}(e_i^J)$ is independent of i . The eigenvalues of each g_i are $q, -1$, so

$$\det \pi_\lambda(g_i) = (-1)^{d_\lambda - \rho_\lambda} q^{\rho_\lambda}. \quad (386)$$

The full twist contains $n(n-1)$ generator factors, an even number. Its determinant is $q^{\rho_\lambda n(n-1)}$. Its scalar is therefore

$$\omega(q) q^{\rho_\lambda n(n-1)/d_\lambda}, \quad \omega(q)^{d_\lambda} = 1, \quad (387)$$

using the positive real power for $q > 0$. The projector trace ρ_λ is a continuous integer, hence constant. The scalar and $\omega(q)$ are continuous, and the latter takes values in a finite set. At $q = 1$, δ_n is an n -cycle in S_n , so its n th power is 1 and $\omega(1) = 1$. Connectedness of $q > 0$ gives $\omega(q) = 1$ everywhere.

The operator $M_\lambda(q) = q^{-\rho_\lambda(n-1)/d_\lambda} \pi_\lambda(\delta_n)$ therefore has n th power I . The polynomial $t^n - 1$ has distinct roots, so M_λ is diagonalizable. For each n th root ζ its spectral projector is explicitly

$$E_\zeta(q) = \frac{1}{n} \sum_{j=0}^{n-1} \zeta^{-j} M_\lambda(q)^j. \quad (388)$$

Its trace is a continuous integer multiplicity, thus constant on $q > 0$. All eigenvalue multiplicities equal their values at $q = 1$. Because $\gcd(m, n) = 1$, the permutation c^m is again an n -cycle. Lemma 7.18 makes the trace of $M_\lambda(q)^m$ zero for nonhooks and $(-1)^k$ for a hook. Restoring its scalar factor and using $\rho_\lambda/d_\lambda = (n-1-k)/(n-1)$ proves (91).

The weighted trace is rational by Theorem 7.17; the resulting hook-only expression is explicitly rational as well. Clearing their denominators, equality on an open real parameter set proves the rational-function identity. This step is not substitution of $q = 1$ into the uncanceled weights. $\square$

Proof of Lemma 7.21

Proof of 7.21. Theorem 3.18 and Lemma 7.6 give (92) for the same construction, including all UNDER-first component initializations. Lemma 3.19 gives the added-circle factor $(t^{-1} - t)/x$ for an actual simple crossing-free component with no mixed crossings, nesting included. The variables in (93) give

$$\frac{t^{-1} - t}{x} = \frac{1 - s^2 \xi^2}{s(\xi^2 - 1)} = K. \quad (389)$$

All divisions are in a rational-function field, not numerical claims at vanishing denominators. Define on this Artin group

$$T_n(\beta) = K^{1-n} s^{-e(\beta)} P_J(\widehat{\beta}; t, x). \quad (390)$$

Lemma 7.8 proves the polynomial equalities for every finite defining-relation substitution and inverse-pair cancellation. Each also preserves the exponent sum, so T_n is well defined on the stated group and extends linearly to its group algebra. Lemma 7.12 supplies the actual last-crossing skein triple for βg_i^2 : switching cancels the adjacent inverse pair and gives β , while oriented smoothing removes only the last gadget and gives βg_i . Thus

$$P_J(\widehat{\beta g_i^2}) = tx P_J(\widehat{\beta g_i}) + t^2 P_J(\widehat{\beta}). \quad (391)$$

After the normalization (390), the two coefficients are $tx/s = q - 1$ and $t^2/s^2 = q$. Thus T_n annihilates every left multiple of the quadratic relation in (69). The full-record cut-slide of Lemma 7.9, together with the unchanged exponent sum, gives $T_n(\beta\gamma) = T_n(\gamma\beta)$, so it also annihilates arbitrary left and right multiples. It descends to a cyclic linear functional on H_n . The group-algebra quotient agrees with the generator presentation because $g_i^{-1} = q^{-1}(g_i - (q - 1))$.

Lemma 7.10 gives the identity diagram's n crossing-free circles and value K^{n-1} , so $T_n(1) = K^{1-n} K^{n-1} = 1$. The same lemma retains every old component and adds exactly one crossing-free cycle for a free last strand. Its factor K cancels the new K^{-1} in the normalization, proving $T_{n+1}|_{H_n} = T_n$. Lemma 7.11 proves the positive-stabilization polynomial equality on the actual closures, including an old crossing-free chosen component, by the local RI and full-record certificates. Exactly one strand and one exponent are added, hence

$$T_{n+1}(\beta g_n) = (Ks)^{-1} T_n(\beta), \quad (Ks)^{-1} = \frac{q - 1}{1 - \lambda_J q} = z_J. \quad (392)$$

Theorem 7.17 gives uniqueness, with $\eta = (q - z_J)/(1 + q)$, so $T_n = \tau_n$. Solving (390) proves (94), exactly the reduced-braid normalization of Jones's Definition 6.1. Only the supplied finite word, local move and full-record scalar certificates were used for the closure evaluations. The uniqueness invoked here is the separately proved finite algebraic trace theorem, not an unknot-only polynomial or a Markov-completeness theorem. $\square$

Proof of Theorem 7.22

Proof of 7.22. Write a hook as $(n - \gamma, 1^\gamma)$ and set $\beta = n - 1 - \gamma$. Its contents are exactly $-\gamma, \dots, \beta$, each once; its hook lengths are $n, 1, \dots, \gamma$ and $1, \dots, \beta$; and $n(\lambda) = \gamma(\gamma + 1)/2$. With $z_J = (q - 1)/(1 - \lambda_J q)$, the parameters of (88) give $\mathfrak{a} = \lambda_J q z_J$, $\mathfrak{b} = z_J$. Thus each box factor is

$$\mathfrak{a} - \mathfrak{b} q^i = \frac{1 - q}{1 - \lambda_J q} (q^i - \lambda_J q). \quad (393)$$

Insert these facts into (77), then multiply by the hook trace (91) of Lemma 7.20. Nonhooks vanish, leaving

$$\tau_n(\delta_n^m) = \frac{1}{1-q^n} \left(\frac{1-q}{1-\lambda_J q} \right)^n \cdot \sum_{\substack{\beta, \gamma \geq 0 \\ \beta + \gamma = n-1}} \frac{(-1)^\gamma q^{\beta m + \gamma(\gamma+1)/2} \prod_{i=-\gamma}^{\beta} (q^i - \lambda_J q)}{[\beta]!_q [\gamma]!_q}. \quad (394)$$

This is Jones's (9.6), now obtained from the explicit trace construction. The exponent sum of δ_n^m is $m(n-1)$. Multiply (394) by $K^{n-1} s^{m(n-1)}$ as required by (94). First, $K = -(1-\lambda_J q)/(s(1-q))$, so all but one of the n copies of $(1-q)/(1-\lambda_J q)$ cancel. Second, the powers of s leave $s^{(n-1)(m-1)} = \lambda_J^{(n-1)(m-1)/2}$. Third, $(-1)^{n-1}(-1)^\gamma = (-1)^\beta$ because $n-1 = \beta + \gamma$. These three separate cancellations give (95). Coprimality prevents both n, m from being even, so $(n-1)(m-1)$ is even and the displayed power of λ_J is integral. $\square$

Proof of Lemma 7.23

Proof of 7.23. Use the normalized positive-braid formula (95) of Theorem 7.22 for this diagram, the positive Artin closure of δ_n^m with δ_n the word of Lemma 7.19. Put $N = n-1$, $g = N(m-1)$ and $y = a^{-2}$. Coprimality prevents n, m from both being even, so g is even. Use the same positive crossing and the same diagram polynomial, with

$$t = a^{-1}, \quad x = \sqrt{q} - 1/\sqrt{q}, \quad \lambda_J = a^{-2} q^{-1} = y q^{-1}. \quad (395)$$

Thus $P_J(D; t, x) = P_D(a, x)$, directly in the printed skein convention. For the algebra one can take an indeterminate square root of q ; the final expressions below lie in $\mathbb{Q}(q)$ coefficientwise. No mirror of the knot or change of writhe is made.

Write $(q; q)_j = \prod_{h=1}^j (1 - q^h)$, with value one at $j = 0$. Formula (95) becomes

$$P_D(a, \sqrt{q} - 1/\sqrt{q}) = a^{-g} q^{-g/2} \frac{1-q}{1-q^n} \sum_{\beta + \gamma = N} \frac{(-1)^\beta q^{\beta m + \gamma(\gamma+1)/2} \prod_{i=-\gamma}^{\beta} (q^i - y)}{(q; q)_\beta (q; q)_\gamma} \frac{1}{1-y}. \quad (396)$$

Every summand has exactly one factor with $i = 0$, namely $1 - y$. Cancel that factor in each summand over the rational-function field. Its remaining polynomial has degree N in y and highest coefficient $(-1)^N$. Thus the whole expression is a^{-g} times a polynomial of degree at most N in a^{-2} . In particular the exponent $-g - 2N = -(n-1)(m+1)$ is its terminal possible exponent. Its coefficient is

$$L(q) = q^{-g/2} \frac{1-q}{1-q^n} \sum_{\gamma=0}^N \frac{(-1)^\gamma q^{(N-\gamma)m + \gamma(\gamma+1)/2}}{(q; q)_\gamma (q; q)_{N-\gamma}}. \quad (397)$$

Here $(-1)^N (-1)^\beta = (-1)^\gamma$, since $N = \beta + \gamma$.

For completeness define $B_q(N, j) = (q; q)_N / ((q; q)_j (q; q)_{N-j})$ for $0 \leq j \leq N$, and zero otherwise. Factoring the common factorial in the two terms shows

$$B_q(N, j) = B_q(N-1, j) + q^{N-j} B_q(N-1, j-1).$$

For an interior index the remaining numerator equality is $(1 - q^{N-j}) + q^{N-j}(1 - q^j) = 1 - q^N$; at either endpoint the one nonzero term is one. Induction by multiplying by the final factor $1 + vq^{N-1}$

therefore proves the finite identity

$$\prod_{h=0}^{N-1} (1 + vq^h) = \sum_{j=0}^N q^{j(j-1)/2} B_q(N, j) v^j. \quad (398)$$

Indeed the two contributions to v^j , after removal of $q^{j(j-1)/2}$, are the two terms in the displayed recurrence. The empty product at $N = 0$ starts the induction.

In (397) factor out $q^{Nm}/(q; q)_N$. The remaining sum is (398) with $v = -q^{1-m}$, because

$$\frac{\gamma(\gamma-1)}{2} + (1-m)\gamma = \frac{\gamma(\gamma+1)}{2} - m\gamma.$$

Consequently

$$\begin{aligned} \sum_{\gamma=0}^N \frac{(-1)^\gamma q^{(N-\gamma)m+\gamma(\gamma+1)/2}}{(q; q)_\gamma (q; q)_{N-\gamma}} &= \frac{q^{Nm}}{(q; q)_N} \prod_{h=1}^N (1 - q^{h-m}) \\ &= (-1)^N q^{N(N+1)/2} \prod_{h=1}^N \frac{1 - q^{m-h}}{1 - q^h}. \end{aligned} \quad (399)$$

For the second equality each numerator is $1 - q^{h-m} = -q^{h-m}(1 - q^{m-h})$. The N signs give $(-1)^N$ and the exponent, including q^{Nm} , is $Nm + \sum_{h=1}^N (h-m) = N(N+1)/2$.

Put $[j]_q = 1 + q + \dots + q^{j-1}$ for $j \geq 1$. Each $m-h$ is positive since $N = n-1 < m$. Cancel the factors $1-q$ in the last product and in $(1-q)/(1-q^n)$ before specializing. We obtain

$$L(q) = \frac{(-1)^N q^{(N(N+1)-g)/2}}{[n]_q} \prod_{h=1}^N \frac{[m-h]_q}{[h]_q}. \quad (400)$$

Its denominator is nonzero at $q = 1$, since $[j]_1 = j$. Therefore

$$\begin{aligned} L(1) &= \frac{(-1)^N}{n} \prod_{h=1}^N \frac{m-h}{h} \\ &= \frac{(-1)^{n-1}}{n} \binom{m-1}{n-1}. \end{aligned} \quad (401)$$

Finally, the supplied braid closure is a knot: its port permutation is the m th power of an n -cycle and has one cycle because $\gcd(m, n) = 1$. By Theorem 3.18, its complete polynomial has nonnegative even z exponents. Substituting $z = \sqrt{q} - 1/\sqrt{q}$ is thus a finite Laurent polynomial in a, q , since $z^2 = q + q^{-1} - 2$. At $q = 1$ it equals exactly the z^0 row of P_D . Taking the single a coefficient commutes with this specialization because both operations act on a finite polynomial. Equation (401) proves the claim. The cancellation of singular factors occurred earlier; no individual singular summand of (396) was specialized. $\square$

Proof of Theorem 7.25

Proof of 7.25. For $r \geq 2$, put $N = 2r + 1$. By Proposition 7.3 only the empty decomposition contributes. Lemma 5.2(i) gives all N turns left, so $\ell(K_r) = N$ and $(-1)^\ell = (-1)^N = -1$, and $d = 1 - N(r-1) - r = 2 - 2r^2$. The corner rounding that produces D_r in Lemma 7.4 takes place in

disjoint discs away from every crossing; it preserves the oriented component, the cyclic order of the crossing visits, their pairing, the over/under bits and the crossing signs of the positive lift used in Definition 3.28. The polygonal positive lift and D_r therefore have the same named decorated record, so Lemma 3.20 equates their source polynomials and the common Laurent substitution equates their local P -polynomials. Lemma 7.4 then equates this polynomial with that of the specified standard positive braid closure. Apply Lemma 7.23 to that closure with $n = r$ and $m = 2r + 1$; these are coprime because any common divisor divides $(2r + 1) - 2r = 1$. Its terminal exponent is $-(r - 1)(2r + 2) = 2 - 2r^2$, exactly the exponent d above. The magnitude of the terminal coefficient is

$$\frac{1}{r} \binom{2r}{r-1} = \frac{(2r)!}{r(r-1)!(r+1)!} = \frac{(2r)!}{r!(r+1)!} = \text{Cat}_r. \quad (402)$$

Here the first equality expands the binomial coefficient, the second uses $r(r-1)! = r!$, and the last is the defining Catalan formula. The sign supplied by Lemma 7.23 is $(-1)^{r-1}$. Theorem 3.18 identifies the diagram's P with the H used in the definition of c , so $c(K_r) = (-1)^{r-1} \text{Cat}_r$ and $C(K_r) = (-1)^r \text{Cat}_r$. This calculation uses the actual diagram records and does not transport a polynomial along the separate S^3 knot equivalence. For $r = 1$: $C(K_1) = (-1)^3 \cdot 1 = -1$ by Lemma 3.70(i). Negative r : Proposition 4.9 with N odd. For K_0 : the empty decomposition is mixed (turns $-, +, +, -$); the decomposition at the unique crossing $\{1, 3\}$ yields two triangles, one all-left and one all-right, each embedded with $c = 1$, so $C(K_0) = (-1)^2 \cdot (-1)^1 \cdot 1 \cdot 1 = -1$. Comparison with Proposition 7.1 gives the equalities. $\square$

Proofs for Section 8

Proof of Lemma 8.2

Proof of 8.2. Translation cancels in the difference of two image vertices, so $f\mu_j - f\mu_i = M(\mu_j - \mu_i)$. The determinant of the two image difference vectors therefore satisfies

$$\det(f\mu_j - f\mu_i, f\mu_k - f\mu_i) = \det(M) \det(\mu_j - \mu_i, \mu_k - \mu_i). \quad (403)$$

Since $\det M$ is positive, the two determinants have the same sign. This proves (99), preserves (G1), and preserves every original corner-turn sign. The left-turn count is consequently unchanged. An affine bijection preserves relative interiors of segments and their incidences. A triple interior concurrence after mapping would pull back to one before mapping, so (G2) is preserved as well (Definition 1.5). Thus fP is generic.

For $0 \leq t \leq 1$, direct substitution into $f(x) = Mx + b$ gives

$$f((1-t)x + ty) = (1-t)f(x) + tf(y). \quad (404)$$

Hence the crossing parameters on each directed edge are unchanged. Injectivity gives a bijection of crossing points, and unchanged parameters preserve the order of visits. The two-visit pairing and interlacement are therefore preserved, identifying exactly the independent crossing subsets of Definition 3.1. At each selected crossing the same incoming visit is reconnected to the same outgoing successor. Applying f to all inherited segments thus produces exactly the corresponding carrier; the incoming-visit assignment is unchanged. Its corners and self-crossings correspond bijectively. For any two successive carrier directions u, v the identity $\det(Mu, Mv) = \det(M) \det(u, v)$ preserves their turn sign. This includes the selected smoothing corners and proves the asserted corner counts and uniformity assertions.

To handle rotation and the state sum, construct a path of positive linear maps from the identity to M . Let a be the first column of M and let R be the rotation whose first column is $a/\|a\|$. Then $U = R^{-1}M$ is upper triangular, with first diagonal entry $\|a\| > 0$. Its second diagonal entry is positive because $\det U = \det M > 0$. Choose an angle θ for R . Rotations through angle $s\theta$, $0 \leq s \leq 1$, join the identity to R . Next join R to $M = RU$ through the matrices

$$R((1-s)I + sU), \quad 0 \leq s \leq 1. \quad (405)$$

The matrix inside parentheses is upper triangular; each diagonal entry is a convex combination of 1 and a positive entry of U . Both entries stay positive, so its determinant stays positive. The two paths concatenate continuously. Parametrize their concatenation by M_t , $0 \leq t \leq 1$, and set $f_t(x) = M_t x + tb$. This is a continuous path from the identity map to f , all of whose linear parts have positive determinant.

The preceding genericity argument applies to every $f_t P$. Consequently this is a path in the generic locus, hence also in the regular locus. Lemma 1.20(ii) proves (101). Every carrier is regular by Lemma 3.5(ii). An invertible linear map preserves nonzero directions; if two image directions were negative multiples, applying its inverse would make the original directions negative multiples. Therefore the whole path $f_t Q$ remains regular even though a carrier need not satisfy the parent's all-triple genericity condition. Applying Lemma 1.20(ii) to this carrier path proves preservation of its rotation.

Finally the generic path $f_t P$ lies in a single chamber. The chamber constancy exported by Proposition 4.2 gives $C(fP) = C(P)$. All constructions respect the inherited cyclic order, proving the last assertion. $\square$

Proof of Lemma 8.3

Proof of 8.3. By Lemma 8.2, we may use the orientation-preserving affine coordinates of Definition 8.1, followed by a translation, so that $\mu_2 = (0, 0)$ and $\mu_1 = (L, 0)$. The crossing data and rotation in the conclusion are preserved.

Along the return chain the first coordinate increases strictly from 0 to L . Consecutive edge projections share only their common endpoint, and nonconsecutive projections are disjoint. The chain is embedded and is the graph of a continuous piecewise-linear function f on $[0, L]$. Every crossing of the polygon must therefore involve E_1 . By (G1), no interior chain vertex lies on the line of E_1 . An interior intersection with that line is transverse: a parallel intersecting chain edge would lie on the line, contradicting the same vertex exclusion. Its first coordinate lies in $(0, L)$, so every such intersection is an actual crossing of E_1 . A chain edge has at most one such intersection.

The chain meets these crossings with increasing first coordinate, while E_1 meets them with decreasing first coordinate. This proves the reverse order and the stated Gauss word. For two distinct labels the two visit pairs are nested, not alternating, so no pair interlaces. Every crossing subset is therefore a decomposition by Definition 3.1.

For $i = 2, \dots, n$, take the direction angle $\theta_i = \arg(\tilde{\lambda}_i) \in (-\pi/2, \pi/2)$. At an interior chain vertex the principal turn is exactly $\theta_i - \theta_{i-1}$, because this difference is in $(-\pi, \pi)$. Their sum telescopes to $\theta_n - \theta_2 \in (-\pi, \pi)$. The two remaining principal turns, at μ_1 and μ_2 , each lie in $(-\pi, \pi)$. Thus total turning is in $(-3\pi, 3\pi)$. Lemma 1.20(i) makes it an integer multiple of 2π , giving $\text{rot}(P) \in \{-1, 0, 1\}$. $\square$

Proof of Proposition 8.4

Proof of 8.4. Lemma 8.3 makes every crossing subset a decomposition; Lemma 3.5(i) gives exactly $k + 1$ carriers for $S_{\max} = X(P)$.

By Lemma 8.2, normalize the coordinates of Definition 8.1 so that $\mu_2 = (0, 0)$ and $\mu_1 = (L, 0)$. This preserves C , $\ell(P)$, the crossing count and the uniformity of all corresponding carriers. Strict increase of the first coordinate makes the return chain a graph $y = f(x)$ on $[0, L]$, with $f(0) = f(L) = 0$. Every interior meeting with the horizontal axis lies in the interior of E_1 . By (G1) it is not a chain vertex and is transverse: a parallel intersecting edge would lie on that line and violate (G1). Thus the interior zeros of f are exactly the crossings of the polygon. Write their first coordinates as $x_1 < \dots < x_k$, and put $x_0 = 0$, $x_{k+1} = L$. On each interval (x_a, x_{a+1}) , continuity and absence of zeros give a strictly constant sign for f . Adjacent intervals have opposite signs because their common crossing is transverse.

Consider a chain excursion between consecutive zeros, including 0 and L . If it lies above the axis, its first directed segment has positive vertical component and its last has negative vertical component. Their angles therefore obey $\theta_{\text{first}} > 0 > \theta_{\text{last}}$. All segment angles lie in $(-\pi/2, \pi/2)$, so the interior principal turns telescope to $\theta_{\text{last}} - \theta_{\text{first}} < 0$. At least one interior original vertex consequently has negative turn. Below the axis the inequalities reverse: the first angle is negative, the last positive, and an interior original vertex has positive turn. In particular an excursion cannot be just one segment. Horizontal segments away from the axis cause no problem for this telescoping argument.

Let $S \subsetneq X(P)$ and choose $y \in X(P) \setminus S$. No crossing interlaces an element of S , so Lemma 3.5(iii) assigns both visits of y to one carrier Q . The two consecutive chain excursions meeting at y have no crossing, hence no smoothing site, in their interiors. Since y is unselected, the oriented reconnections of Definition 3.2 leave both excursion interiors on Q . Their interior original vertices retain their parent turns by Lemma 3.5(ii). The two excursions lie in opposite

half-planes, so they supply one positive and one negative turn on Q . It is mixed, and S is absent from the sum in Definition 3.28.

For $S_{\max}$, smoothing every crossing produces the chain excursion from x_a to x_{a+1} closed by the oppositely directed subsegment of E_1 , for each $a = 0, \dots, k$. Each excursion is embedded and meets that subsegment only at its endpoints, so each resulting carrier is embedded. If they are all uniform, Lemma 3.70(i) gives coefficient 1 for each; Definition 3.28 then leaves $(-1)^{\ell(P)+k}$. If any is mixed, $S_{\max}$ is absent too. These are exactly the two asserted cases. $\square$

Proofs for Section 9

Proof of Lemma 9.1

Proof of 9.1. The map $w \mapsto (w_{B_1}, w_{B_2})$ from E_I onto $\mathbb{R}^2 \times \mathbb{R}^2$ is surjective: prescribe vectors u, v on one label in B_1, B_2 respectively, put $-u - v$ on one label in B_3 , and put zero on the remaining labels. Complete these four quotient coordinates to a linear coordinate system on E_I . In them $f_\pi = ad - bc$. Over any field of characteristic zero this determinant is irreducible: its degree in a is one, so in any factorization one factor is independent of a and must divide both its coefficient d and its constant term $-bc$. Those are relatively prime in the polynomial ring in b, c, d , so that factor is a unit. The same argument works after adjoining the unused coordinates and over $\mathbb{C}$. Polynomial rings over fields are unique-factorization domains; thus this irreducible polynomial is prime as well.

Here is a direct way to recover the partition from its polynomial. Let $V = \{x \in \mathbb{R}^I : \sum_i x_i = 0\}$, and let $\alpha_S(x) = \sum_{i \in S} x_i$. Regarding the two coordinate rows of w as $x, y \in V$, its polynomial is the alternating form

$$f_\pi(x, y) = \alpha_{B_1}(x)\alpha_{B_2}(y) - \alpha_{B_2}(x)\alpha_{B_1}(y). \quad (406)$$

The functionals $\alpha_{B_1}, \alpha_{B_2}$ are independent by the same three-label construction. Hence the kernel of this form in its first argument is exactly $\ker \alpha_{B_1} \cap \ker \alpha_{B_2}$: one may choose y with either prescribed pair $(1, 0)$ or $(0, 1)$ to prove the reverse inclusion. The annihilator of that kernel is therefore the two-dimensional span of $\alpha_{B_1}, \alpha_{B_2}$. A nonzero scalar multiple of the alternating form recovers the same span. Lifting its coefficients from V to $\mathbb{R}^I$ recovers precisely

$$\text{span}(\mathbf{1}, \mathbf{1}_{B_1}, \mathbf{1}_{B_2}). \quad (407)$$

The additional $\mathbf{1}$ appears because it is the annihilator of V . This is the space of all functions constant on each of B_1, B_2, B_3 . Two labels are in the same block exactly when every function in this space has equal values at those labels: distinct blocks are distinguished by an indicator. Thus the space recovers the unordered partition, proving the nonassociation assertion.

Three distinct cut vertices divide a cyclic word into three nonempty consecutive blocks with vector sums u, v, z , where $u + v + z = 0$. Taking the first vertex as origin, the other two are $u, u + v$; their determinant is $\det(u, u + v) = \det(u, v)$. Changing the triple's order changes only its sign. The blocks recover the cut vertices as the boundaries where the block names change along the fixed cyclic word. Thus different cut triples cannot supply the same partition.

For a nonempty proper S , the map $w \mapsto w_S$ is onto $\mathbb{R}^2$: put u on one label in S , $-u$ on one outside it and zero elsewhere. Its kernel therefore has codimension two. Finally the product of the nonzero f_π and the nonzero polynomials $\|w_S\|^2$ is a nonzero real polynomial in linear coordinates on E_I . A nonzero real polynomial cannot vanish on a nonempty open box, by induction on the number of variables as in the proof of Theorem 1.36. Its nonzero set is dense and nonempty, which is exactly the assertion needed for M_I^* . $\square$

Proof of Lemma 9.2

Proof of 9.2. We first establish the polynomial avoidance used to choose the intermediate point. Use any linear coordinates on $E_I \simeq \mathbb{R}^D$. Let P be an endpoint, and let f be one of its nonconstant prime block polynomials, with $f(P) \neq 0$. With an indeterminate direction V and parameter t , put $F(V, t) = f(P + tV)$. Inverting t gives the invertible coordinate change $Z = P + tV$ between the Laurent polynomial rings in (V, t) and (Z, t) . The element $f(Z)$ is still prime after adjoining and inverting t , since the quotient is the Laurent polynomial ring over the integral domain $\mathbb{R}[Z]/(f)$. Consequently F is prime in that localized ring. A factor of F that becomes a unit there has the

form ct^k . Since $F(V, 0) = f(P)$ is a nonzero scalar, no positive power of t divides F . Thus F is irreducible already in $\mathbb{R}[V, t]$. It is primitive as a polynomial in t , since its constant coefficient is the nonzero scalar $f(P)$. Gauss's lemma makes it irreducible in $\mathbb{R}(V)[t]$. Its degree in t is positive: the top homogeneous part of $f(P + tV)$ is a nonzero polynomial in V times a positive power of t . Characteristic zero now gives

$$\gcd(F, \partial_t F) = 1 \quad \text{in } \mathbb{R}(V)[t]. \quad (408)$$

For distinct nonassociate block polynomials f, g , the same coordinate change shows that $f(P + tV)$ and $g(P + tV)$ are nonassociate primes. A possible unit ct^k relating them is ruled out at $t = 0$ unless $k = 0$, which would make f, g associates. Both are primitive as polynomials in t , so Gauss's lemma excludes an associate relation with a nonconstant rational function of V . They are therefore coprime in $\mathbb{R}(V)[t]$ as well.

The Euclidean algorithm for univariate polynomials over that field gives a Bezout identity for each of these coprime pairs. Clearing the finitely many coefficient denominators produces an identity

$$K_1(V, t)F(V, t) + K_2(V, t)G(V, t) = D_P(V), \quad (409)$$

where K_1, K_2 are polynomials, D_P is a nonzero real polynomial, and G is either $\partial_t F$ or the other block polynomial on the line. If $D_P(V) \neq 0$, the pair has no common root for that direction, since evaluating the identity at such a root would give $0 = D_P(V)$. This supplies finitely many nonzero avoidance polynomials for the endpoint P , and likewise for the other endpoint Q .

For every removed codimension-two linear subspace L , the linear span of P and L has dimension $D - 1$, since $P \notin L$. Choose a nonzero linear functional vanishing on that span, and do the same for Q . Choose R where all these functionals, all $f_\pi(R)$, all the above $D_P(R - P)$ and all $D_Q(R - Q)$ are nonzero. Each factor is a nonzero polynomial in R , so their finite product has a nonzero real value by the polynomial nonvanishing argument already used in Lemma 9.1.

The segment from P to R avoids L : if $(1 - t)P + tR \in L$ for $0 < t \leq 1$, solving for R would put it in the span of P, L , contrary to the chosen functional. The same argument treats every L and the segment from Q to R . Both segments therefore lie in M_I . On either segment each block polynomial has a nonzero value at both endpoints and is a nonzero univariate polynomial. It has only finitely many roots. The derivative Bezout identities make every root simple, and the pair identities prevent a shared root of distinct blocks. A simple real root changes sign, since the first nonzero term about that root is its nonzero derivative times the parameter increment. Reversing the second segment yields the stated path from P to Q . $\square$

Proof of Lemma 9.4

Proof of 9.4. We give the algebraic argument to specify precisely this broader interface. Cut at the chosen root and use $x < y < z$, $I_* = [x, z]$, $L = [x, y]$, $R = [y, z]$, and the signs $\delta, \epsilon_L, \epsilon_R$ of (9) and (10). Every other point sign is fixed on both sufficiently close sides; neither gap contains the entire critical triple. Its arrays and closed (G1) coefficients are therefore fixed. Put $c = F_H^{-1}(E)$ using Lemma 2.9. Triangular induction first shows that c_I is unchanged unless I contains I_* . The only changing far coefficient of F_H belongs to coordinate $[x, z]$ and cut y .

Split a composition cutting at y into its two gap compositions. This is a bijection, with inverse concatenation of the cut lists. For another left cut r , collinearity gives $H(x, r, z) = \epsilon_L H(x, r, y)$: the two endpoint vectors along the central line differ by a scalar with sign ϵ_L . For a right cut it gives $H(x, r, z) = \epsilon_R H(y, r, z)$ by the same determinant comparison based at a_z . All determinants compared here are nonzero, so these are also the signs on both close sides. The changing gate of F_H has difference $-\delta$. Summing over the independent gap compositions therefore gives nonunary

difference $-\delta\mathcal{U}_L\mathcal{U}_R$, with exactly the $\mathcal{U}$ convention in (11). The unary coefficient is one and the target E_{xz} is zero. Hence

$$c_{xz}^+ - c_{xz}^- = s_*, \quad s_* = \delta\mathcal{U}_L\mathcal{U}_R. \quad (410)$$

In F_{-H} the changing gate instead has difference $+\delta$; reversing the other far gates exchanges each $\mathcal{E}$ with $\mathcal{B}$. Its direct nonunary difference is thus $+\delta\mathcal{V}_L\mathcal{V}_R$.

For completeness, this source propagates linearly even when other inverse coordinates have changed. In a composition of any larger interval at most one child can contain I_* , since children have disjoint interiors. If one does, the difference of the product is its difference times the unchanged other factors; otherwise it is zero. Contract I_* to a distinguished formal leaf e_* . Compositions with a child containing I_* correspond bijectively to compositions of the contracted interval: erase its interior vertices within that child; conversely expand the unique child containing e_* . Every top cut and endpoint survives, so the far coefficients agree. All other child inverse coordinates agree with those on the contracted word by triangular uniqueness. Induction in interval length in $F_H(c) = E$ now gives

$$c_I^+ - c_I^- = s_* c_{I/e_*}^Q \quad (I \supseteq I_*). \quad (411)$$

The base is (410), since the contracted leaf has inverse value one. For a strictly larger interval the contracted interval has at least two leaves. Its target E is therefore zero, and moving its nonunary sum to the other side gives the displayed induction with the indicated positive factor s_* .

If I_* is proper, apply the same bijection including unary compositions to the full F_{-H} output. Its top coefficients are unchanged, and the difference is $s_* A(Q)$, giving (12). At least one other leaf survives, so Q has at least three vertices. Its vertices omit the critical middle vertex and hence have (G1). If I_* is full, its unary difference is s_* and its direct nonunary difference was calculated above; adding them gives (13). Every step reads only point signs and finite transforms. A crossing concurrence, even at the same central parameter, changes none of the inputs just used. $\square$

Proof of Lemma 9.5

Proof of 9.5. Lemma 9.1 shows that an active word has just one zero point triple, whose cuts are the block boundaries. Distinct vertices are guaranteed by the nonzero proper subset sums. An inactive word has no zero point determinant, so its coefficient is unchanged by Proposition 2.7.

For two singleton blocks, the remaining block has at least two letters. Choose a root inside that remaining block. The critical boundary order is the beginning of u , the beginning of v , and the end of v ; both critical gaps are singleton leaves and the span is proper. A full span would leave only the root outside the two leaves, contradicting the length of the remaining block. Both $\mathcal{U}$ factors are one, regardless of the epsilon signs, by Lemma 2.9(iii). The reverse order in the far sign makes $H = -\text{sgn det}(w_u, w_v)$, so Lemma 9.4 gives the first formula. Contraction replaces precisely u, v by $u + v$. The determinant reverses under their interchange.

For one singleton, root at d . The boundary word traverses U, V , so the critical span is full and both gaps have at least two leaves. Write the endpoints of d as p, p' , and the intermediate cut as m , so the boundary order is p', m, p . On the central line take $p = 0, p' = 1, m = t$, with $t \neq 0, 1$. The two epsilons are $1 - t$ and t in sign. When $0 < t < 1$, both are positive; both $\mathcal{U}$ factors are zero, and both $\mathcal{V}$ factors are the respective half amplitudes. Their closing vectors are $p' - m = -w_U$ and $m - p = -w_V$, the asserted positive portions of d . When $t < 0$ or $t > 1$, exactly one epsilon is negative; each of the two products in the full-span formula contains a nontrivial zero $\mathcal{E}$ factor. The jump is then zero. Finally the far sign is $-\text{sgn det}(w_d, w_U)$, while closure gives $\text{det}(w_d, w_V) = -\text{det}(w_d, w_U)$, proving the reversal assertion.

With no singleton blocks, in any root cut both critical gaps have at least two leaves and the critical span is proper. Indeed its complementary cyclic block has at least two edges, and hence cannot be just the root. At least one epsilon is positive because the two nonzero line displacements cannot both oppose their sum. Its $\mathcal{U}$ factor is a nontrivial $\mathcal{E} = 0$; the proper-span response vanishes. Every lower closed word is, up to translation, a vertex subset of the corresponding central word omitting at least one vertex of its sole zero triple. Thus it has (G1). The displayed block lengths give the stated lower arities. Lemma 9.3 permits every root choice made above, including when (G2) fails. $\square$

Proof of Lemma 9.6

Proof of 9.6. Fix the unordered partition of the wall. Words in which it is not consecutive are unchanged. We group the others using Lemma 9.5. Its signs depend only on the ordered block sums, not on the internal orders of the blocks. Every lower polygon in the following groups satisfies (G1) by that lemma. We refer to membership in X or Y as the two colors.

Two singleton blocks. If both singleton letters belong to the same color, they can be adjacent in an active word only in that color's prescribed order and only if consecutive in its source list. When this condition fails there is no contribution. When it holds, merging the two letters gives a bijection from the active words to all shuffles of the shortened color list and the other color list: expand the merged letter back to the prescribed adjacent pair for its inverse. Both lists are nonempty, the coorientation is common, and the sum of the jumps is the fixed coorientation times this lower shuffle sum. It vanishes by the hypothesis.

If the two singleton letters have different colors and neither is c , interchange the adjacent letters. This preserves both color orders, pairs distinct active words, and is its own inverse. The merged word is the same in the pair and the coorientation is opposite. Each pair cancels.

If one singleton is c , call the other $x \in X$; the case $x \in Y$ is identical with the colors exchanged. Adjacency cyclically to c places x at the first or last position of the linear shuffle. It must consequently be the first or last letter of X , respectively. When $|X| \geq 2$, at most one of those two possibilities is available for this fixed x . If neither is available, x is strictly interior to X , the active family is empty, and its contribution is zero. Otherwise x is the available endpoint of X . Merging x into c then bijects the active family with $\text{Sh}(X \setminus x, Y)$, rooted at the merged closing edge: remove that endpoint for the forward map and insert it at the corresponding first or last position for the inverse. Both lists remain nonempty; the lower sum vanishes. When $X = (x)$, both possibilities occur and leave the same word Y with that merged closing edge. Their adjacent orders are c, x and x, c , so their coorientations are opposite and they cancel. These arguments apply to both flat and cusp spatial orders, since Lemma 9.5(1) gives their same signed merger rule.

One singleton block, equal to c . Call the other two block sets U, V . Inactive exterior contacts have zero individual jump. At an interior contact, in the orientation cUV membership in the family is equivalent to each color's U -restriction being a prefix and its V -restriction the complementary suffix. If the prefix conditions fail, there are no words in this orientation. If they hold, all independent choices

$$u \in \text{Sh}(X_U, Y_U), \quad v \in \text{Sh}(X_V, Y_V) \quad (412)$$

occur exactly once. Indeed concatenation preserves each color order under precisely the prefix conditions, and restriction is its inverse. Here subscripts mean ordered restriction to the indicated block. Thus the sum of products in (106) factors into the two sums on the two half polygons, each closed by its fixed positive fragment of c . If either block contains both colors, that factor is a nontrivial lower shuffle sum and vanishes by hypothesis. If neither does, the two nonempty blocks must be all of X and all of Y . Both cyclic placements cUV and cVU are admissible; the internal

color words are fixed and their half products agree. Their opposite coorientations cancel. This also exhausts the other orientation in cases where its prefix conditions hold.

One singleton block belonging to a color. Suppose it is $d \in X$, and let U be the block containing c , V the other block. Write its ordered word $u = a c b$; the words a, b may be empty. For a word v on V , the two cyclic placements duv and dvu , cut immediately after c , have linear words $bvda$ and $bdva$, respectively. Therefore their exact color constraints are, respectively,

$$X = b_X v_X d a_X, \quad Y = b_Y v_Y a_Y, \quad (413)$$

$$X = b_X d v_X a_X, \quad Y = b_Y v_Y a_Y. \quad (414)$$

Equality here is equality of ordered occurrence lists. These formulas include all restrictions imposed by the position of c .

If V contains both colors, its restrictions $v_X = X_V$, $v_Y = Y_V$ are fixed. For either orientation and every admissible fixed u , all and only the words $v \in \text{Sh}(X_V, Y_V)$ satisfy the displayed constraints: the constraints read the restrictions of v , not their interleaving. Thus each fixed u gives a constant half coefficient times a complete mixed shuffle sum on the other half. That half has its fixed closing fragment of d , and the lower induction makes its sum zero. No completeness assertion about the U -half is used.

If V contains only X letters, then $v = X_V$ is fixed. The first orientation is possible exactly when X_V is a consecutive block immediately before d in X , and the second exactly when it is immediately after d . To see necessity, read the X equations above; for sufficiency, delete that adjacent block from the relevant equation. The equations that remain in either case are $X \setminus V = b_X d a_X$ and $Y = b_Y a_Y$. They describe precisely all shuffles of $X \setminus V$ with Y , with c fixed as closing edge. Replace the displayed d by the closing fragment of d on the U -half. Deleting v , or inserting it adjacent to d in the specified orientation, gives mutually inverse maps between this complete lower family and the active words. Its first color is nonempty because it contains the fragment of d ; its second is nonempty because all of Y lies in U . The V -half coefficient is fixed, so the mixed lower sum on the U -half kills the product.

If V contains only Y letters, $v = Y_V$ is fixed and the two constraints become identical:

$$X = b_X d a_X, \quad Y = b_Y Y_V a_Y. \quad (415)$$

For every admissible u , both placements therefore occur or neither does. Pair them retaining the same u and v . Their two oriented half words and product agree, and their coorientations are opposite. They cancel even when U contains both colors and its admissible family is constrained by the position of c . The case $d \in Y$ follows by exchanging the names of the colors in these identical equations. All one-singleton cases are exhausted.

Finally, a partition with no singleton block has zero jump for each word by Lemma 9.5(3). These groups form a partition of the full finite family, so the total jump is zero. $\square$

Proof of Theorem 9.8

Proof of 9.8. We use strong induction on $n = |X| + |Y| + 1$. For $n = 3$ the two words are (x, y, c) and (y, x, c) . Their turn signs are opposite because $\det(y, x) = -\det(x, y) \neq 0$. Lemma 6.1(i) gives minus that turn sign for either triangle on this full (G1) domain. Their sum is zero.

Let $n \geq 4$ and assume the theorem on the full stated (G1) domain in all smaller arities. First restrict the current edge data to M_I^* . Lemma 9.2 joins any two such data by a path with finitely many simple block events. All point signs of all family words are fixed between events, and their amplitudes are therefore fixed by Proposition 2.7. At each event Lemma 9.6, using precisely the

lower induction hypothesis, gives zero total jump. It follows that $S_c(X, Y)$ has one constant value throughout M_I^* . It remains to determine it.

Write $X = qX'$, using q first as the label of its first occurrence. Choose closed hard data on $J = I \setminus \{q\}$ for which every proper subset sum is nonzero and every cyclic ordering is generic. Here $|J| = n - 1 \geq 3$. Such data exist by a finite polynomial avoidance, which we spell out to include (G2) of the hard parents. For each ordering, taking cumulative sums identifies its closed edge space linearly with vertex space anchored at one vertex. The nonzero point and concurrence polynomials from Lemma 1.35 pull back to nonzero polynomials on this space. Indeed they are translation invariant; a polynomial vanishing on the anchored vertex space would then vanish on every translate, contrary to its nonzeroness. For each nonempty proper hard subset, its sum map is surjective by the construction in Lemma 9.1, so the square norm of its sum is another nonzero polynomial. There are finitely many orderings and subsets. The product of these nonzero polynomials has a nonzero real value by the same polynomial nonvanishing argument. Choose the hard data at such a value, and denote its closing edge by c_0 .

Choose the vector assigned to q outside the finitely many lines parallel to these nonzero proper hard subset sums. Their complement is nonempty: the product of their nonzero linear determinant functionals is a nonzero polynomial in q . This choice is admissible at every insertion gap of every hard ordering, because every hard chord is a proper consecutive hard sum. Now use the common closed family

$$w_q(\varepsilon) = \varepsilon q, \quad w_c(\varepsilon) = c_0 - \varepsilon q, \quad (416)$$

with all other hard vectors fixed.

These data lie in M_I^* for all sufficiently small positive ε . For a proper subset whose hard part is nonempty and proper in J , its sum has a nonzero hard limit. The only other nonempty proper cases are the soft singleton and all hard labels; their sums are εq and $-\varepsilon q$, respectively. For a tripartition whose three hard parts are nonempty, the block determinant has the nonzero hard value from the chosen genericity at $\varepsilon = 0$: any hard tripartition can be made consecutive blocks of a cyclic ordering, so its determinant is among those made nonzero by the choice of hard data. Otherwise its sole empty hard part is exactly the soft singleton. The other two hard parts are nonempty and proper; its block determinant, up to sign, has leading term $\varepsilon \det(q, w_B)$, where B denotes either of those two hard parts (the third block sum is $-\varepsilon q - w_B$, giving the same determinant up to sign), nonzero by the chosen direction of q . Continuity and finiteness give a common tail on which every required sum and block polynomial is nonzero.

Fix a hard word $\tau \in \text{Sh}(X', Y)$ and insert q in every gap allowed by the original order $X = qX'$. Formula (416) redistributes compensation to the same nonsoft occurrence c in every word. Apply Lemma 9.7 in each gap, with root c . If X' is nonempty, let x_2 be its first letter. The allowed gaps start immediately after c_0 and end immediately before x_2 . If the letters before x_2 in τ are $e_1, \dots, e_k$, put $e_0 = c_0$ and $e_{k+1} = x_2$. The sum of the multipliers is

$$\frac{1}{2} \sum_{j=0}^k (h_q(e_{j+1}) - h_q(e_j)) = \frac{h_q(x_2) - h_q(c_0)}{2}. \quad (417)$$

When $k = 0$ the sole gap has exactly the right-hand multiplier. All intermediate terms cancel, and the retained endpoint factor is independent of τ . Deleting q and remembering its gap is a bijection from the original shuffles to these hard words with allowed gaps. Summing their soft equalities therefore gives

$$S_{c(\varepsilon)}(qX', Y) = \frac{h_q(x_2) - h_q(c_0)}{2} S_{c_0}(X', Y) = 0. \quad (418)$$

The last equality uses the lower induction hypothesis, since X' and Y are nonempty and all their hard polygons are generic.

If X' is empty, there is one hard word Y, c_0 and every one of its $n - 1$ cyclic gaps is allowed. Summing $(h_q(r) - h_q(l))/2$ around that complete cycle gives zero, since each hard edge occurs once with each sign. Its common hard amplitude is a fixed factor, so the original shuffle sum is again zero. Finitely many words and gaps allow a common positive tail for every soft equality used above.

We have evaluated the constant value on M_I^* as zero. Finally Lemma 9.1 makes M_I^* dense in the entire closed edge space. At arbitrary data in the stated theorem domain, Lemma 9.3 transfers the identity from a nearby point of M_I^* preserving every required family chirotope. This proves the full-domain induction step and the theorem. $\square$

Proof of Lemma 9.9

Proof of 9.9. We first prove an exact word identity, before taking any quotient. For a nonempty word $X = x_1 \cdots x_p$,

$$\sum_{X=UV} (-1)^{|U|} U^{\text{rev}} \amalg V = 0. \quad (419)$$

Splits include empty prefixes and suffixes. A term is specified by a split index k and an output word in which labels $1, \dots, k$ appear in descending order and labels $k + 1, \dots, p$ in ascending order. Its sign is $(-1)^k$. Keep the output word fixed and change the split as follows. If $0 < k < p$ and label k precedes $k + 1$, replace k by $k - 1$. The new ascending sequence is valid because k precedes all of $k + 1, \dots, p$; the descending sequence is its old restriction. If $k + 1$ precedes k , replace k by $k + 1$. The new descending sequence is valid because $k + 1$ precedes all of $k, k - 1, \dots, 1$; the ascending sequence is its old restriction. At $k = 0$ replace it by 1, and at $k = p$ by $p - 1$.

These moves are inverse pairs. In the first interior case, label k precedes $k - 1$ by the old descending order, so at the new split $k - 1$ the rule moves upward again. In the second case, $k + 1$ precedes $k + 2$ by the old ascending order, so the new split moves downward again. The boundary cases obey the same rule, including $p = 1$, when the two boundary splits pair directly. Thus this is a sign-reversing involution with no fixed terms, proving (419) coefficient by coefficient.

A second exact identity is

$$X \amalg (iY) = \sum_{X=UV} Ui(V \amalg Y). \quad (420)$$

Indeed, before i in a shuffle of X with iY there can only be a prefix U of X . After i there is an arbitrary shuffle of its remaining suffix V with Y . Extracting that prefix and suffix is inverse to concatenating them as on the right. Shuffle multiplication is associative and commutative on disjoint alphabets: both parenthesizations correspond bijectively to words preserving the order of all three input lists, and exchanging two lists preserves this description. We will use this exact fact.

Prove (111) by induction on $|X|$, simultaneously for all suffixes Y and finite alphabets. For X empty it is an identity. For nonempty X , the left side of (420) is a proper shuffle and hence is zero modulo L . Separate the split $U = X$, V empty. For every other split and for each individual word Z in $V \amalg Y$, apply the induction hypothesis to UiZ . It has shorter prefix U and the same full alphabet. Adding these congruences gives

$$0 \equiv XiY + i \sum_{\substack{X=UV \\ V \neq \emptyset}} (-1)^{|U|} U^{\text{rev}} \amalg (V \amalg Y). \quad (421)$$

By exact associativity, the sum after i is the proper-prefix part of (419), shuffled with Y . The omitted last term of that exact identity is $(-1)^{|X|} X^{\text{rev}}$, so the displayed sum is exactly $-(-1)^{|X|} (X^{\text{rev}} \amalg$

Y). Concatenating i to this exact equality and substituting it in (421) proves the desired congruence. We used congruences for actual full words, and concatenation only on an exact identity; no closure property of L under concatenation was required. $\square$

Part III. Inputs and formalization

11 Assumptions and formalization

The proofs distinguish literature inputs, arguments developed here, and the separate proof of triple-wall invariance. Theorem 4.10 supplies the triple-wall condition used in Theorem 6.11; it is not an additional unproved assumption.

11.1 Literature inputs

HOMFLY–PT and Lickorish–Millett. The inputs are Literature inputs 1, 2, and 3. They provide the HOMFLY–PT invariant, the local Lickorish–Millett function, and its uniqueness [16, Section 1, pp. 111–120]. The coefficient change and identification used here are established in Lemma 3.17 and Theorem 3.18. The formal development represents the local function by an existential witness satisfying its clauses; it does not implement the historical construction algorithm. Its uniqueness interface uses polygonal diagrams and the stated move generators; the interpretation as uniqueness of a link invariant uses the PL Reidemeister identification. Smooth descent is the separate admitted clause of Literature input 1 specified below.

Named planar records. The named-record arguments, including Lemma 3.11, are proved here. Carter and Smale provide geometric context [3, 29]; their cited results do not by themselves supply the full named-record statement. The distinctions between geometric realizations and the formal record model are listed below.

Contact formulas. Literature input 5 supplies the front and positive-pushoff formulas with the orientations and signs fixed there [8, Sections 2.1 and 2.4]. Its formal interface concerns the specified single-component fronts and supplied positive pushoffs. It assumes the existence of numbers satisfying those formulas; identification with the geometric contact invariants is the literature interpretation. It does not construct arbitrary pushoffs.

Finite front reduction. Literature input 4 supplies finite front-word reduction [21, pp. 6–8, Figure 1][28, Lemma 3.2]. The formal interface asserts a finite principal chain on encoded closed oriented words with the stated syntactic counts and terminal data. It does not implement the source’s geometric reduction procedure or its arm bound. The individual-front bound and carrier floor are derived here; no additional blanket Bennequin bound is admitted.

Tableaux and circular stars. The tableau input in 6 uses Wenzl’s explicit matrices, branching rules, and the stated finite diagram facts [14, Lemmas 2.1 and 2.11(a,b)]. The normalization qualifications in 6 and Remark 7.26 remain part of this use. This input belongs to Section 7, outside the formalized route.

Whitney’s formula. Whitney’s formula and Polyak’s basepoint formulation [31, 22] explain the discussions at 1.22 and 1.23; the main proof does not consume them.

Prescribed directions. Mehlhorn–Yap is a historical comparison [18]. The prescribed-direction result used here is proved in Theorem 5.5.

The delta move. Murakami–Nakanishi [20] is cited for the historical comparison following Lemma 4.11. No knot-type conclusion from that comparison is used.

11.2 Proof map

The tree sum satisfies the flat, cusp, vertex–edge, triple, and silent-wall laws and the soft theorem (Theorems 2.12, 2.13, 2.14, 2.15, and 2.24). Transport and anchor values then give Theorem 6.7.

On the knot side, the carrier floor and knot clause (Theorems 3.62 and 3.64) support the coefficient arguments for the flat, vertex–edge, and soft laws (Theorems 4.4, 4.6, and 4.8). The separate R proof and its bridge supply Theorem 4.10. These laws and the triangle values meet the hypotheses of uniqueness, yielding Theorem 6.11 and Corollary 6.12.

11.3 Correspondence with the letter

Letter item	Companion location
Theorem 1	Theorem 6.7 gives the generic-locus core; Corollary 6.9 supplies the letter’s wider domain.
Theorem 2 and its proof	Theorem 3.64, with the geometric bound 3.62 and Corollary 3.65; the carrier domain and turn condition are part of the statements.
Enumeration and the five-vertex atlas	Definition 10.1 and Computation 10.2; Sections 10.1–10.3 give the method, certificate scope, and atlas in Figure 13.
Polygonal Whitney–Graustein	Theorem 5.5: cyclic relabelling and the exceptional labelled (4, 0) case are distinguished.
Silent extension walls	Definition 1.29, Theorem 2.15, and Proposition 4.3; continuation is defined separately in 6.5.
W1 state-sum bookkeeping	Theorem 4.6, including the bigon and sliding cases.
W3 cancellations	Theorem 4.10 cites the separate R proof and bridge; Proposition 4.13 gives the local partial cases.
Star kinematics and gate parity	Proposition 7.1 and equation (349).
$U(1)$ decoupling and Kleiss–Kuijf	Theorems 9.10 and 9.11, with their stated generic-family hypotheses.
Proof, dependencies, and formalization	Theorem 6.11 and Sections 11.2–11.5. The full mapped dependency graph accompanies [35] as <code>blueprint/web/dep_graph_document.html</code> .

11.4 The accompanying Lean development

The formalization [35] follows this route on the generic locus $\mathcal{U}_n$, using labelled representatives with the stated cyclic invariance. Its triple-wall theorem is `Bridge.sm.R`; the unconditional comparison is `SM.thm.comparison Bridge.sm.R`. The assembled theorem `SM.corner_laws_and_soft` contains chamber constancy, the silent, flat, vertex–edge, triple, cusp and empty-cusp laws, the soft theorem, reversal, cyclic invariance, and triangle values, on their stated domains.

Five literature inputs are represented by six admitted constants:

Printed input	Lean constant
1	<code>SM.lit_homfly</code>
1 , descent clause	<code>SM.lit_homfly_descent</code>
2	<code>SM.lp_lm</code>
3	<code>SM.lp_lm_uniqueness</code>
4	<code>SM.ng_finite_word</code>
5	<code>SM.src_contact</code>

The descent constant concerns jointly smooth families of oriented spatial embeddings with regular generic end projections and supplied polygonal readings. It includes the isotopy-extension and reading passage used by that interface. The first HOMFLY constant concerns the formal polygonal diagram and move vocabulary. Smooth descent is thus an admitted literature clause, not a missing premise.

The reported axiom footprint of the assembled theorem consists of these six constants together with `propept`, `Classical.choice`, and `Quot.sound`. The accompanying statement comparisons were performed by AI reviewers, not human reviewers. Kernel checking establishes the formal statements relative to their definitions and trusted base; it does not independently establish their correspondence with every geometric formulation in this document. The following table records the material representation distinctions in addition to the five literature interfaces described above.

Statement	Formal representation and limit
3.9 , 3.21	Records use polygonal diagrams. The scalar equality in 3.21 assumes a supplied named-record isomorphism; constructing it from every listed smooth, spherical, or page-chart operation is separate geometric content.
3.33 , 3.34	Cusps and their signs use analytic derivative criteria. Their geometric normal-form and arm-order interpretation is separate. Rounding independence assumes the polygonal marking included in the rounding structure; the row does not construct a reading from every geometric smoothing.
3.36	Nonsingular deformation means a jointly smooth family on fixed parameter circles. Word representation preserves the listed counts and the rounded named record; it does not assert full geometric equivalence of fronts.
3.52 , 3.54	The formal clean cusp neighborhoods are convex compact discs, and the diagram readings are supplied. Endpoint equality uses the specified smooth family and the admitted smooth-descent clause. The broader printed clean- smoothing formulation is retained.
3.46	The formal statement is the Euclidean coordinate case, including the finite-collection conclusion. Transfer to a general smooth manifold by charts is not a clause of this bundle.
3.57	The bound uses the specified finite regular generic projection and a supplied polygonal reading. This row does not assert existence of every such reading.

3.61	The smooth and polygonal outputs are related at the named-record level, with the stated old-point identities and a closed-disc convention. Pointwise carrying of the newly inserted kink is not asserted. The chosen arc has one u -tangency, obtained by shrinking the isolated arc.
3.69	The magnitude and supporting-vertex sign of rotation are formalized. Identification of that sign with traversal around the bounded complementary region is a separate geometric interpretation.
4.5	The identities concern supplied marked representatives, clean joins, and split unions. Existence and well-definedness of the corresponding geometric class constructions are not conclusions of this row.

11.5 Coverage by chapter

The table describes the selected statements covered by the formalization; it does not assert that every sentence of a chapter is machine-checked. In particular, the continuation from $\mathcal{U}_n$ to the weak locus $\mathcal{W}_n$ uses the additional visible-chamber arguments printed here and is outside the formalized comparison theorem.

Chapter	Coverage
1: polygons and walls	The selected generic-locus definitions, chamber geometry, and wall statements.
2: tree coefficient	The selected tree-sum definitions, wall laws, and soft theorem.
3: corner state sum	The selected coefficient, carrier, and front route to the floor theorem, subject to the interfaces and representation distinctions above.
4: laws for C	The selected corner laws; triple-wall invariance also uses the separate R proof and bridge.
5: fixed-count transport	The selected prescribed-direction, transport, soft-rotation, and anchor statements.
6: uniqueness and comparison	Root independence, uniqueness, comparison, and inherited laws on their formalized domains.
7: circular stars and convex polygons	Not machine-checked by this development. The earlier star definitions and triangle normalization used in the main route are covered.
8: decay	Not machine-checked by this development.
9: linear relations	Not machine-checked by this development.
10: computations	Not machine-checked by this development.

References

- [1] Bennequin, Daniel. Entrelacements et équations de Pfaff. *Astérisque*, vol. 107–108, pp. 87–161, 1983.
- [2] Brini, Andrea and Eynard, Bertrand and Mariño, Marcos. Torus knots and mirror symmetry. *Ann. Henri Poincaré*, vol. 13, pp. 1873–1910, 2012.
- [3] Carter, J. Scott. Classifying immersed curves. *Proc. Amer. Math. Soc.*, vol. 111, pp. 281–287, 1991.

- [4] Chmutov, Sergei and Polyak, Michael. Elementary combinatorics of the HOMFLYPT polynomial. *Int. Math. Res. Not. IMRN*, no. 3, pp. 480–495, 2010.
- [5] DeTurck, Dennis and Gluck, Herman. Linking integrals in the n -sphere. 2008. Retained arXiv:0802.0320v1, 4 February 2008; the cited Euclidean background is on page 3. <https://arxiv.org/abs/0802.0320v1>.
- [6] Dowker, C. H. and Thistlethwaite, M. B.. Classification of knot projections. *Topology Appl.*, vol. 16, pp. 19–31, 1983.
- [7] Dunfield, Nathan M. and Gukov, Sergei and Rasmussen, Jacob. The superpolynomial for knot homologies. *Experiment. Math.*, vol. 15, pp. 129–159, 2006. arXiv:math/0505662; the cited section numbering is that of arXiv v2.
- [8] Etnyre, John B.. Legendrian and transversal knots. *Handbook of Knot Theory*, pp. 105–185, Elsevier, 2005. Section and page locators refer to arXiv:math/0306256v2 (22 November 2004).
- [9] Franks, John and Williams, R. F.. Braids and the Jones polynomial. *Trans. Amer. Math. Soc.*, vol. 303, pp. 97–108, 1987.
- [10] Freyd, Peter and Yetter, David and Hoste, Jim and Lickorish, W. B. R. and Millett, Kenneth and Ocneanu, Adrian. A new polynomial invariant of knots and links. *Bull. Amer. Math. Soc.*, vol. 12, pp. 239–246, 1985.
- [11] Geiges, Hansjörg. An Introduction to Contact Topology. *Cambridge Studies in Advanced Mathematics*, vol. 109, Cambridge University Press, 2008.
- [12] Geiges, Hansjörg. Contact Geometry. 2004. Retained arXiv:math/0307242v2, 24 January 2004; all cited locators refer to this 86-page version. <https://arxiv.org/abs/math/0307242v2>.
- [13] Gorsky, Eugene. q, t -Catalan numbers and knot homology. *Contemp. Math.*, vol. 566, pp. 213–232, 2012. arXiv:1003.0916.
- [14] Hans Wenzl. Hecke algebras of type A_n and subfactors. *Inventiones Mathematicae*, vol. 92, pp. 349–383, 1988.
- [15] Jones, V. F. R.. Hecke algebra representations of braid groups and link polynomials. *Ann. of Math.*, vol. 126, pp. 335–388, 1987.
- [16] Lickorish, W. B. R. and Millett, Kenneth C.. A polynomial invariant of oriented links. *Topology*, vol. 26, pp. 107–141, 1987.
- [17] Lickorish, W. B. Raymond. An Introduction to Knot Theory. *Graduate Texts in Mathematics*, vol. 175, Springer, New York, 1997.
- [18] Mehlhorn, Kurt and Yap, Chee-Keng. Constructive Whitney–Graustein theorem: or how to untangle closed planar curves. *SIAM J. Comput.*, vol. 20, pp. 603–621, 1991.
- [19] Morton, Hugh R.. Seifert circles and knot polynomials. *Math. Proc. Cambridge Philos. Soc.*, vol. 99, pp. 107–109, 1986.
- [20] Murakami, Hitoshi and Nakanishi, Yasutaka. On a certain move generating link-homology. *Math. Ann.*, vol. 284, pp. 75–89, 1989.

- [21] Ng, Lenhard. A skein approach to Bennequin type inequalities. *Int. Math. Res. Not. IMRN*, 2008. article rnn116; arXiv:0709.2141v1, 13 September 2007; the page and figure locators cited in this document were verified on that retained 15-page author version and not on the IMRN printing.
- [22] Polyak, Michael. Invariants of curves and fronts via Gauss diagrams. *Topology*, vol. 37, no. 5, pp. 989–1009, 1998.
- [23] Poonen, Bjorn and Rubinstein, Michael. The number of intersection points made by the diagonals of a regular polygon. *SIAM J. Discrete Math.*, vol. 11, pp. 135–156, 1998.
- [24] Przytycki, Józef H. and Traczyk, Paweł. Invariants of links of Conway type. *Kobe J. Math.*, vol. 4, pp. 115–139, 1987.
- [25] Read, Ronald C. and Rosenstiehl, Pierre. On the Gauss crossing problem. *Combinatorics (Proc. Fifth Hungarian Colloq., Keszthely, 1976), Vol. II, Colloq. Math. Soc. János Bolyai*, vol. 18, pp. 843–876, North-Holland, Amsterdam, 1978.
- [26] Reidemeister, Kurt. Knotentheorie. *Ergebnisse der Mathematik und ihrer Grenzgebiete*, vol. 1, Springer, Berlin, 1932.
- [27] Reidemeister, Kurt. Elementare Begründung der Knotentheorie. *Abh. Math. Sem. Univ. Hamburg*, vol. 5, pp. 24–32, 1927. DOI 10.1007/BF02952507; the original statement and proof of the move theorem.
- [28] Rutherford, Dan. The Bennequin Number, Kauffman Polynomial, and Ruling Invariants of a Legendrian Link: the Fuchs Conjecture and Beyond. Retained arXiv:math/0511097v1, 4 November 2005; PDF typeset 19 October 2018. <https://arxiv.org/abs/math/0511097v1>.
- [29] Smale, Stephen. Diffeomorphisms of the 2-sphere. *Proc. Amer. Math. Soc.*, vol. 10, pp. 621–626, 1959.
- [30] Wenzl, Hans Gerd. Representations of Hecke algebras and subfactors. University of Pennsylvania, 1985. UMI 8603724. Page 32 presents the weights of Jones’s Figure 5.3 with w, z interchanged and the same non-standard $[k]!$; this comparison is not used in the proofs.
- [31] Whitney, Hassler. On regular closed curves in the plane. *Compos. Math.*, vol. 4, pp. 276–284, 1937.
- [32] *The polygon coefficient X_1 satisfies the wall-crossing laws: a distilled proof*. Unpublished manuscript, 2026. Sections 1 (setup and event conventions) and 10 (`ax:R`).
- [33] *RIII attachment—the establishing local warrants*, with the companion proofs of generic common-row transports, the generic selected complementary couple, extreme singleton transports, and the extreme selected empty/full couple. Separate unpublished R proof attachment, 2026.
- [34] *Bridge from the frozen R theorem to SM11 Hypothesis R*. Unpublished manuscript, 2026, Sections 2–3.
- [35] *Corner laws: Lean formalization*. Repository `aguevara22/gluon-knot-lean`, 2026, revision c871a41b15ab. <https://github.com/aguevara22/gluon-knot-lean>.

- [36] J. Richter and B. Sturmfels, *On the topology and geometric construction of oriented matroids and convex polytopes*, Transactions of the American Mathematical Society **325** (1991), 389–412. doi:[10.1090/S0002-9947-1991-0994170-3](https://doi.org/10.1090/S0002-9947-1991-0994170-3).
- [37] O. Aichholzer, F. Aurenhammer and H. Krasser, *Enumerating order types for small point sets with applications*, Order **19** (2002), 265–281. doi:[10.1023/A:1021231927255](https://doi.org/10.1023/A:1021231927255).
- [38] G. Rote, *NumPSLA—An experimental research tool for pseudoline arrangements and order types*, arXiv:2503.02336 (2025). <https://arxiv.org/abs/2503.02336>.