Technical supplement: BAC260919

Freedom and hidden order in six-colourings of the triangular tiling

Aidan Quigley
21 September 2026 · BAC260919 · Preprint supplement

This supplement supplies the finite certificates and supporting arguments used in the article. Its notation and numbered theorems refer to the main article. Coordinates are axial throughout. Every companion-file path is relative to the supplementary archive accompanying this preprint on Zenodo; access to the development repository is not required. S.12 records the exact scope of the Lean proofs and distinguishes them from historical computational evidence. S.13 gives further written deductions, explicitly outside the Lean-verified portion.

S.1 The finite all-120 assignment

The companion data file `evidence/serrated-t8.json` specifies the finite assignment. Its `assignment` object maps keys such as `U,-1,2` to integer colours 0 through 5; `centres` gives the 120 hexagonal tile centres. At every listed centre the six incident faces, in the cyclic order printed in the article, contain all six colours. Reducing each six-letter word to its lexicographically least cyclic rotation gives 120 distinct words. Reflections are not identified.

The canonical assignment hash is

`82b25b7171b42d75cf838fae3a959ee842d2f64160fb01fd22728ee15616c6af`

To obtain this hash, sort the parsed face triples lexicographically as `(orientation,i,j)`, form one line `orientation,i,j=colour` per face, join with a single LF between lines and no terminal newline, encode as UTF-8 and apply SHA-256. This hashes the assignment rather than the JSON file's formatting.

Direct coordinate replay gives 720 faces, 120 distinct tile classes, no equal-colour contact between present faces, and perimeter 134 in exposed tile edges. Separate Lean proofs establish legality at every contact between distinct present faces, the exact tile partition, one occurrence of every cyclic type, the exposed-edge perimeter, and an exhaustive classification of exactly ten blocked missing faces anywhere in the integer lattice. Only one blocked face is needed for Proposition 1:

Missing face	Colour	Present face	Shared vertex
$D(-10, 8)$	0	$D(-10, 7)$	$(-9, 8)$
$D(-10, 8)$	1	$D(-11, 9)$	$(-10, 9)$
$D(-10, 8)$	2	$U(-9, 7)$	$(-9, 8)$
$D(-10, 8)$	3	$U(-10, 8)$	$(-9, 8)$
$D(-10, 8)$	4	$D(-9, 7)$	$(-9, 8)$
$D(-10, 8)$	5	$U(-11, 9)$	$(-10, 9)$

The missing face must receive a colour from 0 to 5, and each choice clashes at the shared vertex with its row's present face. Every whole-plane extension contains that missing face. This is the complete local reason that no extension can preserve the exact assignment. It is not an impossibility statement for every colouring of this shape, a recolouring obstruction, or an infinite seed.

The other four serrated assignments are not needed for the main-text result. The certificate above and the supplied assignment suffice to verify its nonextension.

S.2 The complete switch certificate

The periodic background is $B(U(i, j)) = T_U[j \bmod 12][i \bmod 12]$ and $B(D(i, j)) = T_D[j \bmod 12][i \bmod 12]$. Row index is j, column index i. The two tables and five unwrapped supports below have machine-readable form `evidence/switch-certificate.json`.

Checking the six incident values at each of the 144 vertex residues proves background legality on the entire plane. For each support, the finite certificate checks that all faces have a colour in its pair, that its contact graph is connected, and that every same-pair face at each touched vertex belongs to the support. This last check includes neighbours outside the displayed cell.

Distinct oriented residues across all 112 support faces prove disjointness of *every* period translate. This is a residue argument, not an extrapolation from a finite patch. The general star and marker arguments of Theorem 6 then yield the infinite construction and entropy bound.

T_U	T_D
040203051053	324545405304
513121324121	050302132452
324545405304	412130510131
053021324252	304542435405
121305010314	452131521232
304542435405	130203040510
251315212324	424541531435
030203040501	151324252123
424515314345	030401030502
512324251213	451535424341
230410130520	042023051053
513545424314	251312132421

J_0: pair $\{0,1\}$, 32 faces

U(-2,16)	U(-2,17)	U(0,15)	U(0,16)
U(1,13)	U(1,17)	U(2,12)	U(2,16)
U(3,13)	U(3,15)	U(4,10)	U(4,12)
U(5,10)	U(5,13)	U(6,10)	U(6,12)
D(-2,15)	D(-2,17)	D(-1,14)	D(-1,17)
D(0,13)	D(0,17)	D(1,14)	D(1,15)
D(2,11)	D(2,13)	D(3,10)	D(3,14)
D(4,11)	D(4,13)	D(6,10)	D(6,11)

J_1: pair $\{0,1\}$, 20 faces

U(4,6)	U(4,7)	U(4,8)	U(6,4)
U(6,7)	U(7,4)	U(7,6)	U(7,8)
U(8,4)	U(8,7)	D(4,5)	D(4,8)
D(5,4)	D(5,6)	D(5,8)	D(6,5)
D(6,8)	D(8,4)	D(8,5)	D(8,6)

J_2: pair $\{3,4\}$, 16 faces

D(6,12)	D(6,15)	D(7,11)	D(7,13)
D(7,15)	D(9,11)	D(9,12)	D(9,13)
U(6,13)	U(6,14)	U(6,15)	U(8,11)
U(8,13)	U(8,15)	U(9,11)	U(9,14)

J_3: pair $\{3,4\}$, 16 faces

D(9,3)	D(9,6)	D(10,2)	D(10,4)
D(10,6)	D(12,2)	D(12,3)	D(12,4)

```
U(9,4)      U(9,5)      U(9,6)      U(11,2)
U(11,4)     U(11,6)     U(12,2)     U(12,5)
```

J_4: pair $\{3,4\}$, 28 faces

```
D(11,10)    D(11,12)    D(12,9)     D(12,12)
D(13,8)     D(13,10)    D(15,7)     D(15,8)
D(16,6)     D(16,9)     D(17,5)     D(17,7)
D(19,5)     D(19,6)     U(11,11)    U(11,12)
U(13,10)    U(13,12)    U(14,8)     U(14,11)
U(15,9)     U(15,10)    U(17,7)     U(17,9)
U(18,5)     U(18,8)     U(19,5)     U(19,7)
```

Markers and indexing

The marker faces in the main text use period translates of the printed representatives. For J_0, $U(-2,16)$ is translated by $(12,-12)$ to $U(10,4)$. For J_2, $D(6,12)$ is translated by $(0,-12)$ to $D(6,0)$. The other three markers are already in the canonical cell. Such translations reindex the arbitrary bits; no density is lost.

The quadratic colouring uses the *untranslated* representatives printed above. Its radicands, in order of families J_0 through J_4, are $(2,3,5)$, $(7,11,13)$, $(17,19,23)$, $(29,31,37)$ and $(41,43,47)$. Keeping that convention fixes the exact illustrated member of the family.

The observed sixteen-face switch

The unchanged recurring-switch figure concerns a saved finite occurrence on the pair $\{4,5\}$, with support

```
D(-10,-21) D(-10,-18) D(-9,-22) D(-9,-20)
D(-9,-18)  D(-7,-22)  D(-7,-21) D(-7,-20)
U(-10,-20) U(-10,-19) U(-10,-18) U(-8,-22)
U(-8,-20)  U(-8,-18)  U(-7,-22)  U(-7,-19)
```

The finite witness from which this occurrence was extracted has canonical assignment hash

```
d244f9d43baa85b000e03a9cb1dff9617905dea86389080d19a7fc2bb3abac4b
```

The accompanying computational record verifies connectivity, closure, eight faces of each colour and 24 affected rainbow stars. The entropy proof uses the five printed supports independently of this finite illustration.

For a direct support comparison, translation by $(16,33)$ carries the listed occurrence exactly to J_2 above, face by face. This finite coordinate identity is checked again by the companion program. It concerns geometry, not the preservation of raw colour names: the observed occurrence uses $\{4,5\}$ and J_2 uses $\{3,4\}$.

S.3 Directional coverage and exact controls

One support intersection for every line class

For each of the six line keys in Section 5.1, the tables give one support face in each key class modulo 12. An entry $r : Q(i,j)$ means that the face belongs to J_r. The entries are generated directly from the printed supports, not from a sampled colouring.

Key mod 12	A	B	C
0	0: $U(2,12)$	0: $U(0,15)$	1: $U(4,8)$
1	0: $U(1,13)$	0: $U(1,13)$	1: $U(6,7)$
2	0: $D(-1,14)$	0: $U(2,12)$	0: $U(-2,16)$
3	0: $U(0,15)$	0: $U(3,13)$	0: $U(-2,17)$
4	0: $U(-2,16)$	0: $U(4,10)$	0: $U(0,16)$

Key mod 12	A	B	C
5	0: $U(-2,17)$	0: $U(5,10)$	0: $D(-1,17)$
6	1: $U(4,6)$	0: $U(6,10)$	0: $U(1,17)$
7	1: $U(4,7)$	1: $U(7,4)$	2: $D(6,12)$
8	1: $U(4,8)$	1: $U(8,4)$	2: $U(6,14)$
9	4: $D(12,9)$	2: $D(9,11)$	2: $D(7,13)$
10	0: $U(4,10)$	0: $U(-2,16)$	1: $U(4,6)$
11	0: $D(2,11)$	0: $D(-1,14)$	1: $U(4,7)$

Key mod 12	D	E	F
0	0: $U(5,13)$	0: $U(1,13)$	0: $U(1,17)$
1	0: $U(-2,16)$	0: $D(0,13)$	0: $U(4,10)$
2	0: $U(-2,17)$	0: $U(2,16)$	0: $U(5,10)$
3	0: $D(-2,17)$	0: $U(0,15)$	0: $U(2,12)$
4	0: $U(0,15)$	0: $U(0,16)$	0: $U(1,13)$
5	0: $U(0,16)$	0: $U(5,10)$	0: $U(4,12)$
6	0: $D(1,14)$	0: $U(-2,16)$	0: $U(3,13)$
7	0: $U(4,10)$	0: $U(-2,17)$	0: $U(-2,16)$
8	0: $U(1,17)$	0: $U(4,12)$	0: $U(5,13)$
9	0: $U(2,16)$	0: $D(2,11)$	0: $U(-2,17)$
10	0: $U(3,15)$	0: $U(2,12)$	0: $U(3,15)$
11	0: $U(6,10)$	0: $D(2,13)$	0: $U(2,16)$

The direction vectors (α, β) are A $(1,0)$, B $(0,1)$, C $(1,-1)$, D $(-1,2)$, E $(1,1)$ and F $(2,-1)$. Translating by $12(\alpha,\beta)$ preserves the exact line key. It advances 24 face positions, with sign depending on the reading direction. A whole line in the selected residue class is reached by an initial period translation, since the coefficients of each key have greatest common divisor one.

Thus a fixed support face, repeated along those translates, reads the bits of one family on an affine integer line of component addresses. It is enough to find one such readable subsequence for each geometric line; every raw face on that subsequence has the same background colour.

Equidistribution details

The version of Weyl's polynomial theorem used here says that a real polynomial sequence is equidistributed modulo one if some nonconstant coefficient is irrational. The vector version follows by applying that statement to every nonzero integer character. The same criterion works for multivariate polynomials averaged over expanding integer rectangles. These standard facts concern limiting sequences, not finite numerical samples.

For $Q(t) = \lambda t^2 + \mu t + \nu$ with irrational λ, consider the character $aQ(t) + bQ(t+q)$. If $a + b \neq 0$, its quadratic coefficient is $(a+b)\lambda$. If $a + b = 0$ and $(a,b) \neq (0,0)$, its linear coefficient is $2bq\lambda$. Both are irrational. Hence $(Q(t), Q(t+q))$ is equidistributed on the two-torus. The pairs whose half-interval bits differ occupy two rectangles of combined area $1/2$; their boundary has measure zero. The mismatch set therefore has density $1/2$. Replacing t by $-t$ proves the same assertion on the negative tail.

For the global-period part, an integer character of $(\theta(m,n), \theta(m+p, n+q))$ either retains a nonzero irrational quadratic coefficient or is a nonzero multiple of the difference. The latter has linear coefficients $2p\sqrt{a}+q\sqrt{b}$ and $p\sqrt{b}+2q\sqrt{c}$. Linear independence of $1, \sqrt{a}, \sqrt{b}, \sqrt{c}$ over $\mathbb{Q}$ makes at least one irrational whenever $(p,q) \neq (0,0)$. The same two-torus mismatch argument rules out a period of the control field.

The preserved block field has exactly one period residue modulo 12, namely $(0,0)$. The companion checker tests all 144 residues, so every putative raw period lies in $12\mathbb{Z}^2$. This finite observation is what allows the control-field argument to exclude *all* nonzero lattice periods, rather than only the six distinguished directions.

Exact evaluation of a bit

For integer m, n the coefficients (m^2, mn, n^2) are integers. To enclose $\sqrt{d}$, choose $S = 2^b$ and compute $r = \lfloor \sqrt{dS^2} \rfloor$ using integer square root. Then $r/S \le \sqrt{d} < (r+1)/S$. Multiply these intervals by the three coefficients, reversing lower and upper contributions when a coefficient is negative, and add them. Double the bounds. If their floors agree, the parity of that integer is the bit; otherwise double b and repeat.

At the origin the bit is zero. Elsewhere $2\theta(m, n)$ is not an integer by radical independence, so sufficiently narrow rational bounds fix its floor and the algorithm terminates. This proves exact computability of the pictured colouring at every input coordinate. Testing a finite window is only a check of the rendering implementation.

S.4 Phase conventions and finite replay

The four phase formulas in Section 6 are part of the article's definition. They are not chosen by fitting a particular displayed sample. The rotation of a face through 120 degrees sends

$$U(i, j) \mapsto U(-i - j - 1, i), \qquad D(i, j) \mapsto D(-i - j - 2, i).$$

Applying P_0 after one or two such rotations gives the printed P_1 and P_2. The regular label is $(i - j) \bmod 3$ for U and three plus that label for D. Each reference labelling is bijective on every vertex star. A phase field first selects the face of a specified *raw* colour and only then assigns its coordinate reference label.

For the fixed colour-2 class of the switch construction, direct enumeration over $0 \le i, j < 12$ gives the P_0 counts $(24, 12, 24, 24, 24, 36)$. Their sum is 144 and their normalized squared sum is $13/72$. Because colours 2 and 5 never switch, all eight of their named fields have period $(12, 0)$ and $(0, 12)$ without any restriction on the controls. The histogram is a finite identity; its repetition on the whole plane follows from the fixed background.

The old four matched-window images retain their original bytes and labels. They compare an earlier three-family formula with a saved finite extension, not with C_5. A matched window makes the pictured comparison meaningful at that scale; it does not turn one colour's phase image into a result about all 24 fields. The new C_5 figures are computed using the supports and exact controls in S.2–S.3.

S.5 The local ribbon decoder

Fix colour 0, or any other single raw colour. Its faces partition the original vertices. Apply $A(i, j) = (i + j, j)$ and replace each vertex by the unit square with that anchor. The image of each selected triangular vertex triple is an EN or NE ribbon. The full cover becomes a tiling of all square cells by these ribbons.

Use alphabet symbols recording the ribbon shape and the cell's position 0, 1 or 2 along its monotone path. At $v = (i, j)$, the unique incident selected-colour face determines the output at $A(v)$. The complete table is:

Incident selected face	Shape	Cell position	Ribbon anchor relative to $A(v)$
$U(i, j)$	EN	0	$(0, 0)$
$U(i-1, j)$	EN	1	$(-1, 0)$
$U(i, j-1)$	EN	2	$(-1, -1)$
$D(i-1, j)$	NE	0	$(0, 0)$
$D(i, j-1)$	NE	1	$(0, -1)$
$D(i-1, j-1)$	NE	2	$(-1, -1)$

The anchor column describes the *output ribbon anchor*. The reindexed *input face anchors* required to inspect the vertex star instead form the set $K = \{(0, 0), (-1, 0), (-1, -1), (-2, -1)\}$. These are different notions of anchor. Both follow directly from the face formulas, and the output position specifies the correct ribbon cell in all six cases.

After the integer coordinate change, this decoder has radius at most two in the sup norm. It commutes with every integer translation: translating an input by t translates its output by $A(t)$, and A is a bijection of $\mathbb{Z}^2$. Its image lies in the order-three ribbon shift; surjectivity onto that entire ribbon shift is neither needed nor asserted.

For the strong-irreducibility application, lift two occurring output patterns to input patterns on their decoder neighbourhoods. If their output supports are sufficiently far apart, those enlarged input supports are farther apart than the input's gluing gap, because $\|A^{-1}x\|_\infty \le 2\|x\|_\infty$ and conversely $\|Ax\|_\infty \le 2\|x\|_\infty$. Gluing upstairs and applying the decoder glues the output patterns. Thus a nonempty strongly irreducible subshift of X would yield a nonempty strongly irreducible ribbon subshift.

A stationary finitely dependent law similarly pushes forward to a stationary finitely dependent ribbon law: sufficiently separated output sets inspect separated input sets, and A identifies the full translation groups. Theorem 2.35 and Corollary 8.9 of Chandgotia–Thorat, arXiv:2605.02226v2, exclude the respective ribbon objects. These are the external results used in the article's proof. The formal companion supplies the order-three specialization directly, as follows. Nothing in this argument excludes arbitrary mixing laws or unbounded-radius finitary codings.

S.5.1 The checked order-three obstruction

The formal proof retains the ribbon height weights of [4, Example 2.33] and checks them directly from the actual four possible ribbon shapes. Every edge weight has three integer coordinates. Predecessor and successor consistency gives zero sum around each unit square, so finite rectangles telescope exactly. Each coordinate of a transverse horizontal edge has absolute value at most two. Consequently the height sums along two parallel diagonal paths of length n and separation m differ in each coordinate by at most $4m$, independently of n.

If a nonempty ribbon subshift were strongly irreducible with separation radius r, compactness extends its finite-pattern gluing property to arbitrary separated sets. Glue two complete diagonal cell lines at horizontal separation $m = 3(r+1)$. Their diagonal classes agree modulo three. The rectangle bound then compares the length-n sums of any two configurations by one constant independent of n. Translating a configuration one step east cyclically permutes the three height coordinates. A second, width-one rectangle therefore bounds the differences between the three coordinates of every diagonal sum.

The periodicity ingredient is proved from the same hypotheses. The actual one-dimensional diagonal projection is compact and strongly irreducible. In its factorial language, minimize the nonempty finite set of connectors of a fixed length. This produces a synchronizing word. Connect two copies with each sufficiently large connector length, repeat the resulting block, and use compactness to realize the bi-infinite periodic configuration. Thus every sufficiently large integer is a period of some diagonal restriction, including a positive $p \equiv 1 \pmod 3$. A period here need not be the least period.

For such a restriction, the sum over kp steps is k times its sum over one period. The bounded coordinate differences force the three period coordinates to be equal integers. Every diagonal step contributes twice a single coordinate basis vector, so their total over a period is $2p$. Equality of the three coordinates gives $3 \mid p$, a contradiction. This proves the needed order-three ribbon nonexistence theorem without a further rigidity premise.

For probability laws, Lean proves that the actual topological support of a stationary finitely dependent law is a nonempty strongly irreducible subshift: separated positive-measure cylinders have positive intersection by independence. The geometric decoder preserves these properties and almost-sure legality. Finally, measurable equivariant finite-memory maps are represented by measurable block rules; actual iid inputs give stationary finitely dependent output laws. These arguments complete all three conclusions of Theorem 10. The source attribution is unchanged, and no novelty claim is made for this formal proof.

S.6 Full circle-coding argument

Use the model and axial coordinates of the main article, Section 7.2. Fix a raw colour k and its indicator step functions $u = \mathbf{1}_{f_U=k}$ and $d = \mathbf{1}_{f_D=k}$. At vertex (i,j) the faces are $\mathrm{U}(i,j)$, $\mathrm{D}(i-1,j)$, $\mathrm{U}(i-1,j)$, $\mathrm{D}(i-1,j-1)$, $\mathrm{U}(i,j-1)$ and $\mathrm{D}(i,j-1)$. Let $Xh(x) = h(x-\alpha)$, $Yh(x) = h(x-\beta)$ and put

$$\mathcal{A} = 1 + X + Y, \qquad \mathcal{B} = X + Y + XY.$$

Exact single-colour vertex coverage requires $\mathcal{A}u + \mathcal{B}d = 1$ at every point $\theta + i\alpha + j\beta$ modulo one. When $1, \alpha, \beta$ are $\mathbb{Q}$-independent these points are dense. A disagreement on a constancy interval would be sampled, so the identity holds almost everywhere on the circle. Arbitrary endpoint values do not invalidate this step.

The distributional derivatives Du, Dd are finite integer atomic measures, and $\mathcal{A}Du + \mathcal{B}Dd = 0$. The translations of atoms by α, β act freely: $m\alpha + n\beta = 0$ modulo one forces $m = n = 0$. Partition the finitely many derivative atoms into these orbits. Choosing one basepoint in each orbit identifies its finite atomic measures with Laurent polynomials over $\mathbb{Z}$ in independent variables X, Y. The derivative equation splits into

$$\mathcal{A}U + \mathcal{B}V = 0.$$

Every solution has $U = \mathcal{B}W$, $V = -\mathcal{A}W$ for an integer Laurent polynomial W. Indeed multiply U, V by a common monomial to make them polynomials, then substitute $Y = -1 - X$. Since $\mathcal{B}(X, -1 - X) = -(X^2 + X + 1)$ and $\mathbb{Z}[X]$ is a domain, $V(X, -1 - X) = 0$. Division by the monic polynomial $Y + 1 + X$ gives $V = -\mathcal{A}W$ with integer coefficients; substitution and cancellation give $U = \mathcal{B}W$. Undoing the monomial proves the Laurent statement. This is divisibility, not a claim that $\mathcal{A}, \mathcal{B}$ generate the unit ideal.

Reassembling the finitely many orbits gives a finite integer atomic measure w with $Du = \mathcal{B}w$ and $Dd = -\mathcal{A}w$. A periodic derivative has total mass zero, so $0 = \mathrm{mass}(Du) = 3\,\mathrm{mass}(w)$. An integer atomic measure of zero total mass has an integer-valued periodic step primitive h: start at an integer value and add the integer jumps around the circle. Thus $Dh = w$. It follows that $u - \mathcal{B}h$ and $d + \mathcal{A}h$ have zero derivative, and are constants a, b almost everywhere. Their integer values make $a, b \in \mathbb{Z}$. Commutation gives

$$\mathcal{A}u + \mathcal{B}d = \mathcal{A}(\mathcal{B}h + a) + \mathcal{B}(-\mathcal{A}h + b) = 3(a + b),$$

contradicting the required value one. Even one colour's exact coverage is therefore impossible in the independent case.

In the dependent case a nonzero $p \in \mathbb{Z}^2$ has $p_1\alpha + p_2\beta = 0$ modulo one. Both raw orientations have exact period p, including any endpoint values. Translation by $6p$ also preserves all reference labellers, so every named field is periodic and A fails. This proves Theorem 11 for the specified same-circle model.

Rational slopes can still give legal periodic colourings: at $\alpha = 1/6$, $\beta = 1/3$, $\theta = 0$, take $f_U(x) = \lfloor 6\{x\} \rfloor$ and $f_D(x) = \lfloor 6\{x\} \rfloor + 4$ modulo six. The atomic action now has $X^6 = 1$, $Y = X^2$; its quotient ring does not satisfy the free-orbit argument. Arbitrary measurable partitions need not have finite atomic derivatives. Additional finite coordinate factors, position-dependent state and higher-dimensional torus codings are also outside this proof. No worldwide priority claim is asserted. Lean now verifies both branches of the full specified theorem, including Laurent localization, the free circle action, actual finite atomic jumps, integer primitive and constant recovery, and the raw-to-phase period implication. The formal analytic route uses finite floor expressions and one-sided jumps, with exact control of the finite exceptional set. S.12 describes this correspondence and the checked M3 consequence.

S.7 Companion files and computational evidence

The supplementary archive contains the editable manuscript, data, figure generators, proof sources and reproduction instructions. The PDFs are generated from the supplied Markdown sources using Quarto and Tectonic. `assemble.py` concatenates the ordered section files and inserts tables from the finite certificates. From the root of the extracted archive, the finite figures and manuscript can be reconstructed with:

```
python -B make_artifacts.py
python -B make_revision_artifacts.py
python -B assemble.py
```

The finite checkers verify background stars, support colours, connectivity and closure, disjointness of translates, line coverage, block-period residues, phase histograms and the displayed assignment's six blocker contacts. These checks use exact coordinate arithmetic and small finite matrices and graphs. They verify finite premises; the infinite conclusions require the arguments printed in the article and supplement.

Complete saved assignments, input hashes, verifier outputs and figure manifests are included in the archive. The source material for S.8–S.11 is indexed by `evidence/revision-inputs.json`; the source arguments and AI reviews for S.13 are indexed by `evidence/september21-inputs.json`. Both indexes identify copies included in the archive. Original source snapshots retain their dates and wording, while the manuscript states the corrected mathematical conclusions. The proofs and qualifications printed here do not require access to private development notes.

The full two-ring weighted evaluation, exhaustive survivor scan and native LRAT replays have separately dated computational records. Rebuilding the PDFs or checking their input hashes does not rerun those computations. Supplement S.12 identifies which conclusions also have Lean proofs and which remain outside that formal verification.

S.8 The four-colour inventory criterion

Theorem 2 concerns multisets of cyclic permutations of six named colours. There are 120 types, represented uniquely by putting 0 first and ordering the other five colours. For a word w of four distinct colours, its two cyclic completions t, t' give the equation $m_t + m_{t'} = 2$. Adjacent interchanges connect the completion graph. The cycle

$$012345 \to 023451 \to 034512 \to 045123 \to 051234 \to 012345$$

has five edges, each preserving a consecutive block of four. Thus the equations $y_t + y_{t'} = 0$, where $y_t = m_t - 1$, force every $y_t = 0$. The argument uses ordinary integer multiplicities. Conversely, a fixed word has two completions in the exact inventory. In fact the 360 equations alone imply 120 tiles: each tile contributes six words, so $6 \sum_t m_t = 720$.

For an explicit failure of lower-order balances, start with one copy of every type and make these changes:

Double these types	Omit these types
013452	013542
021354	021345
023145	023154
031542	031452

All labels denote cyclic orders. These are respectively zero-based lexicographic IDs $9, 25, 30, 53$ and $11, 24, 31, 51$. Every directed pair still occurs 24 times and every directed triple six times, but 24 quadruple totals differ from two. The companion checker enumerates the words directly and checks the displayed perturbation, all 360 completion edges, connectivity and the five-cycle. It also applies the criterion to the saved serrated assignment. No geometric realisability is claimed for the perturbed multiset.

The graph argument above proves the result without a numerical rank calculation.

S.9 Seed counts and their exact finite-depth boundary

Geometry and transfer correspondences

Tile centres are $(i, j) = (u + 2v, u - v)$ for $\max(|u|, |v|, |u + v|) \le r$. Each centre carries its six incident faces in the order of Section 1. The centre is fixed to $(0, 1, 2, 3, 4, 5)$; tile types may repeat. These are 42- and 114-face regions at radii one and two, distinct from the vertex-star regions H_r in S.10.

For the neighbour centred at $(1,1)$, write its oriented state as $p = (p_0, \dots, p_5)$. Its precise restrictions from the centre are

$$p_2 \notin \{0,1\}, \qquad p_3 \notin \{0,1,5\}, \qquad p_4 \notin \{0,5\}.$$

There are $3(3 \cdot 3 - 2)3! = 126$ such permutations. Put $\rho(i,j) = (-j, i+j)$ and $(\tau p)_j = p_{j-1} + 1$ modulo six. At centre $c_i = \rho^i(1,1)$, the actual state is $\tau^i p_i$. This simultaneous rotation and colour renaming preserves the central rainbow. Let M_{pq} test every shared-vertex contact between p at c_0 and τq at c_1. The wheel geometry then gives

$$N_1 = \operatorname{tr}(M^6) = 1,060,600.$$

The walk is rooted at a fixed physical position; no division by six is made. In the present revision, a direct raw-vertex implementation reconstructs the 126 states, 1,264 transitions and this integer trace.

For two rings, the outer centres are radial corners $A_i = 2c_i$ and joining tiles $B_i = c_i + c_{i+1}$. A sector state (p,a) records the inner tile and its compatible radial neighbour, giving $126^2 = 15,876$ states before pruning. Define

$$W_{(p,a),(q,c)} = M_{pq} \#\{z : \ z \text{ at } B_0 \text{ is compatible with } p, a, \tau q, \tau c\}.$$

Here z ranges over all 720 oriented rainbow tuples, and compatibility means every point and edge contact. The six joining tiles do not touch one another. Conditional on the six sector states their choices therefore multiply. Every legal two-ring patch gives exactly one such closed sector walk with a joining-tile choice on each edge, and conversely. Hence

$$N_2 = \operatorname{tr}(W^6) = 1,051,975,931,234.$$

The inherited exact evaluation retained 8,121 cyclic sector states and used checked integer arithmetic. An independent raw-geometry builder and different matrix contraction gave the same value. That large evaluation is historical and is not rerun here. The complete definition, pruning justification, overflow bounds and independent evaluation are retained in the copied seed-count source.

The Lean companion now proves $N_2 = \operatorname{tr}(W^6)$ for the literal 114-face assignments, including every contact and uniqueness of the reconstruction. It does **not** yet prove the displayed numerical evaluation of this weighted trace. That equality retains the historical computational evidence just described.

Why the side-test survivors have a second ring

Deleting inner transitions whose joining tile has no compatible state leaves M_+, with 1,136 edges. A direct finite reconstruction gives $\operatorname{tr}(M_+^6) = 599,799$. Thus 460,801 seeds fail a necessary test. Necessity alone would not classify the others.

For each surviving edge $p \to q$, the Boolean relation $R_{pq}(a,a')$ is true when some joining tile is compatible with the two inner states and the two chosen radial corners. A complete second ring exists exactly when

$$\operatorname{tr}_{\mathrm{B}}(R_{p_0p_1} \cdots R_{p_5p_0}) > 0.$$

A positive closed lift supplies the corners, then each true relation independently supplies its joining tile; every second ring determines such a lift. The exhaustive scan evaluated this criterion on every survivor rank $0, \dots, 599798$, finding 599,799 positive and zero negative traces. Its independent verifier regenerated the geometry and all twelve block digests. The copied primary and independent JSON receipts retain that historical attribution. Fresh checks here compare their completion fields and block bindings; they do not

repeat the exhaustive scan. This separates the proved correspondence, the completed finite evaluation and the still-open infinite question.

In the new formal layer, Lean verifies the filtered trace 599,799 and proves that each failed joining test excludes a full second ring. A separate completed run checks actual 114-face completions for all 50,720 entries of a representative list, preserving their 42-face inner seeds; physical rotations and reflections transport these completions. The remaining formal task is to prove that the orbits of this exact certified list exhaust all 599,799 survivors. A complete Lean proof connecting these orbits to the full survivor set is not yet available. Therefore the representative proofs are positive evidence for their stated seeds, not yet a universal Lean classification.

A portable seven-tile obstruction

The seed displayed in the main article is given completely by this table. Each tuple uses the incident-face order of Section 1.

Centre	Tuple
$(0,0)$	012345
$(1,1)$	013245
$(-1,2)$	201345
$(-2,1)$	021345
$(-1,-1)$	012345
$(1,-2)$	012345
$(2,-1)$	013245

It has 42 distinct faces and 31 incident vertices, all with distinct present colours. At $(0,2)$ the present faces $D(-1,1), D(0,1), U(-1,2), U(0,1)$ have colours $5,1,2,3$. The two missing faces $x = D(-1,2)$ and $y = U(0,2)$ must use $0,4$. But x touches the present colour-0 face $D(-2,2)$ at $(-1,3)$, and y touches the present colour-0 face $U(1,1)$ at $(1,2)$. Both would have to be 4 although they share a vertex. This is a deductive nonextension certificate for this exact seed.

The supplementary archive contains the machine-readable classification records and nonextension certificates, indexed as described in S.7. Colour relabelling gives 720 times the fixed-centre counts in fixed coordinates; quotienting by physical rotations gives 120 times those counts, since the central rainbow makes the rotation action free. Reflections remain distinct. None of these operations enforces an all-120 inventory or classifies infinite extension.

S.10 Annulus transfer and certified lookahead traps

Complete local interface

Let H_r be the union of stars centred at axial vertex distance at most r, so $|H_r| = 6(r+1)^2$ faces. Put $n = r+1$, $L = 6n$, and order the boundary vertices $v_0, \dots, v_{L-1}$ cyclically starting at $(n,0)$. A vertex is a corner when its index is divisible by n. Let M_k be its present-colour set and $C_k = \{0, \dots, 5\} \setminus M_k$ its missing colours. Any exterior assignment is compatible with the fixed interior precisely when it has distinct colours on its own contacts and avoids M_k at every shared vertex. This is why the masks are an exact interface.

The next annulus has $2L+6$ faces. Let E_k be its face containing the edge $v_k v_{k+1}$. At a side vertex there is one further face $B_k = F_k$; at a corner there are two, with B_k adjacent to the outer vertex of E_{k-1} and F_k adjacent to that of E_k. On the first side at $v_k = (i,j)$, $i,j \geq 1$, $i+j = n$, the faces are

$$E_{k-1} = D(i, j-1), \quad B_k = F_k = U(i,j), \quad E_k = D(i-1, j).$$

At $(n,0)$ they are respectively $U(n,-1), D(n,-1), U(n,0), D(n-1,0)$. Rotations cover all six sides, including $n = 1$ where every boundary vertex is a corner. The outer vertex of E_k sees E_k, F_k, B_{k+1}; an outer corner sees B_k, F_k. Thus the inner rainbow conditions leave only the extra inequalities $F_k \neq B_{k+1}$ to check.

Retain $q_k = (c(E_k), c(F_k)) \in Q = \{(a, b) : a \neq b\}$, with $|Q| = 30$. The local transition $(a, b) \to (c, d)$ obeys:

Vertex	Exact transition condition
Side	$\{a, c, d\} = C_k$ with three distinct colours, and $d \neq b$
Corner	a, c, d distinct in C_k; its remaining colour t satisfies $t \neq b$

At a side put $t = d$. The recovered local colours are $E_{k-1} = a, B_k = t, F_k = d, E_k = c$. The last inequality checks the outer vertex on the incoming edge. Shared E_k values agree between consecutive transitions, including the cyclic seam. All new contacts are exactly those already listed, so closed paths give legal annuli, and every annulus gives one closed path. This proves the bijection in Theorem 4. The integer trace counts rooted, named-colour assignments without any rotation quotient.

There are 20 possible side masks and 15 corner masks. Their tables have respectively 24 and 96 nonzero entries: $3! \cdot 4$ and $4! \cdot 4$. Propagating 30-bit reachability sets takes $O(L)$ fixed-width operations and constant working memory beyond the input; retaining predecessors constructs a witness with $O(L)$ output space. Integer path counts require $O(L)$ bits, so this is not a linear bit-complexity assertion for exact counting.

Each transition emits the next boundary masks $\{a, b, t\}$ and, at a corner, $\{t, d\}$. Rotate the emitted word left by one symbol to restore the canonical start. Composing the first-layer transition with the one or two second-layer transitions reading those masks retains a pair of states in Q^2. Closing both paths decides two-ring feasibility with at most 900 states. Boolean composition merges witnesses; the implemented two-ring decision is neither a count nor a filling constructor.

The fresh figure check reconstructs the contact graph for the displayed annulus directly from vertices and compares it with this decomposition. The all-radius proof is the coordinate argument above, inherited from the copied ring-transfer source, not the finite figure check.

Finite depth is not closure

For frozen extension define $K_m(P)$ as in Section 3. Restricting an $(m + 1)$-ring completion to its first ring, and conversely viewing an m-ring completion of a successor as an $(m + 1)$-ring completion of its predecessor, proves

$$K_{m+1}(P) \iff \exists Q\,[P \to Q \wedge K_m(Q)].$$

The two recorded H_{12} traps each have 1,014 fixed faces and a compatible H_{14} witness with 1,350 faces. The impossible H_{15} region has 1,536 faces. To encode completion, give each face six Boolean variables with exactly one true; fix all old colours; and forbid each equal-colour pair at every shared vertex. This has a model exactly when a completion exists: one true variable specifies the colour, and the pair constraints enforce all point and edge contacts, including partial boundary stars. No symmetry breaking or selected first-ring continuation restricts the formula.

The saved radius-15 formulas have 9,216 variables and 78,600 clauses. The historical independent replay reconstructed their full clause multisets from raw geometry and checked the LRAT refutations with both `cake_lpr` and `lrat-check`. A proof check establishes the CNF contradiction; the preceding encoding equivalence supplies its geometric meaning. Those historical native checks alone were not an end-to-end proof-assistant formalisation. The separate 21 September Lean proof now establishes the exact geometric traps directly, as described below and in S.12.

Record 0015-1

```
r12_assignment_sha256:
5d85f99f6404d479583a041954f506925606e956e8309158a98a2f52154093ea
r14_assignment_sha256:
```

```
c42d29059130e8dc2a763f9a3197920bb80385034e1842a8ead7ed774bba191c
cnf_sha256:
91f2f1882826657e0ea23edc55ef838f2341403893ea52b3c23ea26bf3b2a56d
lrat_sha256:
dbad8fbc09d757c8bbb6024e3c0c8fc3c6225c1369e99d0f80604ed6002e189e
```

Record 0018-1

```
r12_assignment_sha256:
1069e2ff2d894d4ea27dda3e791d2867d9f740e11723457de3f1ba0f79600ac8
r14_assignment_sha256:
3d225fc9d9ba3b64da3b0f962834d3f8cf240eaf686bf8313b200ca5972c817b
cnf_sha256:
58870cb46ad376ca8ebc01aea329497e38b1f27bd4272cf8b655c958f901ab9b
lrat_sha256:
9cffc99e556876f69de30cda939976932c8c77ca25441671e53f6a0c8227e058
```

For the new Lean proof, an untrusted exporter reads selected LRAT dependency hints and emits ordinary disjunction-elimination proofs. Lean derives each initial constraint from six-colour exhaustiveness, fixed input colours or a raw-coordinate touching pair, and checks every subsequent deduction. The saved two-ring witnesses are also checked against geometric legality and preservation of every old face. Thus the formal conclusion requires neither a trusted SAT parser nor an assumed CNF-to-geometry equivalence. The original native-checker records retain their historical dates.

The main-text implication is universal over successors: if one successor of a trap were in K_2, the trap would be in K_3. Thus *every* successor fails K_2, even though some extend by one further ring. The tree illustration is deliberately schematic and counts no branches.

The companion evidence preserves both frozen inputs, both H_{14} witnesses, both CNFs and LRAT refutations, the native checker logs and the original final replay receipts. Fresh checks in this revision verify the file bindings, raw witness legality and preservation of the frozen interiors; native proof replays remain attributed to the historical receipts. No solver or lookahead campaign is rerun. The selected records are counterexamples, not an exhaustive census or an estimate of a trap probability. They impose no impossibility conclusion on the original 42-face seeds, recolouring, longer lookahead, other boundary families or M3.

S.11 A periodic field with small domains and no finite edits

Formula and whole-plane validity

Let (a_i, b_i) be the period-five sequence $((0,4),(2,5),(0,4),(3,0),(4,3))$. Put $Q = (0\ 2)(1\ 4)(3\ 5)$ and $g = (0\ 4\ 3)(1\ 2\ 5)$. For all signed integers i, k, define

$$C_{\mathrm{fr}}(U(i,2k)) = g^k(a_{i+k}), \qquad C_{\mathrm{fr}}(D(i,2k)) = g^k(b_{i+k}),$$
$$C_{\mathrm{fr}}(U(i,2k+1)) = g^kQ(b_{i+k}), \quad C_{\mathrm{fr}}(D(i,2k+1)) = g^kQ(a_{i+k+1}).$$

The sequence has period five and $g^3 = 1$, so $(30,0)$ and $(0,60)$ preserve every raw colour and every reference phase formula. Checking the 1,800 stars on this rectangle, including wrap contacts, proves legality everywhere by periodicity. No minimum period is claimed. The companion checker reconstructs the formula and every star directly.

Finite quotient versus finite domains

Join nearest-neighbour vertices with equal phase label. Each edge of the rectangular quotient retains its actual displacement across the period cell; loops and parallel edges are kept. A quotient component lifts to finite whole-plane components if and only if every closed walk has zero deck translation. Indeed, in that case integer potentials assign one consistent lift to each quotient vertex, independently of path, so each lifted component

has exactly the quotient size. If some closed walk has nonzero translation h, repeating it reaches infinitely many translates by nh.

Direct reconstruction finds zero winding in every component of all 24 fields. On the 1,800-vertex quotient the per-colour counts are:

Original colour group	Reference	Component size: number
$\{0,3,4\}$	P_0	3: 120; 6: 60; 18: 60
$\{0,3,4\}$	P_1	6: 120; 9: 120
$\{1,2,5\}$	P_0	6: 120; 9: 120
$\{1,2,5\}$	P_1	3: 120; 6: 60; 18: 60
All six	P_2	3: 120; 6: 60; 18: 60
All six	R	3: 360; 6: 120

Each row applies separately to each colour in the stated group. The histogram masses sum to 1,800. The quotient proof converts this finite certificate to the bound 18 on the plane. It is not a boundary-cut sample estimate. The figure includes whole components intersecting its marked observation window, so the highlighted component sizes are also exact.

Row determination and rigidity

For a candidate lower row write $u_i = C(U(i,j))$, $d_i = C(D(i,j))$. A fixed upper row determines

$$T_i = \{u_i, d_{i-1}, d_i\}$$

as the complement of its three faces at vertex $(i, j+1)$. Also u_i, d_i, u_{i+1} must be distinct since those faces meet at $(i+1, j)$. The two canonical parities are:

Row parity	Pairs (u_i, d_i), $i = 0, \dots, 4$	Complement sets T_i
0	04, 25, 04, 30, 43	034, 245, 045, 034, 034
1	10, 32, 15, 21, 52	012, 023, 125, 125, 125

For each parity construct states (i, a, b), with i modulo five and distinct $a, b \in T_i$. Put an edge to $(i+1, c, d)$ exactly when

$$\{b, c, d\} = T_{i+1}, \qquad |\{a, b, c\}| = 3.$$

Every legal bi-infinite row gives a bi-infinite path in this graph. Simultaneously deleting vertices with no incoming or no outgoing edge cannot delete any occurrence on such a path. Both graphs have 30 states and 28 edges. The successive surviving state counts are $16, 10, 7, 5$; the survivors form the canonical five-cycle, with one state at each column residue. The small graph can be reconstructed completely from the displayed table and rule. Its verification must prune sources and sinks, not merely discard noncyclic strongly connected components, since bi-infinite paths can otherwise have transient connectors.

For $C_j(i) = (C_{\mathrm{fr}}(U(i,j)), C_{\mathrm{fr}}(D(i,j)))$, the formula gives

$$C_{2k+\varepsilon}(i) = g^k C_\varepsilon(i+k), \qquad \varepsilon \in \{0,1\}, \quad k \in \mathbb{Z}.$$

A column shift and colour permutation preserve the graph constraints, so the two calculations cover every signed height. Hence a complete canonical row forces the entire row below, then every lower row. A nonempty finite set of changed faces has a highest changed row. The row above is canonical and forces that row to be canonical too, a contradiction. This proves finite-support rigidity even when all six colours may change.

The component bound implies the M3 tail clause T, but the common period makes the disagreement at $(30, 0)$ zero in every field, violating A. This control therefore fails M3. The rigidity conclusion concerns this background only; infinitely supported changes and other starting fields are outside it.

Sources: the copied periodic phase-control and rigidity notes of 9 September, their original JSON certificates, and the periodic-control and rigidity proofs included in the source archive. The finite period data and graph-pruning certificates can be checked by the companion program. The infinite consequences follow from the periodic lift and row arguments just given.

S.12 Lean verification: correspondence and limits

The verification date is 21 September 2026. **Complete** means that the formal statement proves the entire numbered assertion for the mathematical objects and hypotheses printed here. **Partial** means that actual ingredients have been proved but some obligations remain. It does not mean that a final theorem was accepted with the missing argument supplied as an axiom. The eleven numbered results comprise eight theorems and three propositions; definitions, adjacent computations, external theorems and assertions inside proofs create further obligations beyond that count.

S.12.1 The ten complete numbered results

Proposition 1. The formal assignment contains all 720 saved face-colour entries. Lean checks the six displayed blockers of $D(-10, 8)$ against their actual vertices and colours. Six-colour exhaustiveness then excludes every whole-plane no-touch extension, and the proved incidence bridge gives the rainbow formulation. The quantified conclusion concerns preservation of this exact assignment.

Separate supporting proofs establish perimeter 134 by counting distinct actual exposed tile edges; they prove that the 120 tile stars partition the 720 entries and contain every rotation-only cyclic type exactly once. For missing faces, a complete contact certificate reduces the unrestricted integer lattice to a 264-face collar. Checked available colours exclude every unlisted collar face, while the ten listed faces have all six blockers. This proves exactly ten blocked missing faces globally.

A further Lean proof establishes local legality of the entire saved assignment: every two distinct present faces sharing any lattice vertex have different colours. For each saved face, a checked contact row contains the complete six-face stars at all three vertices. Colour checks on those rows, together with proved collar absence and uniqueness of each saved face's colour, establish the unrestricted pairwise statement. This includes contacts between different tiles and at boundary vertices. The separate receipt freshly rebuilds the supporting geometry and counting proofs as well as this legality proof; it does not alter the earlier receipts or claim an extension of the finite assignment.

Theorem 2. The formal cyclic type is a six-colour order modulo rotation, with reflection distinguished. Lean proves completeness and uniqueness of the 120 canonical representatives, evaluates actual consecutive cyclic occurrences, and verifies both completions of every distinct-colour four-word. A connected spanning tree and an odd cycle force every natural-valued multiplicity to one. No total-size hypothesis is assumed. The reverse implication is also proved. Additional proofs verify the eight changed types of S.8, total size 120, pair counts 24, triple counts six, exactly 24 imbalanced four-word totals, and the size-120 consequence of four-word balance. These are multiset statements, without a placement or extension conclusion.

Theorem 4. The formal regions are the actual unions of vertex stars, and legality tests all raw face pairs sharing any vertex. For every radius, including zero, Lean classifies the old, boundary and new contacts, proves complete injective enumeration of the new faces, and reconstructs precisely the legal annuli from closed 30-state paths. The proof includes both cyclic seams and the outer point contacts; checking only completed boundary stars would be insufficient.

An exact bijection identifies geometric assignment cardinality with the natural-number trace of the ordered transfer product. Executable Boolean code stores fixed-size transition and predecessor tables. Its actual returned state list is converted to a raw face-colour witness preserving the interior. Counting fixed-width probes, table writes, stored lookups, recovery scans and output writes gives a bound of $29711L + 36$, where $L = 6(r + 1)$; materialising each site's table adds a fixed cost per site. The bound concerns the specified

operation model, not arbitrary-precision integer arithmetic or a runtime benchmark. Exact next-mask emission and the shifted incoming seam give the two-ring decision with state type of cardinality 900. Its Boolean positivity is not asserted to count two-ring assignments. The separate constant-working-memory decision implementation mentioned in S.10 is not the proved memory bound of this witness-producing solver.

Proposition 5. Both saved radius-12 inputs are represented on the exact geometric H_{12}. The complete H_{14} witnesses are legal at every touching pair, including partial boundary stars, and preserve every old face. The radius-15 refutations are geometric Lean deductions derived from selected historical LRAT hints: each initial disjunction is justified from actual face coordinates, six-colour exhaustiveness or fixed colours. Thus no external unsatisfiability premise or trusted SAT encoding is needed. Restriction and composition of completions prove that every one-ring successor fails two-ring feasibility. This proves the two specified traps, without asserting a general obstruction at every depth.

Theorem 6. Lean proves whole-plane validity of the periodic background, closure and disjointness of all period translates of the five supports, and validity for arbitrary binary controls on all of $\mathbb{Z}^2$. Marker readings give an injective encoding. Each support is connected and has exactly two legal fillings with its exterior fixed; the exterior forces the two-colour restriction rather than assuming it.

The formal box consists of both face orientations with $0 \le i, j < n$ and has cardinality $2n^2$. Admissible patterns are precisely restrictions of legal whole-plane colourings. Lean proves $N(\Lambda_{12n}) \ge 2^{5n^2}$ and the corresponding normalized logarithmic bound. To establish the limit, it chooses a whole-plane extension of each admissible pattern and records its restrictions to a padded block cover. This gives, for positive k,

$$N(\Lambda_n) \le N(\Lambda_k)^{(\lfloor n/k \rfloor + 1)^2}.$$

The associated density is nonnegative and bounded above by $(1 + k/n)^2$ times the k-box density. An infimum argument and the vanishing boundary factor prove convergence of the exact printed square-box density. The final entropy bound has no missing limit-existence premise.

Theorem 7. The exact five radical triples and real floor-parity controls are formalised. Lean proves irrationality of each leading quadratic coefficient in every nonzero integer direction, the whole-plane construction, the exact block-period reduction, and the ordered sampling for every line in A–F, every origin and both signs.

The formal proof replaces the stronger Weyl route in Section 5 with an elementary first-difference argument. Irrational linear differences are dense modulo one. A suitable triple of successive values contradicts the assumption that an entire quadratic tail stays in one closed semicircle. Thus the half-interval bit cannot be constant along a nonzero arithmetic tail. Any proposed periodic ray would force just such a constant arithmetic subsequence. This proves both absence of nonzero raw translation periods and absence of eventually periodic rays. It does not formalise the adjacent mismatch density 1/2, the general Weyl theorem or termination of the integer-square-root drawing program.

Proposition 8. For arbitrary five-field switch controls, the proof identifies the actual unique incident face of each fixed colour, 2 and 5. It applies all four printed reference labellers to those faces and proves invariance under $12\mathbb{Z}^2$ at every vertex. This gives all eight periodic phase fields. The colour-2 P_0 histogram is checked on the actual 144-vertex cell and holds for every control assignment. The adjacent value 13/72 is verified by exact integer cross-multiplication, without floating-point arithmetic.

Theorem 9. The signed-coordinate formula in S.11 is formalised on the whole plane. Lean proves its rainbow validity and the stated raw and phase periods. For each of the 24 fields, checked unwrapped vertex lists cover every root modulo the proved skew-period lattice and are closed under all six equal-phase neighbour steps. Finite-path induction and translation then put every whole-plane component inside a list of at most 18 vertices. This is not a torus-only bound. Separately, exact row constraints and predecessor/successor pruning force the canonical row below any canonical row. The resulting rigidity theorem quantifies over every legal colouring agreeing with the witness off any finite set of faces. It is a proof against arbitrary finite edits, not a scan of bounded edits. The earlier 1,800-vertex histogram and its multiplicities retain their historical computational status.

Theorem 10. All three numbered conclusions are proved. The exact decoder, coordinate shear, compact subshift image, arbitrary-support strong irreducibility, stationarity and independence of complete coordinate restrictions are connected to the literal order-three ribbon tiling space. Its labels specify a unique ribbon root, shape and position at every cell, including all four monotone three-cell shapes.

The needed nonexistence theorem is now proved in Lean. Actual integer edge weights have zero unit-square curl, giving a rectangle bound independent of diagonal path length. Strong irreducibility glues two separated complete lines; translation invariance cycles the height coordinates. A checked synchronizing-word and compactness argument supplies an actual periodic diagonal with positive period congruent to one modulo three. The height identities force every such period to be divisible by three, a contradiction. S.5.1 gives the argument. No height primitive, general cocycle-rigidity theorem or external nonexistence assertion is left as a premise.

For the probability branch, the actual topological support of a stationary finitely dependent law is proved nonempty and strongly irreducible. The final iid statement covers literal measurable bounded-radius equivariant maps on arbitrary iid source alphabets and probability spaces; a measurable block-rule representation is derived. The formalization proves the specific order-three consequences needed from Chandgotia–Thorat [4], while retaining their attribution. It does not claim to formalise their general all-order theorem.

Theorem 11. Both branches are proved for the literal finite-interval readouts, including arbitrary assigned endpoint values. The polynomial proof is extended to integer Laurent polynomials by localization and clearing common monomials. In the independent case Lean proves the free integer-lattice circle action, dense orbit, and coverage identity outside finitely many endpoints. It represents the actual finite integer atomic measures by orbit polynomials, applies the syzygy and reassembles them.

The analytic bridge uses finite expressions of floor functions and their actual one-sided jumps. Zero total mass yields an integer periodic step primitive; equality of jumps recovers constants outside a finite exceptional set. Literal interval partitions are converted to these expressions, and operator transport yields the contradiction $1 = 3(a + b)$ at a point avoiding the required translated endpoints. This implements the finite atomic argument without assuming a primitive or an analytic identity and without claiming a general formalisation of distribution theory.

In the dependent case a rational relation gives a nonzero integer raw period. Its sixfold multiple preserves all 24 named phase fields, including endpoint values. Lean checks the exact hexagonal window and its cardinality $1 + 3r(r + 1)$, as well as fixed translates; agreement is one, contradicting the printed positive-lower-mismatch clause. Thus the M3 consequence is also proved. The model restrictions in Section 7.2 remain essential.

S.12.2 The remaining partial numbered result

Theorem 3. Lean proves completeness of the 126-state domain and the exact bijection between raw assignments on the 42 distinct faces of the one-ring tile ball and the rooted six-step transfer walks. Every transition is checked against actual contacts. Bounded arithmetic certificates establish exactly 1,060,600 legal one-ring assignments with the oriented central tile fixed.

The joining-tile test is proved equivalent to an actual legal state on each joining tile, with the cyclic gauge and seam included. The filtered sixth trace is 599,799. The remaining 460,801 assignments are proved to have no complete second ring: a failed joining test contradicts any actual radius-two extension. The original displayed 42-face obstruction also remains verified.

Lean also proves the exact weighted count formula for the 114 distinct faces of the radius-two tile ball. Each of the six sectors contains an inner tile and a radial corner; the edge fibre retains every joining-tile choice satisfying all four surrounding contacts. A bijection reconstructs every legal patch and proves uniqueness, so the sixth matrix trace counts actual fixed-centre patches without identifying rotations.

An additional completed run verifies actual complete second rings for all 50,720 entries of a supplied representative list. Shared adjacency and joining-state certificates reconstruct literal legal 114-face assignments and preserve each 42-face inner seed. The geometric rotation and reflection actions are also proved to preserve completions. This run rebuilt 132 local modules, recorded 7,269 axiom reports and rejected three deliberately false controls. These are checked completions for the listed words; the list size itself does not establish exhaustive orbit coverage.

Two obligations remain: the numerical weighted trace 1,051,975,931,234, and the final connection between the certified representative list and every one of the 599,799 survivors. A separate orbit run accepted its positive modules but failed to complete the required rejection-control checks, so it issued no accepted receipt; the complete list bridge was not replayed. Those unfinished runs are excluded from the coverage claimed here. The historical exhaustive census remains the evidence for the full printed classification and numerical two-ring total. The verification receipts distinguish completed proofs from incomplete attempts.

Part of Theorem 3	Current evidence
Exact 42-face and 114-face geometric correspondences	Lean verified
1,060,600 legal one-ring seeds	Lean verified
599,799 pass the joining tests; 460,801 fail	Lean verified
Every failed seed has no complete second ring	Lean verified
Actual two-ring count equals the weighted sixth trace	Lean verified
Completions for the supplied 50,720 representative entries	Lean verified
Those representatives cover all 599,799 survivors	Historical exhaustive evidence; final Lean coverage/bridge unfinished
Weighted trace equals 1,051,975,931,234	Historical exact evaluations; Lean arithmetic unfinished

S.12.3 Trust, reproduction and attribution

The `formal/` directory of the supplementary archive contains the proof sources, runners, reproduction guide and claim ledger. The ledger lists the checked declarations and remaining obligations. The file `formal/proof-inputs.json` maps each verification run to its exact archived inputs. Earlier verification records are retained with their original scope. Individual receipts record exact source hashes, executable version and hash, commands, output, dependency reports and deliberately false controls.

The new aggregate audit checks all receipts, unchanged proof/data inputs and the current paper correspondence. Successful runs retain their original times and exact manuscript snapshots. The final annotations are bound separately; all eleven numbered mathematical statement texts are unchanged from the preserved pre-upgrade paper. Historical inputs are not silently replaced by the new annotations. The portable companion provides the exact per-receipt input mapping needed to reproduce each run in a fresh checkout view.

All accepted local proof modules were compiled from source with Lean 4.24.0, `--trust=0` and warnings as errors. The analytic, algebraic, topological and probabilistic developments use pinned Mathlib v4.24.0. Each relevant receipt binds the source commit and dependency-cache hashes before and after its run. The audit permits only `propext`, `Quot.sound` and, where actually used, `Classical.choice`. No proof hole, additional mathematical axiom or native decision bypass is accepted. Deliberately false controls must be rejected for mathematical falsehood; an import error, timeout or resource failure is insufficient.

The exporters, source-to-paper correspondence review and packaging remain outside Lean's logical kernel. Exact finite objects and their mathematical deductions are checked inside it. Hashes establish provenance, not semantic correspondence by themselves. The correspondence reports describe the definitions, coordinates and alternate proof routes used for that review.

Lean verification establishes implications from explicit hypotheses and accepted foundations. It does not establish novelty or priority. The rule, Kempe switches, transfer matrices, entropy-limit methods, equidistribution, polynomial divisibility and external ribbon obstruction retain the attribution in the article. The model-specific constructions, reductions and counterexamples are the proposed contributions. Ten complete numbered results do not mean that every adjacent assertion or external dependency has been formalised. The additional written deductions in S.13 are explicitly not Lean verified. Full formal verification of the paper remains unfinished.

S.13 Further finite-growth deductions of 21 September

Verification status. The arguments in this section are written proofs and are **not Lean verified**. They supplement the eleven numbered results and do not supply a construction satisfying M3. The companion archive includes their source drafts and AI reviews under `evidence/september21-sources/`; these reviews are not independent human peer review. The authoring and AI-use disclosure appears in the main article.

S.13.1 Component tails without a frequency-limit assumption

Let F be any finite-alphabet vertex field, let Y be its shift-orbit closure, and put

$$\nu_r = \frac{1}{N_r} \sum_{v \in B_r} \delta_{\tau_v F}, \qquad \mathcal{V}(F) = \{\text{weak subsequential limits of } \nu_r\}.$$

The convention is $(\tau_v F)(u) = F(u+v)$. For $s \geq 1$, set $E_s = \{G \in Y : |C_G(0)| \geq s\}$ and $E_\infty = \bigcap_s E_s$. Then

$$\boxed{\lim_{s\to\infty} \limsup_{r\to\infty} M_{r,F}(s) = \max_{\mu \in \mathcal{V}(F)} \mu(E_\infty).}$$

Proof. The finite alphabet makes Y compact and metrizable. Its probability measures form a compact metrizable space, so $\mathcal{V}(F)$ is nonempty and compact. The relative symmetric difference of B_r and any fixed translate tends to zero, proving shift invariance of every member of $\mathcal{V}(F)$.

The event E_s is determined by the radius-$(s-1)$ ball: a connected component with at least s vertices contains a connected set of s vertices containing its root, found within that ball. Thus E_s is clopen. Since $\nu_r(E_s) = M_{r,F}(s)$, continuity gives

$$a_s := \limsup_r M_{r,F}(s) = \max_{\mu \in \mathcal{V}(F)} \mu(E_s).$$

Write $L = \inf_s a_s$. The sets $K_s = \{\mu \in \mathcal{V}(F) : \mu(E_s) \geq L\}$ are nested, nonempty and closed. Compactness supplies one μ_* in their intersection. Continuity from above gives $\mu_*(E_\infty) \geq L$, whereas every μ satisfies $\mu(E_\infty) \leq a_s$ for every s. This proves the identity and attainment of the maximum. □

Consequently T holds exactly when **every** empirical limit measure gives zero probability to an infinite component at the origin. One favourable subsequential limit is insufficient. No convergence of fixed-threshold frequencies, unique ergodicity, effective description of $\mathcal{V}(F)$ or rate is assumed. The identity applies separately to all 24 phase fields and preserves M3's allowance for exceptional infinite components of zero density.

S.13.2 Finite guards and the exact cost of sealing components

Take a realised legal H_r, so every phase label on B_r is determined, and write $D_r = B_r \setminus B_{r-1}$ for $r \geq 1$. In the induced equal-phase graph on B_r, call a component **open** if it meets D_r, and **sealed** otherwise. Permanently fixing H_r makes every sealed component permanent under extension.

For a field F, let $h_{r,F}(t)$ count sealed components of size t. Define

$$G_{r,F}(s) = \sum_{t<s} t\, h_{r,F}(t), \qquad U_{r,F}(s) = N_r - G_{r,F}(s).$$

The charged mass U includes every vertex of every open component, including its interior vertices, and all sealed components at or above the cutoff. Under an extension fixing the old core, $\Delta G_F(s) \geq 0$ and

$$\Delta N = N_{r+2} - N_r = 12r + 18,$$
$$U_{r+2,F}(s) = U_{r,F}(s) + \Delta N - \Delta G_F(s).$$

For a fixed tolerance ϵ, put $b_r = \epsilon N_r - U_{r,F}(s)$. The endpoint condition $U_{r+2,F}(s) \le \epsilon N_{r+2}$ is equivalent to

$$\Delta G_F(s) \ge (1-\epsilon)\Delta N - b_r.$$

Integer mass therefore requires at least $\max(0, \lceil (1-\epsilon)\Delta N - b_r \rceil)$ newly sealed-small vertices. Newly sealed components contribute their whole mass once, after mergers, and only if the resulting size is below s. The identity does not promise that this amount can be attained.

Here is a sufficient finite-to-infinite criterion. Fix cutoffs $s_j \uparrow \infty$, tolerances $\epsilon_j = 2^{-j}$, finite activation radii $\rho_j \uparrow \infty$ with $\rho_j \ge 2j$, and positive rational floors δ_j. At every published radius $r_n = r_0 + 2n$, permanently retain all earlier cores and require, for every active j with $\rho_j \le r_n$ and every named field,

$$\begin{aligned} U_{r_n,F}(s_j) &\le 2^{-j} N_{r_n}, \\ L_{r_n,F}(t) &\ge \delta_j N_{r_n} \qquad (0 < \operatorname{dist}(t) \le j), \\ \left| \frac{C_{r_n,F}(m,a,\ell)}{N_{r_n}(m,a)} - \frac{1}{6} \right| &\le 2^{-j} \qquad (1 \le m \le j). \end{aligned}$$

Here $L_{r,F}(t)$ counts mismatching ordered pairs with both endpoints in B_r; its denominator remains the **whole** N_r. The quantities $C_{r,F}(m,a,\ell)$ and $N_r(m,a)$ count label-ℓ vertices and all vertices in $B_r \cap (a + m\mathbb{Z}^2)$, respectively. The last guard ranges over every residue a and all six labels.

Conditional conclusion. Any indefinitely continuing nested legal sequence meeting these guards has an M3 union. Indeed, final components of size at least s_j can contain no previously sealed component below that cutoff, so their mass fraction at publication radii is at most 2^{-j}. The internal mismatch count is a lower bound for the full mismatch count in A. A fixed translation retains its positive floor from a fixed active level. The address guards give limiting frequency $1/6$ on every fixed residue class. Between successive publication balls there are $O(r)$ additional vertices, compared with $\Theta(r^2)$ vertices in each fixed coset. Squeezing transfers all three conclusions to every radius. Every finite-index subgroup contains some $m\mathbb{Z}^2$, so its cosets are finite unions of these classes. Fixed shifts of the observation centre have the same vanishing boundary error.

This is a sufficient criterion with an unfulfilled continuation premise. A concrete two-stage implementation can store a checked next annulus in each intermediate H_{r+1} state and consume it to reach the next admitted H_{r+2}. The missing theorem is that **every admitted realised source** has such a pair of annuli under the same schedules and restrictions. A realised root, effective schedules, sound admitted-state checks and a selector remain to be supplied. Decidable verification of a supplied certificate does not make its unrestricted existential projection decidable.

S.13.3 Trimer congruence and a stronger simultaneous radius bound

Every vertex of a raw colour-k triangle has the same phase $\Phi_{k,H}$ for each reference family H: its phase is the global reference label of that very triangle. In a sealed phase component, all three vertices of each incident colour-k triangle lie in the component. Distinct colour-k triangles have disjoint vertex sets by rainbow legality. They partition the component into triples. Hence every sealed component has size divisible by three, and

$$G_{r,F}(s) \equiv 0, \qquad U_{r,F}(s) \equiv 1 \pmod 3.$$

For $r \ge 1$, every frontier vertex belongs to an open component and $|D_r| = 6r$. Thus $X_{r,F}(s) := U_{r,F}(s) - 6r$ is nonnegative and congruent to one modulo three. For a level $j \ge 1$, put $q = 2^j \ge 2$. The mass guard necessarily satisfies

$$q(6r+1) \le 1 + 3r(r+1), \qquad r \ge 2q, \qquad j \ge 1.$$

For the last implication, $q6r \le N_r$ first excludes $1 \le r \le 2q-2$; at $r = 2q-1$ the stronger inequality has slack $1 - q < 0$. Radius zero also fails the guard. At the next integer radius $r = 2q$, its per-field allowance

permits only $X \in \{1, 4\}$. These are necessary values, not realised examples. The restriction $j \geq 1$ is essential: for $j = 0$, $U_{r,F}(s) \leq N_r$ holds automatically. In particular, at $r = 1$ every phase component meets the frontier, so $U_{1,F}(s) = N_1 = 7$; the guard does not force $r \geq 2$.

The common raw colouring gives a stronger bound for simultaneous guards. Let A_k consist of the vertices of D_{r-1} whose incident colour-k face reaches D_r. For $r \geq 2$, an inner corner has four incident faces reaching D_r; an inner side vertex has three. Distinct incident faces have distinct colours, and double counting yields

$$\sum_{k=0}^{5} |A_k| = 6 \cdot 4 + 6(r-2) \cdot 3 = 18r - 12.$$

For $r = 1$, the inner set is the origin and all six colours reach D_1, giving the same formula. Some colour k has $|A_k| \geq 3r - 2$. Its vertices are open in **all four** reference fields of that colour. They are disjoint from D_r, so

$$U_{r,(k,H)}(s) \geq 9r - 2 \quad \text{for every cutoff } s \text{ and every } H.$$

For $j \geq 1$, simultaneous level-j guards therefore require $q(9r - 2) \leq N_r$. If $2q - 1 \leq r \leq 3q - 2$, its slack is

$$1 + 3r(r + 1 - 3q) + 2q \leq 1 - 3r + 2q \leq 4 - 4q < 0.$$

Consequently their necessary actual publication radius is

$$\boxed{r \geq 3 \cdot 2^j - 1 \quad (j \geq 1).}$$

The selected colour may vary with the core and radius. The condition concerns actual active publications, including the first $r_0 + 2\max(0, \lceil(\rho_j - r_0)/2\rceil)$, rather than an unconditional lower bound on the declared ρ_j. It proves neither sharpness nor feasibility of any remaining radius, and it does not rule out M3.

S.13.4 Growth families must respect the scope of their symmetries

For the exact ring-extension relation R, a cycle of families $I_0, \ldots, I_{d-1}$ with $I_i \cap E \subseteq \mathrm{Pre}_R(I_{i+1})$ (indices modulo d) is sufficient for d-ring return from realised starts in I_0. It is not necessary for a particular chosen intermediate family cycle. Here E is the exact boundary envelope; membership alone is not a realised raw patch. A periodic control demonstrates the two-phase rule without satisfying M3.

A useful negative test transports known finite traps. The reflection $\sigma(i, j) = (i + j, -j)$ preserves radius and sends

$$\sigma U(i, j) = D(i + j, -j - 1), \qquad \sigma D(i, j) = U(i + j + 1, -j - 1).$$

It bijects every H_r and preserves contacts and colours. Applying its inverse to an extension proves that a patch and its reflection have identical unrestricted extension possibilities at every depth. This does not identify membership in an oriented candidate family unless that family is also transported or proved reflection invariant.

For example, the recorded forward-only boundary predicate avoids all colour relabellings of the consecutive mask triple $(22, 38, 41)$, including the cyclic seam. A mask m represents the set of colours whose bits occur in m. It admits the reflected saved H_6 trap, although it excludes the original. The retained finite diagnostic verifies this membership change and the boundary reversal; the reused exhaustive obstruction states that the trap has first annuli but no second ring. Thus this full family cannot have universal d-ring return for any $d \geq 2$. It cannot have universal one-ring closure either: applying that universal property twice would construct two rings. This latter argument does not assert that the same reflected H_6 has no first-ring successor in the family. Proper subfamilies and other candidates remain undecided. The application retains its historical computational premise and is not Lean verified.