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

Aidan Quigley
21 September 2026 · BAC260919 · Preprint

Abstract

Colour the triangular faces of the regular triangular tiling with six colours, requiring all six colours at every vertex. An explicit construction has five independently switchable binary choices per 288 triangles, giving topological entropy at least $5 \log 2/288$. Quadratic-phase controls produce a whole-plane colouring with no translation period and no eventually periodic ray in six canonical line families, yet eight associated phase fields remain periodic. We connect this separation of freedom and hidden order to exact finite results: a four-colour characterization of the complete 120-type inventory; seed counts and a complete second-ring extension classification; a 30-state annulus transfer; and proof-certified patches that extend two rings but not three. A separate periodic control has phase domains of at most 18 vertices in all 24 fields, although no nonempty finite legal recolouring is possible. Finally, a trimer-to-ribbon factor applies a theorem of Chandgotia and Thorat to exclude stationary finitely dependent laws, and an integer jump-measure argument excludes a specified class of finite-interval circle codings. These results distinguish local legality, indefinitely sustainable extension, extensive freedom and phase disorder. The phase-domain existence problem M3 remains open. Lean verifies ten of the eleven numbered results completely and part of the remaining census theorem; additional finite-growth deductions are not Lean verified.

Keywords: triangular tiling; precise colouring; topological entropy; Kempe switches; aperiodicity; trimer covers; ribbon tilings.

1. One local rule, several different questions

A picture can suggest a mathematical question before it suggests a definition. The motivating question [1] asked how to arrange six colours on triangular faces so that triangles of the same colour never touch, even at a corner, while avoiding an obviously repeating pattern. Around a vertex there are exactly six triangles. The requirement is therefore equivalent to saying that each vertex sees each colour exactly once. Equivalently, on the dual honeycomb graph this is a six-colouring of vertices in which every hexagonal face is rainbow.

The local rule is easy to draw and unusually restrictive. A face must disagree with every other face that shares one of its three vertices. Yet it does not determine a unique continuation from a small neighbourhood. The difficulty is to understand the kind of freedom that remains. A large collection of legal finite pictures, a formula for an infinite nonperiodic colouring, and an infinite colouring without persistent phase order are three different accomplishments.

This article connects finite extension to infinite freedom. Exact inventories and a seed census expose local boundary obstructions; an annulus algorithm and certified traps explain why several successful growth steps do not establish indefinite continuation. We then construct exponentially many plane-admissible patterns and one explicit nonperiodic colouring, before

reading their colour classes to expose retained order. A frozen small-domain control and two further obstruction arguments show why additional conditions are needed.

1.1 Coordinates and legality

Use axial vertices $(i, j) \in \mathbb{Z}^2$, embedded in the Euclidean plane as

$$(i, j) \longmapsto (i + j/2, \sqrt{3}\, j/2).$$

The two orientations of face are

$$U(i, j) = \{(i, j), (i + 1, j), (i, j + 1)\},$$
$$D(i, j) = \{(i + 1, j), (i, j + 1), (i + 1, j + 1)\}.$$

A colouring assigns a value in $\{0, 1, 2, 3, 4, 5\}$ to every face. We call it **legal** if the six incident faces at every vertex have distinct values. Let X be the space of all legal whole-plane colourings. A finite assignment is locally legal when any two of its present faces that share a vertex have different colours. At an incomplete boundary star this tests the faces present; it supplies no missing faces.

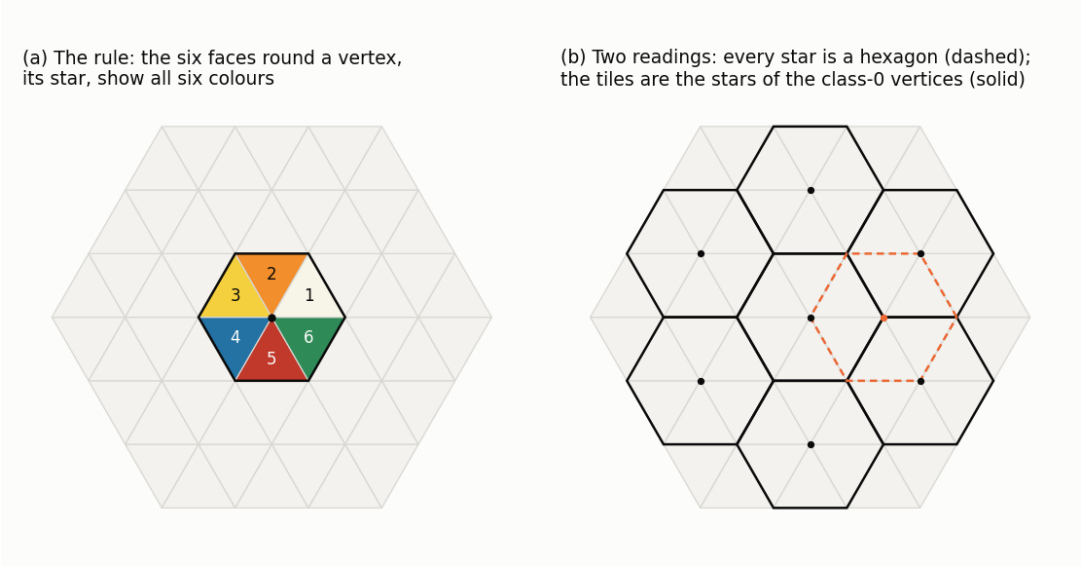


Figure 1: The local rule and the two readings of a hexagon: a disjoint tile inventory and overlapping vertex stars. The first panel uses display digits 1–6, equal to colour index plus one. The distinction between readings matters when a finite patch is said to contain all 120 cyclic classes.

The six incident faces at (i, j), in cyclic order, are

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

There are $6!/6 = 120$ cyclic orders up to rotation, with reflections distinguished. One can inspect these orders at every vertex, allowing overlapping stars, or choose only the centres with $i - j \equiv 0 \pmod 3$. The stars at those centres partition the faces into disjoint six-triangle hexagonal tiles. An inventory of 120 such tiles concerns this second reading. It should not be confused with a count of arbitrary overlapping vertex stars.

1.2 The contributions and their scope

The main statements are these. All numerical censuses below refer to their specified finite spaces.

1. A periodic background contains five families of finite two-colour components whose period translates are disjoint. Arbitrary binary switches on them remain legal, giving

$$h(X) \geq \frac{5 \log 2}{288} = 0.012033805218\ldots$$

   nats per triangular face (Theorem 6).
2. Quadratic irrational controls on these switches give a specified colouring C_5 with neither a nonzero translation period nor an eventually periodic ray in any of six lattice line families (Theorem 7).
3. Every colouring in this switch family retains eight periodic phase fields. In particular C_5 does not satisfy the phase-domain target M3 (Proposition 8).
4. A local, translation-equivariant trimer-to-ribbon map transfers the ribbon obstruction of Chandgotia–Thorat [4] to X: X contains no nonempty strongly irreducible subshift and supports no stationary finitely dependent probability law (Theorem 10).
5. A separate algebraic obstruction rules out M3 for a shared-circle coding with finite-interval readouts and no additional address state (Theorem 11).
6. Exact finite results connect inventories to extension: the four-colour inventory criterion, the one- and two-ring seed counts and complete second-ring classification, and the 30-state annulus transfer (Theorems 2–4). Checked finite traps show that two-ring feasibility is not closed under continuation (Proposition 5).
7. An explicit periodic control has phase domains of at most 18 vertices in every named field but admits no nonempty finite legal recolouring (Theorem 9).

The all-120 example of Section 2 is an assignment-specific finite nonextension certificate. The seed classification stops at a second ring, and the traps fix their entire saved interiors. The entropy theorem counts a family of whole-plane colourings; it does not assert positive orbit entropy for the single deterministic example C_5. The ribbon nonexistence result is an application of an external theorem, with the factor map supplied here. The circle result concerns its stated model, not every quasiperiodic colouring. These qualifications are part of the results.

1.3 Relation to earlier work

Precise colourings of triangular tilings have a substantial history. Rigby [2] discusses the same six-colour setting, strip constructions and the possibility of a more random-looking arrangement. We make no claim to have introduced the rule, its elementary periodic examples

or its trimer interpretation. The thesis of Yaz (1997), cited by Rigby, was not available to us; no priority claim for symmetry classification is made here.

For a fixed colour, its triangles cover every lattice vertex exactly once. Thus each colour class is a triangular-trimer cover, and the six covers fit together to partition all triangular faces. Verberkmoes and Nienhuis [3] solved a triangular-trimer model. Their single-cover entropy does not by itself count six mutually compatible labelled covers, so we prove the lower bound for the coupled colouring rule directly.

The local exchanges used below are Kempe switches: exchanging two colours on an entire connected component preserves legality. The useful information is the explicit packing of five independent families, their density, and the ability to read each switch from a marker face. Transfer matrices, graph potentials and finite proof certificates are established methods; the contributions are their exact model-specific reductions, constructions and counterexamples. Our obstruction to finite dependence retains the attribution to the versioned ribbon theorem in [4]; the formal companion proves the needed order-three case. Supporting arguments and certificates appear in Supplement S.1–S.13. The additional component-tail, growth-certificate and frontier deductions in S.13 have written proofs and recorded reviews, but no Lean verification.

2. Complete inventories and a local obstruction

Before counting patterns in X, it is worth seeing why a locally legal picture need not belong to its language. A search may return a large valid assignment whose boundary has already made continuation impossible. Even possessing every cyclic colour order is no safeguard against this failure.

The finite assignment supplied as `evidence/serrated-t8.json` in the supplementary archive colours 720 triangular faces arranged as 120 disjoint hexagonal tiles. It contains each of the 120 cyclic colour classes once. Its perimeter, measured in exposed tile edges, is 134. Direct verification checks contacts between every pair of present faces, including contacts at boundary vertices. It is a locally legal assignment.

The obstruction is much smaller than the picture. The absent face $D(-10, 8)$ has vertices $(-9, 8)$, $(-10, 9)$ and $(-9, 9)$. For each possible colour there is a present face of that colour touching one of those vertices:

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

Proposition 1 (assignment-specific nonextension). The saved serrated assignment cannot be extended to a legal colouring of the plane while preserving its present colours.

Lean status. Complete for this nonextension statement. Separate Lean proofs also verify legality at every contact between distinct present faces, perimeter 134, each of the 120 cyclic types exactly once, and exactly ten blocked missing faces for the saved assignment. See S.12.

Proof. Any whole-plane extension assigns one of the six colours to $D(-10, 8)$. The corresponding row of the table supplies a face with the same colour at a shared vertex. This violates legality. □

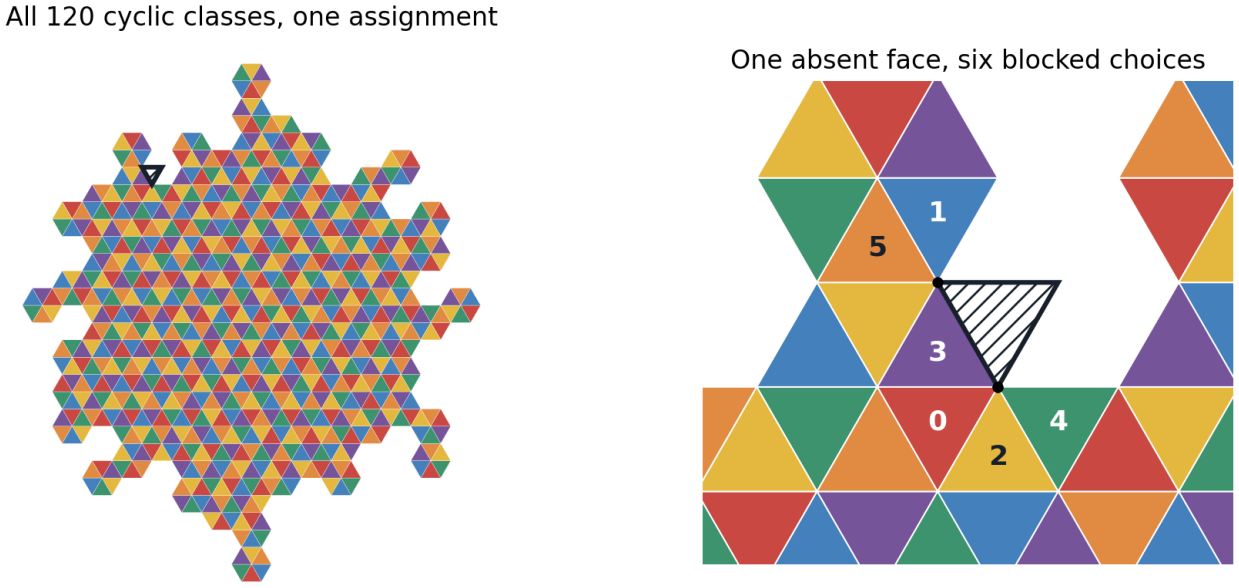


Figure 2: The saved serrated assignment and a close view of its blocked missing face, shown hatched. Digits identify one blocker of each colour; the black points are the shared vertices. All colours shown belong to the saved assignment.

This proof does not say that the shape is intrinsically unextendible. Another colouring of the same faces could have different boundary possibilities, and recolouring some of the displayed faces may remove the obstruction. The example does not decide a minimum-perimeter inventory problem. It establishes a property of one exactly specified assignment, identified by its canonical hash in Supplement S.1.

The companion data contain five serrated assignments, with perimeters from 134 to 170; each has an assignment-specific six-colour obstruction. Only the smallest-perimeter one is needed here. The complete certificate for its displayed obstruction is the six-row table, together with the raw assignment. No machine-local verification receipt is needed to understand the implication.

2.1 Recognising all 120 types through four colours

There is an exact way to recognise the inventory without testing 120 type multiplicities separately. On each cyclic tile read its six directed consecutive words of length four. There are $6 \cdot 5 \cdot 4 \cdot 3 = 360$ possible words with distinct colours. The direction matters: reflection remains distinct from rotation.

Theorem 2 (four-colour inventory criterion). A multiset of cyclic six-colour types contains every type once if and only if every ordered quadruple of distinct colours occurs exactly twice among its consecutive words.

Lean status. Complete for arbitrary natural-valued type multiplicities; the lower-order counterexample in S.8 is also verified. See S.12.

Proof. Fix a word (a, b, c, d). The two unused colours e, f give exactly two completions, (a, b, c, d, e, f) and (a, b, c, d, f, e), understood cyclically. If m_t is the multiplicity of type t, the word condition is $m_t + m_{t'} = 2$.

Join these two completions for every word. The resulting graph on the 120 types is connected: its moves swap adjacent colours, and adjacent swaps generate every circular order. It contains the five-edge cycle shown in Figure 3. Writing $y_t = m_t - 1$, each edge forces $y_{t'} = -y_t$. Alternating signs around that odd cycle gives $y_t = -y_t$, hence $y_t = 0$. Connectivity then forces every multiplicity to one. Conversely, an inventory with one copy of each type has exactly the two completions of each word. □

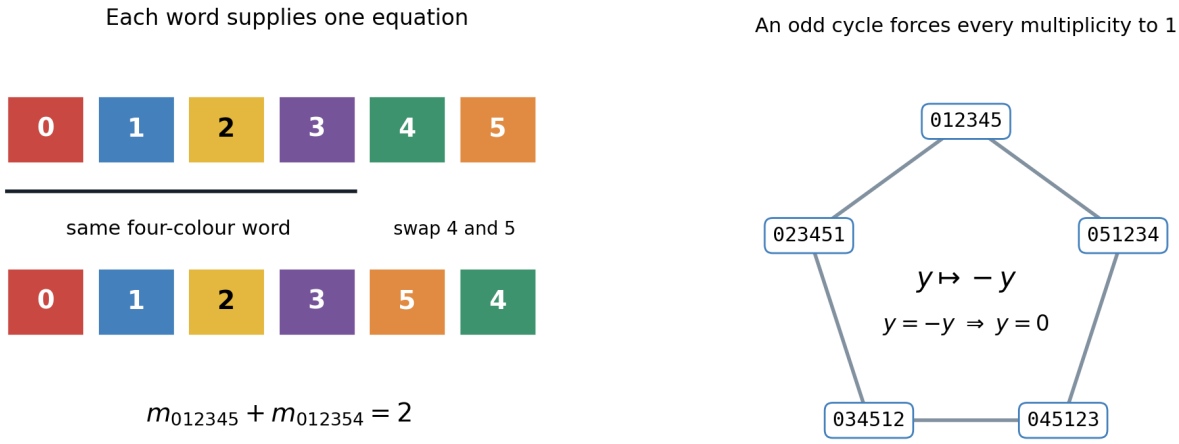


Figure 3: Four consecutive colours leave two cyclic completions (left). An odd cycle in the completion graph (right) forces the alternating multiplicity deviations to vanish. Each six-digit label is a cyclic order, written with 0 first. The diagram is combinatorial; it does not prescribe positions of tiles in a legal patch.

Pair and triple balances are insufficient. Supplement S.8 gives an explicit multiset with four types omitted and four doubled, yet with every directed pair occurring 24 times and every directed triple six times. The theorem concerns the type multiset only. It neither solves the

geometric placement constraints nor proves extension of a placement; Proposition 1 remains an example of that distinction.

3. From a seed census to the boundary of extension

3.1 Exactly what the finite census counts

Consider the centre tile and its first one or two complete rings of disjoint tiles. In tile coordinates (u, v), take $\max(|u|, |v|, |u + v|) \leq r$, and map tile centres to axial vertices by $(i, j) = (u + 2v, u - v)$. The two footprints contain 7 and 19 tiles, hence 42 and 114 triangular faces. Fix the central colour tuple to $(0, 1, 2, 3, 4, 5)$ in the cyclic convention of Section 1. Repeated tile types are allowed here.

Theorem 3 (finite seed census and second-ring classification). With this oriented centre fixed, there are exactly 1,060,600 locally legal one-ring seeds and 1,051,975,931,234 locally legal two-ring seeds. Of the one-ring seeds, exactly 460,801 admit no complete second ring; all remaining 599,799 admit one.

Lean status. Partial. Lean verifies the exact geometry, all 1,060,600 one-ring assignments, the 599,799 joining-test survivors, noncompletion of all 460,801 failed tests, and the equality of the actual two-ring count with its weighted trace. It also verifies complete second rings for a supplied list of 50,720 representative seeds and their geometric symmetry transports. Two gaps remain: the numerical weighted trace 1,051,975,931,234 and the final proof that those representatives cover every survivor. The full printed census retains its historical computer-assisted evidence; it is not claimed fully Lean verified. See S.9 and S.12.

The result combines exact transfer correspondences with completed finite evaluations. It is not a count of infinite extensions. In particular, the two-ring total counts full 114-face assignments, while 599,799 counts their distinct 42-face inner seeds; an inner seed can have many completions.

The first transfer is small enough to explain directly. A neighbouring tile has 126 oriented states compatible with the fixed centre. Rotation together with a cyclic renaming of colours identifies the six positions with one state set. Put $M_{pq} = 1$ when consecutive neighbours in states p, q have no colour clash, and zero otherwise. There are 1,264 allowed transitions. The six neighbours form a wheel around the centre; there are no contacts between nonconsecutive outer tiles. Thus legal seeds correspond exactly to rooted closed walks of length six, giving $\mathrm{tr}(M^6) = 1,060,600$.

Each adjacent pair of inner tiles faces one joining tile in the second ring. Reject a transition if that joining tile has no legal state. This removes 128 transitions and leaves a matrix M_+ with $\mathrm{tr}(M_+^6) = 599,799$. Passing these six separate tests is necessary. Its sufficiency at this particular depth is the additional computer-assisted result: the exhaustive Boolean-relation census found a complete second ring for every survivor, and an independent implementation reproduced every block. Supplement S.9 gives the correspondences, the weighted two-ring count and the reproducible computational records. The partial Lean coverage of this computation is stated above.

The local contradiction in Figure 4 can be read without a solver. The two missing faces

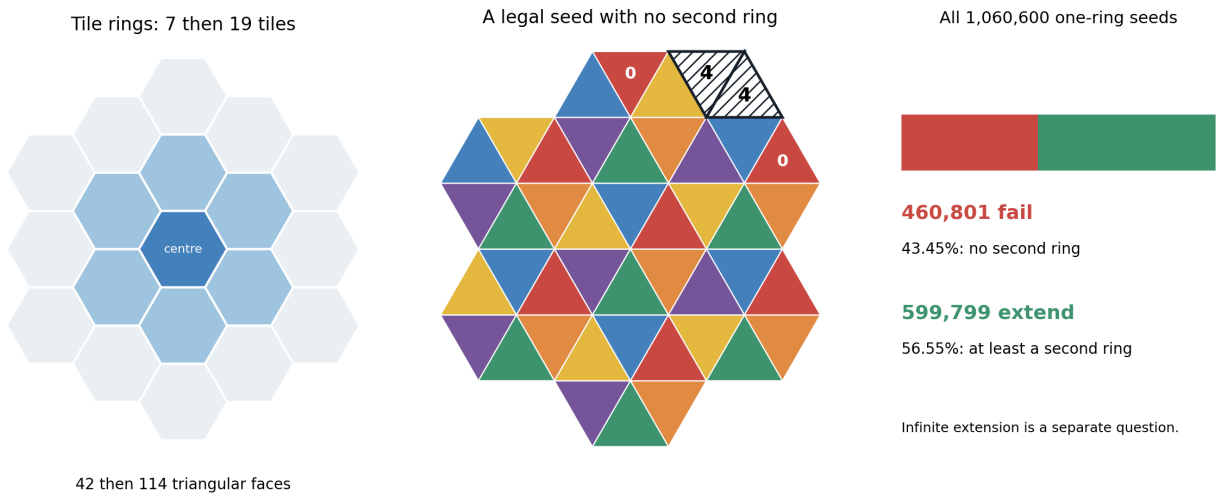


Figure 4: The seed census uses rings of disjoint tiles: 7 tiles at radius one and 19 at radius two (left). The middle panel displays an exact legal 42-face seed. Its two hatched missing faces must each be colour 4, although they touch, so no second ring is possible. The right panel shows the exact one-ring classification, not a sampling estimate or an infinite-extension probability.

$D(-1,2)$ and $U(0,2)$ meet at $(0,2)$, where four present faces already use colours $1,2,3,5$. The missing pair must therefore use $0,4$. Each also touches a present colour-0 face at another vertex, so both are forced to 4. The exact seven tile tuples and contacts are in S.9. This is the boundary mechanism behind the failed side tests, not merely a low success rate in an experiment.

3.2 An exact annulus interface

For growing general patches it is useful to change footprints. Write H_r for the union of all vertex stars centred at axial distance at most r. Then H_r has $6(r+1)^2$ faces. In particular, H_1 has 24 faces, not the 42 faces in the tile-ring seed census. The notation and the two constructions must not be interchanged.

At a boundary vertex retain the set of colours already present. These masks contain every constraint that the fixed interior imposes on an exterior face: it must avoid the colours at each shared vertex. Interior colours away from the boundary are irrelevant to that extension decision.

Theorem 4 (exact annulus transfer). For a fixed legal H_r, a cyclic transfer with 30 states decides whether it extends to H_{r+1}. The Boolean decision and construction of a witness take a linear number of fixed-width operations in boundary length. The integer trace counts its legal next annuli. A composed transfer with at most 900 states decides whether two further rings exist.

Lean status. Complete for the numbered theorem: all-radius geometry, the closed-path bijection and seam, integer trace, a linear-work Boolean decision and returned geometric witness,

and the composed two-ring decision. The supplement's separate constant-memory implementation is outside this proof's memory claim. See S.12.

The state records two different colours, one on an edge face and one on its forward singleton: there are $6 \cdot 5 = 30$ choices. At each boundary vertex a local transition completes its rainbow star and checks the remaining outer point contact. A path must close at the final seam. For the resulting local matrices,

$$H_r \text{ extends to } H_{r+1} \quad \Longleftrightarrow \quad \mathrm{tr}_{\mathrm{B}}(T_0 \cdots T_{L-1}) > 0, \qquad L = 6(r+1).$$

Supplement S.10 proves the geometry and reconstructs the annulus from a closed path. The transfer also emits the next masks, allowing two layers to be composed. Its Boolean two-ring API decides existence; it is not an integer count or a witness-producing API. Likewise, the linear Boolean bound is not a linear bit-complexity claim for arbitrarily large exact counts.

3.3 Two successful rings still need not lead anywhere

Let $K_m(P)$ mean that the exact patch P extends through m further complete rings, with all its existing colours fixed. Write $P \to Q$ for a legal one-ring extension. Restriction and concatenation give the elementary identity

$$K_{m+1}(P) \quad \Longleftrightarrow \quad \exists Q\, [P \to Q \text{ and } K_m(Q)].$$

Proposition 5 (certified failure of two-ring closure). There are exact legal H_{12} assignments in $K_2 \quad K_3$. Consequently, the family of patches admitting two further rings is not closed under one-ring continuation: every one-ring successor of either assignment fails K_2.

Lean status. Complete for both exact saved inputs, including the legal two-ring witnesses, three-ring impossibility and the universal successor conclusion. See S.12.

Proof with finite certificates. The retained records `0015-1` and `0018-1` each have a legal H_{14} witness preserving their 1,014 original faces. For each, an exact no-touch encoding of completion on H_{15} has a checked LRAT refutation. The historical independent replay reconstructed the geometry and clause multiset and accepted each refutation with two proof checkers. Thus K_2 holds and K_3 fails. If some successor had two further rings, the displayed identity would give three further rings for the original patch, a contradiction. The witnesses, refutations, encoding argument and precise source hashes are retained with S.10. □

This result excludes a proposed invariant, not the underlying one-ring seeds, which have separately certified patterned infinite extensions. A construction may choose a different exterior, exclude these particular snapshots, or recolour old faces. Nor do two counterexamples exclude every fixed lookahead depth or every boundary family.

For a fixed seed, extensions at *every* finite radius would give an infinite colouring by the finitely branching tree argument: repeatedly choose a child with descendants at arbitrarily large depths. A finite cutoff supplies no such premise. An effective growth proof instead needs a realised seed in a family whose members always have a successor in that family; a repair

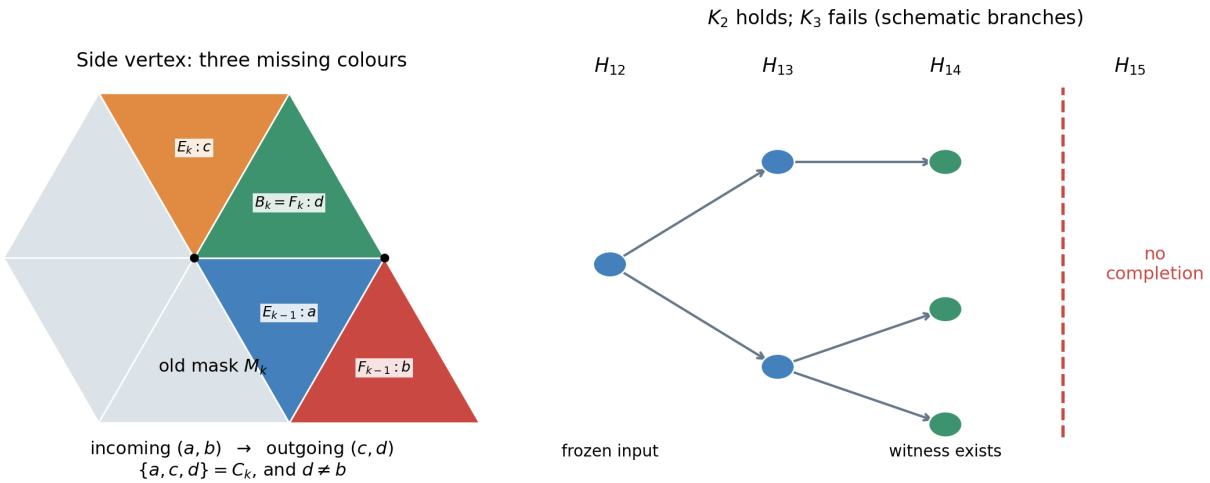


Figure 5: At a side vertex of the annulus, the old mask and the incoming/outgoing two-colour states capture every required contact (left). The extension tree (right) illustrates the quantified trap: some paths reach radius 14, but none preserving the frozen radius-12 assignment reaches radius 15. The branching is schematic, not a count of successors.

process also needs permanently fixed cores exhausting the plane. Phase-domain behaviour requires further arguments. The next section takes a route where whole-plane validity is established before any patterns are counted.

4. Five local choices and a positive entropy bound

The six-colour rule uses exactly as many colours as there are faces in a vertex star. Such a tight count might suggest rigidity. Indeed, a proper three-colouring of the *vertices* of the triangular lattice is determined by the colours at two adjacent vertices: the third vertex of each elementary triangle is forced. Propagation yields exactly six colourings, all periodic. That comparison is motivational; our faces and their contact graph are different objects.

To measure freedom in the face rule, let $N(\Lambda)$ be the number of patterns on a finite face set Λ that occur in some colouring of X. Set

$$\Lambda_n = \{U(i, j), D(i, j) : 0 \leq i, j < n\}.$$

The topological entropy per face is

$$h(X) = \lim_{n \to \infty} \frac{\log N(\Lambda_n)}{2n^2}.$$

The limit exists by the standard subadditivity argument for rectangular pattern counts. Logarithms are natural. Counting only restrictions of whole-plane colourings is essential: a finite locally legal assignment such as Proposition 1 must not inflate the count.

4.1 A binary switch with its exterior fixed

For two colours a, b, form the graph whose vertices are the faces of those colours, with an edge when the faces share a lattice vertex. A connected component of this graph is an $\{a, b\}$-Kempe component.

At every vertex touched by a complete component, both the a-face and the b-face belong to it. They meet at that vertex and so lie in the same component. Exchanging a and b on the whole component therefore exchanges two entries of a rainbow star and leaves the other four unchanged. Legality survives at every vertex, including the component's boundary.

For a finite complete component, the exchange is a local binary choice. More precisely, when its exterior is fixed there are exactly two legal fillings of the component. At each touched vertex the four exterior faces fix the other four colours, so a component face must still use a or b. Its connected pair graph is bipartite under the original colouring; one face determines all the others. Either initial choice works.

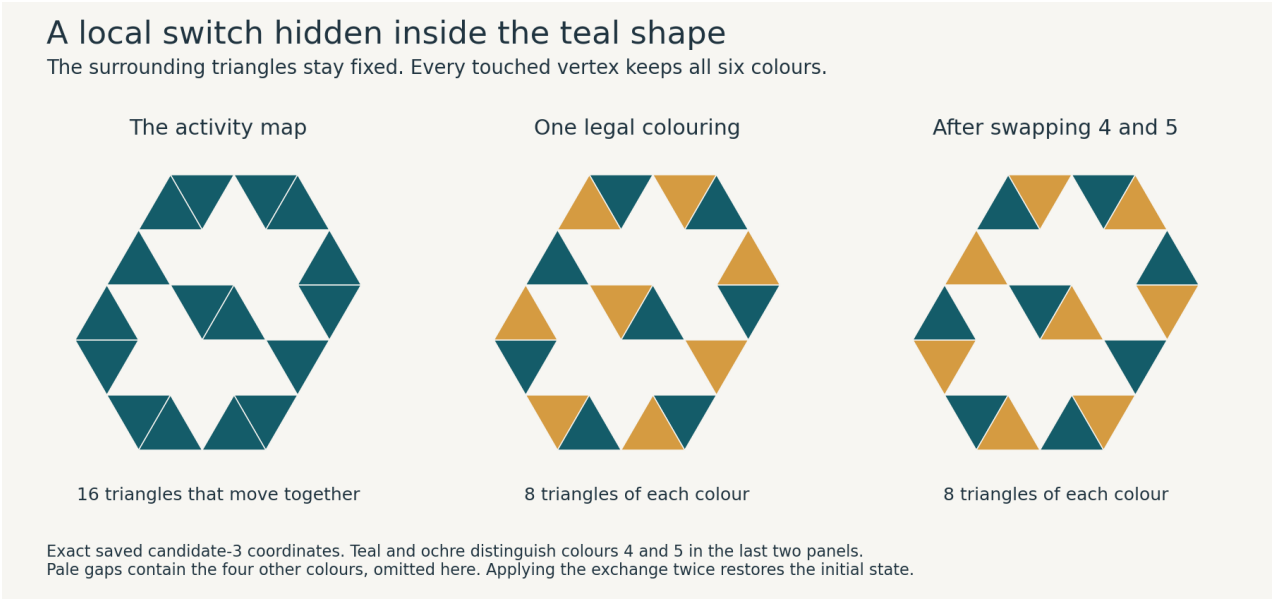


Figure 6: A sixteen-face switch observed in a separately grown finite witness. The two right-hand panels show its legal states with the exterior fixed; pale gaps contain faces of other colours. The same uncoloured support occurs, after a lattice isometry, among the periodic construction's switch components. This finite occurrence illustrates the mechanism and does not establish a density of independent choices.

Finding one switch does not prove positive entropy. It might occur only once, or repeated switches might interact. The construction must supply a positive density of choices, certify that they can all be made independently, and show that different choices produce different observable patterns.

4.2 A periodic background and five compatible families

Let B be the colouring of periods $(12, 0)$ and $(0, 12)$ given by the two digit tables in Supplement S.2. Its 144 vertex stars are all rainbow, which proves legality everywhere by periodicity. Its fundamental cell contains 288 triangular faces.

In this background choose the five finite supports $J_0, \dots, J_4$ listed in that supplement. Each is a complete connected component for its stated colour pair. Their sizes and marker faces are:

Family	Colour pair	Faces	Marker in a period cell
J_0	$\{0, 1\}$	32	$U(10, 4)$
J_1	$\{0, 1\}$	20	$U(4, 6)$
J_2	$\{3, 4\}$	16	$D(6, 0)$
J_3	$\{3, 4\}$	16	$D(9, 3)$
J_4	$\{3, 4\}$	28	$D(11, 10)$

The 112 support faces have distinct **oriented residues** (U or $D, i \bmod 12, j \bmod 12$), both within and across families. This finite fact proves that all the sets

$$J_r + (12m, 12n), \qquad r \in \{0, \dots, 4\}, \quad (m, n) \in \mathbb{Z}^2,$$

are pairwise disjoint. If two translated supports met, their original faces would have the same oriented residue. Completeness in the infinite background is checked by listing every same-pair neighbour of each support face, rather than by finding a component inside a cropped window.

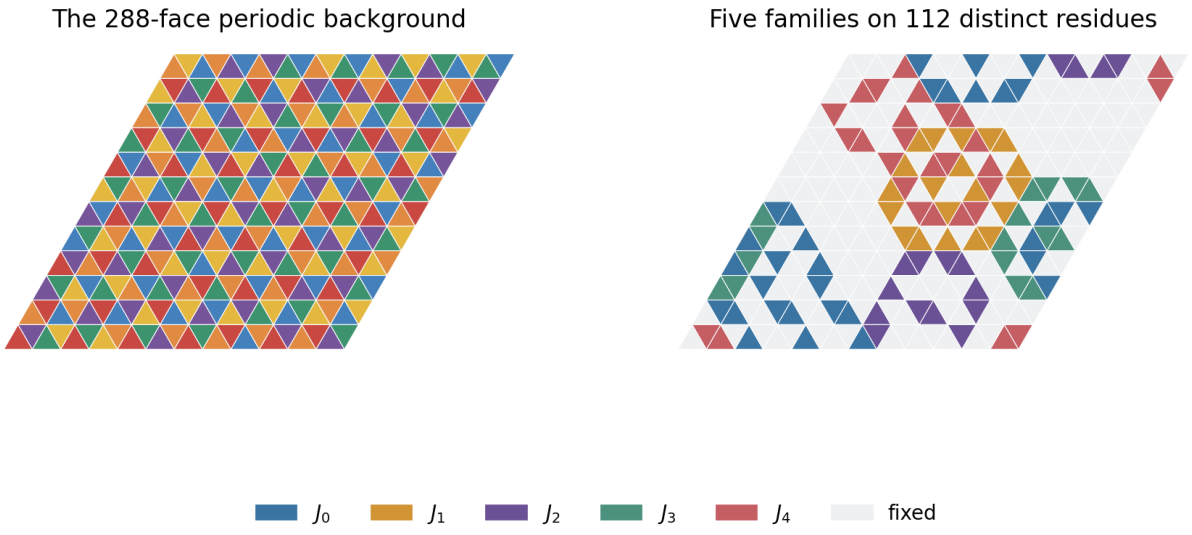


Figure 7: One fundamental cell of the background and the residues of its five switch families. The right-hand colours distinguish families, not raw face colours. Some unwrapped components cross cell boundaries; the supplement retains their full coordinates.

Assign an arbitrary bit $\varepsilon_r(m,n)$ to each translated component, and exchange its colour pair exactly when that bit is one. Denote the resulting whole-plane assignment by C_ε.

Theorem 6 (independent switches and entropy). Every total assignment of the bits produces a colouring $C_\varepsilon \in X$, and

$$h(X) \geq \frac{5 \log 2}{288}.$$

Lean status. Complete. The whole-plane switch family, actual admissible-pattern count, square-box entropy limit and stated lower bound are verified. Limit existence is proved by a finite block cover and a vanishing boundary factor. See S.12.

Proof. Disjoint supports give each face an unambiguous final colour. A support touching a vertex contains both faces of its pair. Two disjoint supports touching the same vertex cannot share a colour, because that vertex has only one face of each colour in B. Thus all switches acting on that star are disjoint transpositions of its six colours. Every star remains rainbow, for arbitrary choices on the entire plane.

For each family, translate its representative by a background period so that its chosen marker is in the canonical cell. This only reindexes that family's bits. The marker has its background colour when its bit is zero and the other colour of its pair when its bit is one. Other switches do not touch it.

The n^2 macrocells comprising Λ_{12n} contain $5n^2$ such markers. Fix every uncounted bit and vary these $5n^2$ bits independently. The markers distinguish all 2^{5n^2} restrictions, and each is the restriction of a legal whole-plane colouring. Hence

$$N(\Lambda_{12n}) \geq 2^{5n^2}, \qquad \frac{\log N(\Lambda_{12n})}{288n^2} \geq \frac{5 \log 2}{288}.$$

Take the entropy limit. Components crossing the window boundary cause no loss: the count reads markers inside the window, while their switches are performed on the plane. □

4.3 What this number measures

The lower bound is a guaranteed rate of independent information compatible with the rule: five binary choices per 288 faces. It does not measure a colour's frequency, a solver's diversity, or a visual randomness score. The exact entropy and the optimal packing density remain undetermined.

The finite part of the proof is small enough to print. It consists of the 288 background digits, five coordinate lists, component closure and residue disjointness. The infinite conclusion follows from the star argument and the marker count. A larger solver run is not a substitute for either step, and no solver is used in the replay accompanying this article.

There is also a cost to the construction's convenience. Every switch keeps a face within one block of

$$\{0,1\} \mid \{2\} \mid \{3,4\} \mid \{5\}.$$

The corresponding block field never changes, even when the bits are chosen independently. We will use this invariant twice: first to prove that a particular choice of controls has no translation period, and then to identify the periodic order that survives in its colour-class geometry.

5. An explicit colouring without periodic rays

The bits in Theorem 6 may be arbitrary. To obtain one exactly specified colouring, we choose them by quadratic irrational phases. This makes the absence of repeated tails a theorem about sequences, rather than an inference from a long picture.

For $r = 0, \dots, 4$, use respectively the triples

$$(a_r, b_r, c_r) = (2,3,5),\ (7,11,13),\ (17,19,23),\\ (29,31,37),\ (41,43,47).$$

On the unwrapped supports of Supplement S.2 put

$$\theta_r(u,v) = \sqrt{a_r}\,u^2 + \sqrt{b_r}\,uv + \sqrt{c_r}\,v^2,\\ \varepsilon_r(u,v) = \lfloor 2\theta_r(u,v) \rfloor \pmod 2.$$

Equivalently, the bit is one when the fractional part of θ_r is at least $1/2$. Let C_5 be the resulting colouring. It is legal by Theorem 6. The value at any face is computable from its oriented residue and a single control bit. The companion program evaluates that bit with integer-square-root interval bounds, refined until the floor is fixed; it does not rely on floating-point decisions near half-integers.

5.1 Six ways to read a line

Three families read rows parallel to tile edges, and three read altitude lines. A line is the set of faces with the indicated key. Faces are read in their geometric order along it; reversal does not change eventual periodicity.

Family	Line key on $U(i,j)$	Line key on $D(i,j)$	Cell direction
A	j	j	$(1,0)$
B	i	i	$(0,1)$
C	$i+j$	$i+j+1$	$(1,-1)$
D	$2i+j+1$	$2i+j+2$	$(-1,2)$
E	$j-i$	$j-i$	$(1,1)$
F	$i+2j+1$	$i+2j+2$	$(2,-1)$

An eventually periodic ray has a positive period p and a starting position beyond which its raw-colour letters agree with their p-step translates. Testing these rays is a different question from testing a global translation period. A nonperiodic plane may still contain many periodic lines, and a finite defect in a periodic plane leaves every sufficiently remote ray periodic. The theorem below establishes both properties separately.

Theorem 7 (the quadratic construction). The colouring C_5 has no nonzero translation period. Neither ray of any line in any of the six families A–F is eventually periodic.

Lean status. Complete for the specified quadratic colouring and both numbered nonperiodicity assertions. The formal proof uses irrational first-difference density and a three-value argument in place of the stronger Weyl argument below. The adjacent mismatch-density value and evaluator termination are not formalised. See S.12.

Proof. First consider a line. The support certificate shows that every one of the 72 classes consisting of a family and a line key modulo 12 meets at least one support. Fix such an intersection. Period translates of its face lie on the given line at positions $\ell_0 \pm 24t$, and their component addresses have the form

$$(u_0 + \alpha t, v_0 + \beta t), \qquad t \in \mathbb{Z},$$

where (α, β) is that family's cell direction. The existence of an intersection on every line in the class follows by translating the support in the transverse direction as needed. The line keys have primitive integer coefficients, so every multiple-of-12 change of key is available.

At those faces the background colour is constant and the bit is readable. Its phase is a quadratic polynomial $Q(t)$ whose leading coefficient is

$$\lambda = \alpha^2\sqrt{a_r} + \alpha\beta\sqrt{b_r} + \beta^2\sqrt{c_r}.$$

The numbers $1, \sqrt{a_r}, \sqrt{b_r}, \sqrt{c_r}$ are linearly independent over $\mathbb{Q}$, and $(\alpha, \beta) \neq (0,0)$. Thus λ is irrational. For any integer $q > 0$, Weyl's polynomial equidistribution theorem gives joint equidistribution of $(Q(t), Q(t+q))$ modulo one: a nonzero integer combination has irrational quadratic coefficient unless its two multipliers sum to zero, in which case it has nonzero irrational linear coefficient $\pm 2q\lambda$ times an integer.

The two half-interval bits therefore differ for a set of t of density $1/2$. This holds on both tails, since replacing t by $-t$ preserves the argument. If a raw-colour ray had eventual period

p, its sampled subsequence at spacing 24 would be eventually p-periodic in t: a p-step move in the subsequence is a $24p$-step move in the original word. Its readable bits contradict the conclusion with $q = p$.

For global periods, the block field preserved by every switch has period lattice exactly $12\mathbb{Z}^2$. This is checked by testing all 144 possible residues of a translation on its 288-face cell. Any period of C_5 must therefore be $(12p, 12q)$. Reading the bit of family J_0 at a marker would then make (p, q) a period of its quadratic control field. For $(p, q) \neq (0, 0)$, the pair

$$(\theta_0(u, v), \theta_0(u + p, v + q))$$

is jointly equidistributed over growing integer rectangles. In the character test, an uncancelled quadratic part has an irrational coefficient; if it cancels, the difference has linear coefficients $2p\sqrt{a_0} + q\sqrt{b_0}$ and $p\sqrt{b_0} + 2q\sqrt{c_0}$, at least one irrational. The two bits differ on a positive-density set, contradicting a period. This proves the global assertion. □

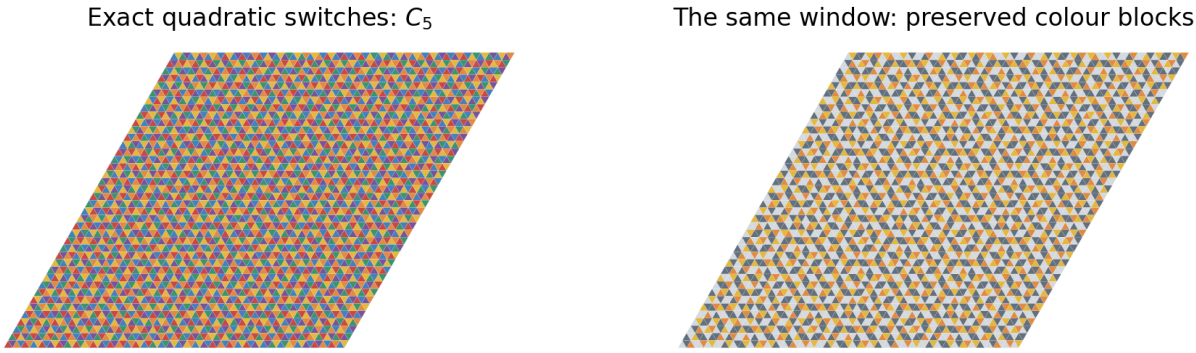


Figure 8: The exact quadratic colouring on a 40-by-40 anchor window, and the same faces recoloured by the four preserved blocks. The right panel is periodic even though no raw-colour ray in A–F has an eventually periodic tail. These are two readings of the same generated data.

5.2 Infinite proof, finite implementation

The Lean companion proves the same numbered conclusion by a different analytic route. Along every nonzero arithmetic direction the leading quadratic coefficient is irrational. Density of its linear first differences supplies a triple of successive values that contradicts confinement of the entire tail to one closed semicircle modulo one. Consequently the corresponding half-interval bit cannot be constant on an arithmetic tail. Applying this to a putative periodic subsequence, together with the exact ordered-line sampling and block-period reduction, excludes every proposed ray or translation period. This argument establishes the required nonperiodicity without proving the stronger mismatch density 1/2 used above.

The equidistribution argument applies to every proposed period, every line in the specified families and both infinite tails. It is not a claim about all possible directions through a lattice colouring, and it does not mean that finite words never repeat. Repetition of finite patches is compatible with the conclusion.

Only three finite facts enter the proof: legal complete supports, coverage of the 72 line classes, and the exact period lattice of the preserved block field. Supplement S.3 supplies one support intersection for each line class and explains the integer arithmetic used to draw the example. The program verifies those finite hypotheses; the polynomial argument supplies the quantifiers over arbitrarily long rays and arbitrarily large periods.

Theorem 6 and Theorem 7 also answer different counting questions. The first varies arbitrary bits at positive spatial density. The second selects a single low-description-complexity assignment of those bits. No positive orbit-entropy assertion about that one assignment is needed or made. The next section makes an equally consequential distinction between a nonperiodic raw word and a nonperiodic geometric arrangement of a selected colour.

6. Reading the colour classes exposes hidden order

Fix a colour k. Its triangles partition the lattice vertices into triples: each vertex lies in exactly one triangle of colour k. A picture of all six face colours superimposes six such covers. A phase reading separates one cover and compares it with periodic reference covers, making long aligned regions easier to see.

At a vertex v, let $f_k(v)$ be its unique incident face of colour k. For a reference labelling H of faces, define the vertex field

$$\Phi_{k,H}(v) = H(f_k(v)).$$

We use four reference labellings, each assigning all six labels once around every vertex. All formulas in the first three columns are modulo six; in the final column $(i - j) \bmod 3$ is taken before the displayed addition.

Face	P_0	P_1	P_2	R
$U(i,j)$	$i + 2j$	$i - j - 1$	$-2i - j - 2$	$(i - j) \bmod 3$
$D(i,j)$	$i + 2j + 4$	$i - j + 2$	$-2i - j$	$3 + (i - j) \bmod 3$

The first three are rotations of a published mod-six trimer labelling. The fourth distinguishes face orientation and the regular three-class address. Six raw colours and four references give 24 named fields. The display hues in a phase plot represent these reference labels, not the original face colours.

6.1 An exact obstruction inside the entropy construction

Proposition 8 (fixed-colour phase order). For every choice of the switch bits in Theorem 6, all eight fields $\Phi_{k,H}$ with $k \in \{2, 5\}$ and $H \in \{P_0, P_1, P_2, R\}$ are $12\mathbb{Z}^2$-periodic. For colour 2 in P_0, the label counts on a 144-vertex period cell are

$$144\pi = (24, 12, 24, 24, 24, 36).$$

Lean status. Complete, uniformly over every switch assignment, for the eight periodic fields and the printed histogram. See S.12.

Proof. The switched pairs are $\{0,1\}$ and $\{3,4\}$. No switch changes any face of colour 2 or 5. Their selected incident faces therefore satisfy

$$f_k(v+12u) = f_k(v) + 12u \qquad (u, v \in \mathbb{Z}^2).$$

Every reference labelling has coordinate periods dividing 12. Applying it gives $\Phi_{k,H}(v+12u) = \Phi_{k,H}(v)$. Counting the selected colour-2 face at the 144 vertices of the background cell gives the displayed histogram. □

For any positive integer m, the agreement of the colour-2 P_0 field with its translate by $(12m, 0)$ is exactly one. Its independent-marginal baseline would be

$$\sum_{s=0}^{5} \pi_s^2 = \frac{13}{72}.$$

The discrepancy is exact and persists at arbitrarily distant lags. Increasing the number of sampled faces cannot remove it. Nor can changing the five control fields: every possible choice leaves this obstruction intact.

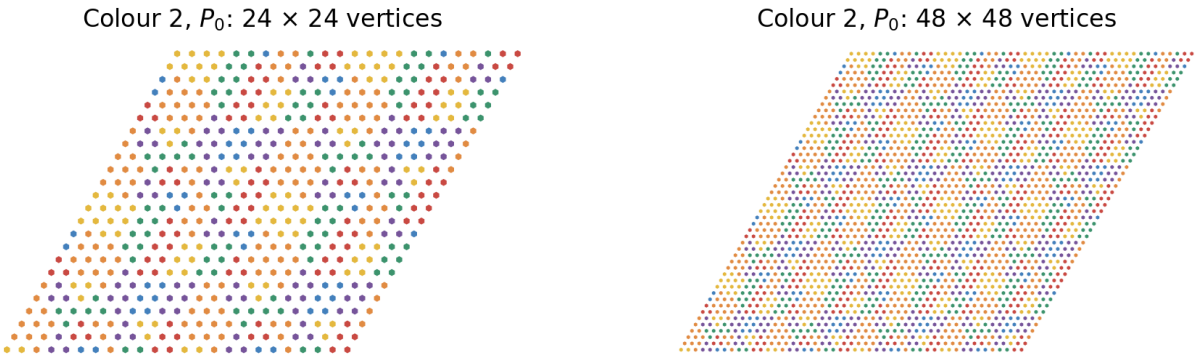


Figure 9: The colour-2 P_0 field of C_5 at two window sizes. Each plotted point is a vertex; its hue is a phase label. The repetitions come from the fixed colour class, not from numerical rounding or from truncating the quadratic controls.

This is the central separation. The space contains exponentially many plane-admissible patterns. A specified member avoids every raw-colour periodic ray in the six canonical families. Nevertheless its selected colour classes retain exact periodicity. None of those statements contradicts either of the others.

6.2 Small domains can coexist with complete local rigidity

The previous example retains periodic phase fields because two raw colours never move. A different control shows why bounding phase-domain sizes alone is also insufficient. Its phase

pictures break into small islands, but the entire pattern repeats, and no finite legal change is possible.

Theorem 9 (a frozen bounded-domain control). There is an explicit legal periodic colouring C_{fr} whose equal-phase components have at most 18 vertices in every one of the 24 named fields. Nevertheless no nonempty finite recolouring, using any subset of the six colours and keeping the exterior fixed, is legal.

Lean status. Complete for the explicit witness, including every whole-plane phase component and every finite legal edit. See S.12.

The field is generated from the repeating row of pairs

$$((0,4),(2,5),(0,4),(3,0),(4,3))^{\mathbb{Z}}$$

by a reversible reflection formula printed in S.11. Its raw colours and all phase labels have periods $(30,0)$ and $(0,60)$. A translation-labelled quotient checks every domain: a nonzero accumulated translation around a quotient cycle would produce an infinite lifted component. Here every such translation vanishes, and the largest component has 18 vertices. This establishes a whole-plane bound using periodicity and the lift argument, not by measuring islands cut off at a finite window.

For rigidity, fix one complete row of the control. The rainbow stars fix three-colour complement sets for the row below. After normalising by parity, the possible lower-row pairs form one of two 30-state directed graphs. Repeatedly remove states with no predecessor or no successor. In each graph the surviving state counts are

$$30 \longrightarrow 16 \longrightarrow 10 \longrightarrow 7 \longrightarrow 5.$$

The remaining five states form precisely the canonical row cycle, one state per column residue. Any bi-infinite lower row must survive every pruning round, so that row is forced. If a finite recolouring existed, take its highest changed row. The unchanged row immediately above would force it to be unchanged too, a contradiction. S.11 gives the graphs, signed-coordinate formula and quotient proof.

This supplies a sharp contrast with the finite switch components of Section 4. A bounded-region resampling procedure started at C_{fr} cannot make a first nontrivial legal move, even if it can use all six colours. The theorem concerns this particular background, not every periodic colouring or every repair method. Its bounded domains satisfy the tail clause of M3, but its exact periods violate the translation clause. It is a control explaining a missing condition, not a construction meeting M3.

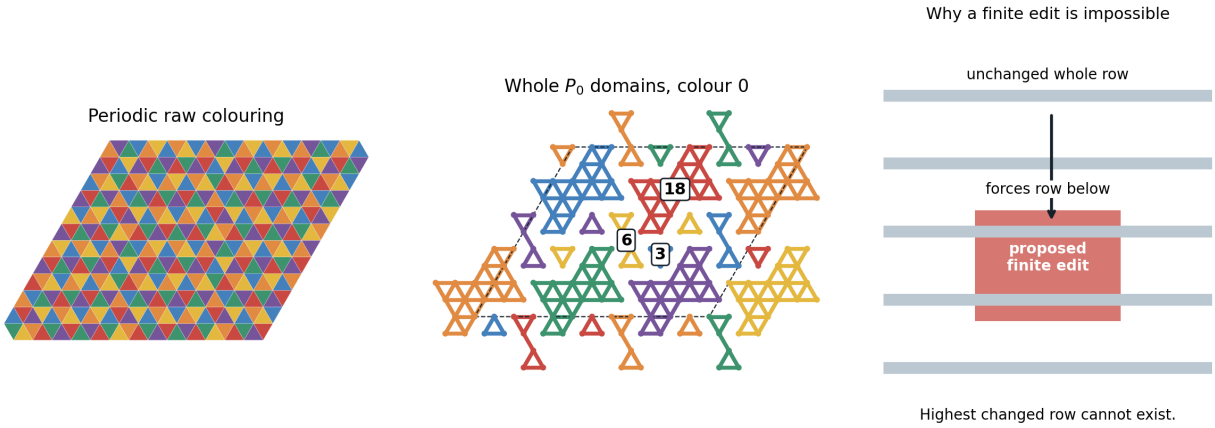
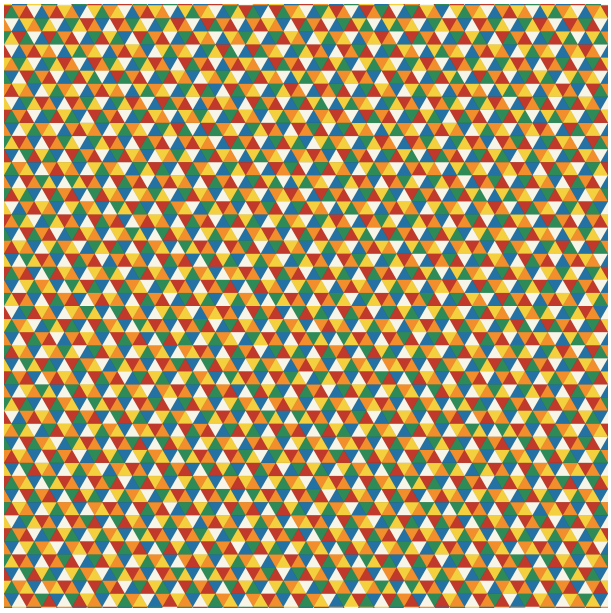


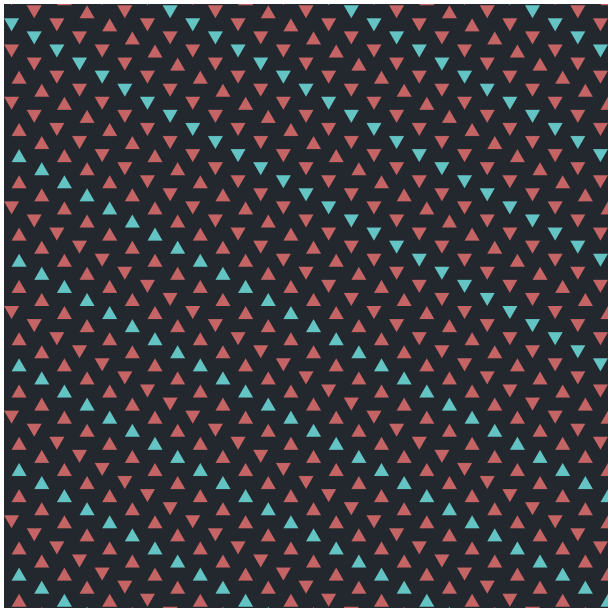
Figure 10: The frozen control in raw colours and in its colour-0 P_0 phase field. The phase panel shows whole finite components intersecting the marked window, including their exterior vertices; highlighted components have 3, 6 and 18 vertices. Small islands coexist with exact spatial repetition. The final panel shows why a changed region with fixed exterior cannot have a highest changed row.

6.3 What the finite experiments contribute

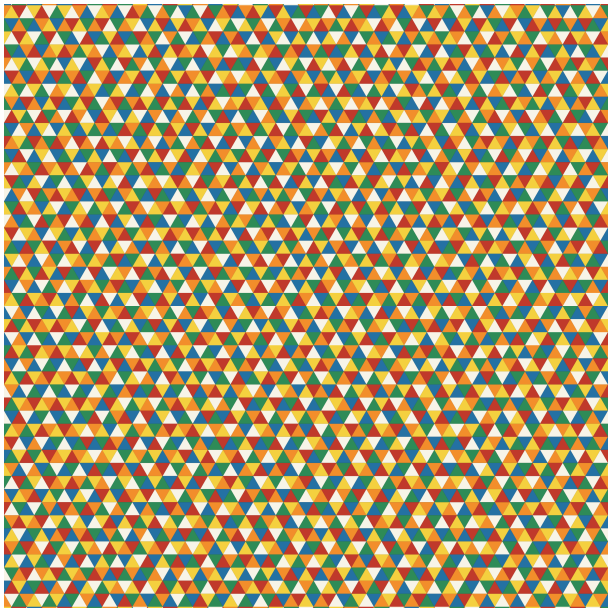
The finite experiments supply useful contrasts and reveal where a proposed diagnostic is insensitive. They do not certify an infinite absence of order. The four panels below use the same origin-centred 40-by-40-unit window. The formula is an earlier three-family infinite construction; it is **not** the five-family colouring C_5. The grown example is a saved finite radius-100 extension of a repaired hand sample. Its exterior was chosen by a constraint solver.



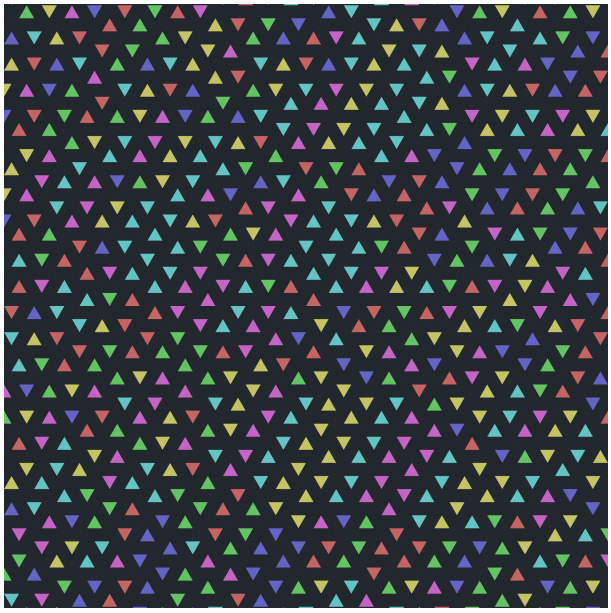
(a) Earlier three-family formula: raw colours



(b) Earlier three-family formula: colour-0 phases



(c) Saved finite extension: raw colours



(d) Saved finite extension: colour-0 phases

Figure 11: Matched-window views of the recorded finite configurations. The phase panels use P_0 and colour 0. Their hues describe one cover's phases. A visual difference at this scale supplies neither an all-field conclusion nor a whole-plane extension theorem.

The distinction between a construction and a sample matters in interpreting such images. The formula is known on the entire plane and can be interrogated at any scale. The saved extension is known only on its finite domain. A solver finding that domain is not a uniform sampler of valid colourings, and the proportions of its observed outcomes are not probabilities under a specified invariant law.

The recurring switch in Figure 6 is another useful experimental connection: a finite grown witness contains, up to a lattice translation, a component geometry already used in the entropy argument. That observation identifies a mechanism shared by the constructed and grown examples. It does not show that the grown witness contains an independent positive-density packing of those mechanisms, or that its present boundary permits continuation.

There is no need to choose a single scalar "randomness score" to see the limitation. A phase histogram can look balanced while the labels repeat at a fixed displacement; a raw-colour string can avoid periodic tails while a selected phase field is periodic. The appropriate next question must say which fields it observes, which translations it tests, and how it treats large domains. Section 8 gives one such target, with existence left open.

7. Two routes to disorder that are obstructed

The previous construction fails to remove phase order for an explicit internal reason. There are also restrictions that apply beyond this construction. We include two because they bear directly on natural alternatives: choosing colours by a bounded local rule from independent random data, and reading a quasiperiodic circle orbit through finitely many intervals.

7.1 From one colour class to an order-three ribbon tiling

A ribbon of order three is a connected union of three square cells, reached by two north-or-east steps. Each cell lies in a different diagonal class $x + y \pmod 3$. The two L-shaped ribbons EN and NE are among the allowed order-three ribbons.

Fix one raw colour and regard each vertex of its triangular-trimer cover as a cell anchor. Apply the lattice automorphism

$$A(i, j) = (i + j, j), \qquad A^{-1}(x, y) = (x - y, y).$$

Then the two face orientations become

$$\begin{aligned} A(U(i, j)) &= A(i, j) + \{(0, 0), (1, 0), (1, 1)\}, \\ A(D(i, j)) &= A(i, j) + (1, 0) + \{(0, 0), (0, 1), (1, 1)\}. \end{aligned}$$

Attach a unit square to each transformed vertex. Since the selected triangles partition the original vertices, the resulting L-ribbons partition all square cells. Nothing is omitted or counted twice.

The map is local: to find the ribbon containing the cell anchored at $A(v)$, inspect the six faces incident to v and select the one of the chosen colour. The output records its ribbon shape and the position of that cell in the ribbon. After reindexing by A, the input anchors needed lie in

$$K = \{(0, 0), (-1, 0), (-1, -1), (-2, -1)\}.$$

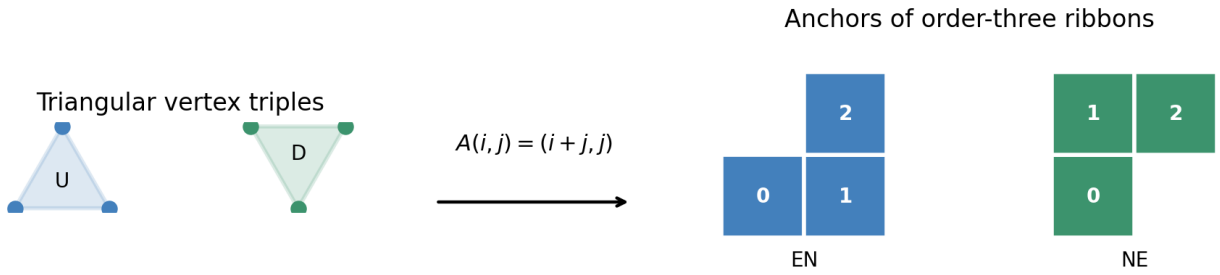


Figure 12: The integer coordinate change turns the vertices of an up or down triangle into the cell anchors of an EN or NE ribbon. Digits in the ribbon panels are the three successive diagonal classes relative to the ribbon's anchor. The map acts on a whole trimer cover, not just on the illustrated pair.

This is a finite-memory map equivariant under the full translation group. It is not merely equivariant under a periodic sublattice. Supplement S.5 prints all six decoding cases.

Theorem 10 (ribbon-factor obstruction). The space X contains no nonempty strongly irreducible subshift and supports no stationary finitely dependent probability law. Consequently a bounded-radius translation-equivariant function of independent identically distributed data cannot produce a legal colouring of X almost surely.

Lean status. Complete for all three conclusions, including the literal measurable bounded-radius iid map. The formal companion proves the required order-three ribbon obstruction from explicit integer edge weights, rectangle identities and periodic diagonal configurations; it does not assume the external nonexistence theorem. See S.5 and S.12.

Proof. A subshift is strongly irreducible if there is a uniform finite separation beyond which any two individually occurring finite patterns can be jointly realised. Strong irreducibility passes to finite-memory factors: lift the two output patterns to input patterns on their enlarged supports, use the separation property there, and project. The coordinate change A and its inverse distort distances by only a constant factor. Thus any nonempty strongly irreducible subshift of X would give one inside the order-three ribbon tiling shift.

Chandgotia and Thorat [4, Theorem 2.35, version 2] rule out nonempty strongly irreducible subshifts of ribbon tiling shifts of order at least two. Their Corollary 8.9 also rules out stationary finitely dependent laws on these ribbon shifts. A finite-memory equivariant factor of a stationary finitely dependent law retains both properties, so the same map proves the second assertion. Finally a bounded-radius factor of iid data is finitely dependent: outputs on sets farther apart than twice the radius inspect disjoint independent inputs. □

The new ingredient in this application is the explicit triangular-cover factor and its translation behaviour. The argument above uses the external ribbon theorem with its original attribution. The formal companion additionally proves the order-three specialization needed here by a direct height and periodicity argument, outlined in S.5; it does not formalise the general all-order theorem. The conclusion excludes exact independence beyond a fixed finite distance.

It does not exclude ordinary mixing with correlations tending to zero, a finitary factor with unbounded coding radius, or a deterministic M3 colouring.

There is no conflict with Theorem 6. Independent switches are installed on a selected periodic background and a selected macrocell grid. This is not a bounded-radius construction equivariant under every lattice translation from iid data alone. Averaging over background origins can create a stationary law, but the shared origin can retain long-distance dependence. The entropy argument never requires that dependence to disappear.

7.2 Finite-interval readouts of one circle

A different candidate is quasiperiodic from the start. Fix real α, β, θ and two circle functions f_U, f_D taking values in six colours. Suppose each function is constant on the intervals of some finite partition, and define

$$C_U(i,j) = f_U(\theta + i\alpha + j\beta),$$
$$C_D(i,j) = f_D(\theta + i\alpha + j\beta),$$

with arguments modulo one. There is no additional periodic address class or hidden state.

Theorem 11 (the specified circle model). If $1, \alpha, \beta$ are linearly independent over $\mathbb{Q}$, the model cannot be legal. If they are dependent, it has a nonzero raw translation period. In particular no model of this form satisfies M3.

Lean status. Complete for the literal finite-interval model, including arbitrary endpoint values, both rational-dependence cases and failure of M3. The formal proof implements the finite atomic derivative argument with actual one-sided jumps of floor expressions. See S.12.

Proof of the independent case. Fix one colour and let u, d be the indicator functions for that colour in the two orientations. Write $Xh(x) = h(x - \alpha)$ and $Yh(x) = h(x - \beta)$. Exact coverage of each vertex requires

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

The orbit is dense, so the identity holds almost everywhere: a failure on a constancy interval would be sampled. Differentiating in the sense of distributions gives an identity between finite integer atomic measures. Translation by α, β acts freely on each orbit of atoms. On such an orbit, finite atomic measures can therefore be represented by integer Laurent polynomials in independent X, Y.

The resulting equation $\mathcal{A}U + \mathcal{B}V = 0$ implies $U = \mathcal{B}W$, $V = -\mathcal{A}W$ with integer Laurent W. One way to see this is to clear negative powers and substitute $Y = -1 - X$; then $\mathcal{B} = -(X^2 + X + 1)$ and monic division shows that $\mathcal{A}$ divides V. Reassembling the orbits gives an integer atomic measure w with $Du = \mathcal{B}w$ and $Dd = -\mathcal{A}w$.

Its total mass is zero, so it has an integer-valued periodic step primitive h. Consequently $u = \mathcal{B}h + a$ and $d = -\mathcal{A}h + b$ almost everywhere for integers a, b. Substitution yields

$$1 = \mathcal{A}u + \mathcal{B}d = 3(a + b),$$

which is impossible. Supplement S.6 supplies the divisibility and primitive steps in full.

In the dependent case, clearing a rational relation gives a nonzero integer vector p with $p_1\alpha + p_2\beta = 0 \pmod 1$. It is an exact raw period, including the assigned endpoint values. Its sixfold multiple also preserves every reference labelling, hence is a period of all phase fields. This violates the positive-mismatch clause of M3. □

The two cases fail for different reasons. Irrational independence conflicts with legality itself; rational dependence leaves a phase period. The restriction to finite-interval readouts is crucial in the first argument, because their derivatives are finite integer atomic measures. Arbitrary measurable partitions, additional finite coordinate factors and higher-dimensional torus codings are outside the theorem. It would be incorrect to infer that quasiperiodic M3 constructions are impossible.

8. The remaining existence question

The results identify several freedoms and several restrictions without resolving whether a legal plane can have the desired phase-domain behaviour. To make that gap precise, we retain the working target **M3**, defined in terms of the same 24 fields as Section 6. The name is a project label; the quantifiers below are its mathematical content.

For a field $F = \Phi_{k,H}$, connect equal-labelled vertices by nearest-neighbour lattice edges. Let $C_F(v)$ be the whole-plane connected component containing v. Use the hexagonal windows

$$B_r = \{(i, j) : \max(|i|, |j|, |i + j|) \le r\}, \qquad N_r = 1 + 3r(r + 1).$$

Define the large-component mass and translation agreement by

$$M_{r,F}(s) = \frac{|\{v \in B_r : |C_F(v)| \ge s\}|}{N_r},$$
$$A_{r,F}(t) = \frac{|\{v \in B_r : F(v + t) = F(v)\}|}{N_r}.$$

An infinite component counts at every finite threshold. The second endpoint $v + t$ is read in the whole field, even when outside the window. Write $p_{r,F}(\ell)$ for the fraction of label ℓ in B_r. For a coset q of a finite-index subgroup $\Lambda \le \mathbb{Z}^2$, let $p_{r,F,q}(\ell)$ be the corresponding fraction in $B_r \cap q$, normalized by $|B_r \cap q|$.

Definition (M3). A legal whole-plane colouring satisfies M3 if each of its 24 named fields satisfies:

$$
\begin{aligned}
\mathrm{T}: &\quad \lim_{s\to\infty} \limsup_{r\to\infty} M_{r,F}(s) = 0;\\
\mathrm{A}: &\quad \text{for every } t \neq 0, \quad \liminf_{r\to\infty}(1 - A_{r,F}(t)) > 0;\\
\mathrm{P}: &\quad p_{r,F,q}(\ell) - p_{r,F}(\ell) \longrightarrow 0
\end{aligned}
$$

for every fixed finite-index subgroup, coset and label in P. The statements also use $B_r + z$ for any fixed centre z; no uniformity over moving centres is required.

T says that increasingly large components carry a vanishing fraction of vertices. It does not require every component to be finite: infinite components are permitted only when their union has zero upper density. A requires positive lower mismatch at each fixed nonzero translation, but that mismatch may become arbitrarily small along a sequence of longer displacements. P prevents a persistent single-label bias on a periodic address class. Each subgroup and label is fixed before the window grows.

This target allows quasiperiodic near-returns and differs from the stronger historical requirements of asymptotic decorrelation and nonperiodicity of every phase ray. It is a declared working interpretation of the desired appearance, not a proved weakest characterisation of it. Neither a favourable picture nor a finite collection of agreement and component measurements proves its three limiting clauses.

Open problem. Does there exist a colouring in X satisfying M3?

Proposition 8 rules out every member of the five-family switch construction, because some named fields have exact periods. Theorem 11 rules out the specified same-circle model. Theorem 10 restricts stationary probabilistic construction methods, while leaving weaker forms of dependence available. These are obstructions to stated routes, not an impossibility theorem for M3.

A future positive answer needs one concrete whole-plane colouring and proofs of T, A and P for all its named fields. For a hierarchical proposal, finite local legality must be accompanied by actual support and a whole-plane exhaustion; large finite patches alone do not supply either. The tail clause needs control of the final components, including components that join beyond the observed window. A and P require their full families of fixed translations and periodic address classes. This article supplies no candidate meeting those obligations.

There are smaller questions that are also left open: the exact entropy of X and a packing improving on five independent bits per 288 faces. None is needed for the theorems proved here. The current contribution is complete with those questions open: it exhibits extensive freedom and explicit directional nonperiodicity, then identifies exactly which hidden order they retain and which alternative constructions are obstructed.

8.1 Finite progress towards the M3 obligations

The following deductions make parts of this construction problem more precise. They are **not Lean verified**; S.13 gives their proofs.

First, for any finite-alphabet field let $\mathcal{V}(F)$ be the compact set of subsequential limits of its empirical shift measures. If E_∞ is the event that the origin belongs to an infinite equal-phase component, then

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

Thus the component-tail clause can be checked through all empirical limit laws without assuming that finite-pattern frequencies converge. This reduces the proof obligation; it does not identify those laws for a new colouring.

Second, in a fixed legal H_r, let $G_{r,F}(s)$ be the mass in phase components smaller than s that do not meet the vertex frontier, and put $U_{r,F}(s) = N_r - G_{r,F}(s)$. These sealed components cannot later merge. A two-ring extension therefore obeys the exact balance

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

Decreasing charged-mass bounds, persistent positive mismatch floors and increasingly accurate residue-class label counts form a sufficient finite certificate scheme for M3. A stored second annulus can certify an intermediate step's return to the required endpoint conditions. The difficult missing premise is a realised starting state and a continuation rule meeting those conditions indefinitely.

The same accounting gives necessary finite restrictions. Every sealed named-field component is a union of whole raw colour triangles, so $G_{r,F}(s) \equiv 0 \pmod 3$ and $U_{r,F}(s) \equiv 1 \pmod 3$. Counting incident colour faces on the inner frontier forces, for some colour and all four of its reference fields, $U_{r,F}(s) \geq 9r - 2$. For levels $j \geq 1$, simultaneous guards $U_{r,F}(s) \leq 2^{-j} N_r$ therefore require

$$r \geq 3 \cdot 2^j - 1, \qquad j \geq 1.$$

At $j = 0$ the mass guard is automatic and gives no radius bound. For $j \geq 1$ the displayed bound is necessary at each active publication radius; no sharpness or feasibility is claimed. S.13 also explains why a cycle of suitable intermediate families is sufficient for block return, and how reflection of a saved finite trap refutes one proposed oriented family. None of these deductions closes the M3 existence question.

9. Reproducibility and evidence

The technical supplement gives the supporting proofs, coordinate tables, finite nonextension certificates, inventory and transfer correspondences, switch supports, phase formulas, ribbon decoder and circle-coding argument. The supplementary archive accompanying this preprint on Zenodo contains editable sources, figures, exact assignments, computational certificates and formal proofs, together with instructions for rebuilding and checking them. File references in this article and the supplement are relative to that archive, which can be used without access to the development repository.

The exploratory web companion *Freedom & Hidden Order* (version 1.1.1) is available at freedom-and-hidden-order.vercel.app. It provides interactive raw-colour and phase views of selected examples, with exports identifying their source and display settings. A separately labelled radius-100 preview illustrates forthcoming work outside this article's reported examples. Its appearance and finite measurements do not establish M3 or infinite extendability. The mathematical and Lean-verification claims remain those stated in this article and supplement.

The rule diagram, recurring-switch drawing and matched-window panels depict recorded configurations. The other figures are generated from the stated coordinates, formulas and finite data. Schematic extension-tree branches are distinguished from counted configurations. A finite extension is shown only as a visual comparison; its boundary is not asserted to extend further.

The seed classification and numerical two-ring total have exhaustive computer-assisted evidence described in S.9. Their Lean coverage is given below. The finite-premise checkers use direct coordinate arithmetic and small matrix and graph calculations. Recorded exhaustive computations retain their dates, input data and verifier outputs, separately from the document build. Section 1.3 gives source attribution; the result statements and S.12 distinguish written proofs, external results, finite certificates and Lean verification. No infinite M3 theorem is claimed.

9.1 Scope of the Lean verification

The original **11 numbered results comprise eight theorems and three propositions**. Ten have complete Lean proofs of their printed numbered assertions; Theorem 3 has checked ingredients and two explicit outstanding obligations. The added written deductions in Section 8.1 and S.13 are not Lean verified. This count does not include every assertion or formula in the article.

Result	Lean coverage	Principal outstanding obligation
Proposition 1	Complete	None for the numbered nonextension assertion
Theorem 2	Complete	None for the multiset equivalence
Theorem 3	Partial	Numerical weighted trace and exhaustive coverage from the 50,720 certified representatives
Theorem 4	Complete	None for the numbered transfer and operation claims
Proposition 5	Complete	None for the two specified traps
Theorem 6	Complete	None for the switch family, entropy limit and bound

Result	Lean coverage	Principal outstanding obligation
Theorem 7	Complete	Alternate proof; adjacent mismatch density excluded
Proposition 8	Complete	None for the periodicity and histogram assertions
Theorem 9	Complete	None for the explicit frozen witness
Theorem 10	Complete	None; the order-three obstruction is proved directly
Theorem 11	Complete	None for the stated finite-interval model

S.12 explains the mathematical correspondence, remaining hypotheses and reproduction procedure. All local proof modules are compiled from source with Lean 4.24.0 at trust level zero. Dependency reports admit only the standard logical foundations actually used: propositional extensionality, quotient soundness and, where needed, classical choice. No missing mathematical argument is installed as an extra axiom. The analytic, algebraic, topological and probabilistic modules additionally use a pinned Mathlib release. Source hashes and deliberately false controls accompany the build receipts. Earlier receipts retain their original times and manuscript snapshots; the final annotations are bound separately, and all eleven numbered statement texts are unchanged.

The formal contribution is verification of the specified mathematics. It does not establish novelty, historical priority, the correctness of an unformalised external proof, or the infinite M3 target. The source attribution in Section 1.3 remains essential. Theorem 3 remains partially formalised, and the deductions in Section 8.1 and S.13 have written proofs only.

Acknowledgements and AI use

The author acknowledges the Lean and Mathlib communities for the proof-assistant infrastructure. OpenAI's GPT-6 Astra (model identifier `gpt-6-astra`), used through OpenAI Codex and referred to as Astra during development, assisted with mathematical exploration, proof drafts, programming, Lean formalisation, manuscript preparation and review. Anthropic's Claude Fable 5.1 (model identifier `claude-fable-5-1`), referred to as Fable, provided additional mathematical and manuscript review. Both are AI models. The author is responsible for the final manuscript. AI review is not independent human peer review; claims of Lean verification are limited to the scope stated in Section 9.1 and S.12.

References

[1] Aidan Quigley. "Arranging coloured triangles in hexagons so that they don't touch one another." *Mathematics Stack Exchange*, question 4221077, 10 August 2021. Original question.

[2] J. F. Rigby. "Perfect precise colourings of triangular tilings, and hyperbolic patchwork." *Symmetry: Culture and Science* **8** (1997), 265–288. Scanned volume. The strip discussion is on p. 266 and the random-colouring remark on p. 283.

[3] A. Verberkmoes and B. Nienhuis. "Triangular trimers on the triangular lattice: an exact solution." *Physical Review Letters* **83** (1999), 3986–3989. doi:10.1103/PhysRevLett.83.3986; arXiv:cond-mat/9904343.

[4] Nishant Chandgotia and Aditya Thorat. *Finitely Dependent Processes on Subshifts*. arXiv:2605.02226v2, 29 May 2026.