Cut tires form a tree (under depth nesting)

The claim

Proposition (Cut tires form a forest)**.** *For each side i of a 6-edge cut of G', the cut tires of G'_i, parameterised by pairs (d, f) with $d \geq 1$ and f a face of H_d, form a* forest *under the parent–child relation*

$$\text{parent}\big(T_{d+1}^{(f')}\big) := T_d^{(f)}$$

where f is the unique face of H_d in whose planar interior f' lies in the inherited embedding of G'_i.

The forest's roots are the cut tires at depth 1 (one per face of H_1); their "virtual parent" is the cut C itself.

Sketch. H_{d+1} is a subgraph of G'_i with the inherited planar embedding. Each face of H_{d+1} is a maximal connected open region of $|\Pi| \setminus E(H_{d+1})$ in the plane.

In particular, every face of H_{d+1} lies inside some face of H_d (since H_d has fewer edges and so larger faces). "Lies inside" means: the open face region of H_{d+1} is a subset of an open face region of H_d. This containment is unique because the faces of H_d partition $|\Pi| \setminus E(H_d)$.

Hence parent is well-defined and unique. No face of H_{d+1} is its own parent (because $d+1 > d$). The relation defines a forest.

The roots are the depth-1 cut tires. Their "virtual parent" is the depth-0 pendant configuration, i.e. the cut C itself. □

Why this matters for the chain half

Chain pigeonhole asks whether the per-tire S_3-orbit structure composes coherently through the chain. With a tree structure on the cut tires, this becomes a **tree dynamic-programming problem**, not a general graph compatibility problem:

- Process tires from leaves to root.
- At each leaf: $\pi(T_{\text{leaf}})$ has known structure (e.g. S_3-orbits) from the per-tire half.
- Internal node $T_d^{(f)}$ combines:
  - Its own internal $\pi(T)$ structure.
  - Compatibility with each child $T_{d+1}^{(f')}$ via the bijection $\{\text{in spokes of } T_d^{(f)}\} \leftrightarrow \{\text{face boundary edges of } T_{d+1}^{(f')}$
- Root: $T_1^{(\cdot)}$ projects its out-spoke colours to $\sigma_i \in \mathcal{R}_i$.

Tree DP is well-understood: $|\mathcal{R}_i|$ can be computed exactly in linear time in the tree size (with size-$|\pi|$ tables at each node). Whether the resulting $\mathcal{R}_0$ and $\mathcal{R}_1$ intersect is a finite check at the cut.

The tree structure is also a **strong topological constraint** on the chain pigeonhole obstruction: any counterexample to chain pigeonhole at the cut must come from a tree-DP failure, which is much narrower than a general-graph obstruction.

Broader empirical sweep

Run on 7 test graphs (script: `tree_structure_sweep.py`; data: `tree_structure_sweep_data.txt`):

graph	$\lvert V\rvert$	$\lvert E\rvert$	# 6-edge cuts found	trees on both sides
HM #0	38	57	128	128/128
HM #1	38	57	127	127/127
HM #2	38	57	122	122/122
HM #3	38	57	123	123/123
HM #4	38	57	101	101/101
HM #5	38	57	97	97/97
Dodecahedron	20	30	45	45/45

Totals:

- 743 distinct 6-edge cuts examined.
- 1486 $(\text{graph}, \text{cut}, \text{side})$ triples tested.
- 11,477 cut tires examined.
- 0 **tree-structure failures** (no cycles in the parent–child relation under the vertex-overlap heuristic).

The data spans:

- The 6 Holton-McKay non-Hamiltonian 38-vertex cubic plane graphs (their duals are 21-vertex maximal planar graphs of minimal degree 4 and vertex-connectivity 3).
- The dodecahedron (20-vertex cubic plane graph, dual of the icosahedron, which is a 12-vertex 5-regular maximal planar graph with vertex-connectivity 5).

Although neither family is strictly "min degree 5 with vertex connectivity 6" (which is incompatible with the maximal-planar upper bound on average degree of $6 - 12/|V|$), the test covers duals of:

1. Several internally non-trivial maximal planar graphs (HM duals).
2. A min-degree-5 maximal planar graph (icosahedron).

This is broader than the typical chain pigeonhole test bed.

Empirical demonstration on Holton-McKay #0 (detailed)

G_1' side ($|S| = 28$, depths 0 to 7)

Two depth-1 roots:

- Root $(1,0)$: face length 12, no children (the outer "shell" of H_1).
- Root $(1,1)$: face length 4, with substantial subtree:
  - $(2,0)$ $|f| = 7$
    - $(3,0)$ $|f| = 2 \Rightarrow (4,0)$ $|f| = 4 \Rightarrow (5,0)$ $|f| = 14$
    - $(3,1)$ $|f| = 2 \Rightarrow (4,1)$ $|f| = 8 \Rightarrow (5,1)$ $|f| = 2 \Rightarrow (6,0)$ $|f| = 12 \Rightarrow (7,0)$ $|f| = 2$
    - $(3,2)$ $|f| = 2$
  - $(2,1)$ $|f| = 7$

G'_0 side ($|S| = 10$, depths 0 to 2)

Two depth-1 roots:

- Root $(1,0)$: face length 9, with one child $(2,0)$ ($|f| = 6$).
- Root $(1,1)$: face length 9, no children.

Caveats on the empirical parent identification

The empirical demonstration used a vertex-sharing heuristic to identify parents: a face f' of H_{d+1} shares vertices with a face f of H_d, and we picked the parent as the one with smallest face length. This gives ambiguous candidates in some cases (8 ambiguous cases observed in G'_1) because vertex sharing does not fully determine geometric containment.

A rigorous parent test would use *point-in-region* containment: pick a point in the open face of H_{d+1} (e.g., the centroid of its boundary walk), determine which face of H_d that point lies in (via the planar embedding's face structure). This always gives a unique answer.

The ambiguity in our empirical run doesn't reflect a violation of the proposition — it's an artifact of the heuristic. Despite the ambiguity, the resulting tree structure looked sensible in both G'_0 and G'_1.

Consequence: the chain half becomes tractable

With the tree structure established (or assumed), the chain half of the loose chain pigeonhole conjecture reduces to:

Reformulated chain half (tree DP form). For each leaf cut tire T_{leaf}, $\pi(T_{\text{leaf}})$ is non-empty and S_3-closed. Propagating bottom-up through the parent–child relation preserves S_3-closure and non-emptiness. At the root depth-1 tires, $\mathcal{R}_i$ is the join of the root tires' out-spoke projections. If $\mathcal{R}_i$ is S_3-closed and contains a full S_3-orbit on each side, then $\mathcal{R}_0 \cap \mathcal{R}_1 \neq \emptyset$ (containing a common orbit by S_3-equivariance).

The remaining questions:

1. Is non-emptiness preserved through parent-child propagation?
2. Is S_3-closure preserved? (Yes, by S_3-equivariance of the proper edge 3-colouring constraint.)
3. Does the join of root projections contain a full S_3-orbit?

Each of these is now a finite tree DP claim, much more tractable than the original "compose through the chain" formulation.

Next step

1. Prove Proposition rigorously using the point-in-region containment definition of parent.
2. Implement the tree DP empirically on the Holton-McKay graphs and confirm $\mathcal{R}_0 \cap \mathcal{R}_1 \neq \emptyset$ at the cut.
3. Attempt an analytical bound: $|\mathcal{R}_i| \geq somefunctionoftreesize$, ensuring $\mathcal{R}_0 \cap \mathcal{R}_1 \neq \emptyset$ in general.