Chain half of the loose conjecture: tree DP and where it gates

Recap

The loose chain pigeonhole conjecture (k$\geq$ 2 form, `cut_depth_label.tex`) has two halves:

Per-tire half. For every cut tire T with ≥ 2 in/out spokes total, the joint projection $\pi(T) \subseteq \{1,2,3\}^k$ is non-empty, S_3-closed, and contains a full S_3 orbit of size 6. Proven for spoke-only cut tires (T'_{ann} = simple cycle C_n, $n \geq 3$) via Prop 1.13 of `paper.tex`.

Chain half. Composing per-tire projections through the cut-tire forest (`cut_tire_tree_structure.tex`, rigorously proved) yields $\mathcal{R}_i \neq \emptyset$ on each side i of the cut, with $\mathcal{R}_0 \cap \mathcal{R}_1 \neq \emptyset$ containing a common S_3-orbit at the cut.

Tree DP formulation of the chain half

The high-side cut tires of G'_i form a forest (rigorously proved). Process tires bottom-up:

(1) Leaves. A leaf cut tire $T^{(f)}_{d_{\max}}$ has no children. Its "achievable projection" onto its out spokes (depth-($d_{\max}-1$) direction) is

$$A(T) := \pi_{\text{out}}(T),$$

the projection of the per-tire projection $\pi(T)$ onto the out spokes alone (in spokes are unconstrained at leaves). By per-tire, $|A(T)| \geq 6$ with full S_3-orbit.

(2) Internal nodes. For a cut tire $T_p = T^{(f)}_d$ with children $T_{c_1}, \ldots, T_{c_r}$, define:

$$A(T_p) := \left\{ \begin{array}{c} \sigma_{\text{out}}(T_p) : \exists\, \chi \text{ proper edge 3-coloring of } T_p \text{ such that} \\ \forall\, j : \chi|_{\text{in-spokes corresponding to } T_{c_j}} \in A(T_{c_j})' \end{array} \right\},$$

where $A(T_{c_j})'$ is the "transferred-back" achievable projection of child T_{c_j} onto the corresponding parent in-spoke positions.

The transfer back: an in spoke of T_p at parent boundary vertex v corresponds to a depth-($d+1$) edge e^{**} in G'_i. This e^{**} is a face-boundary edge of T_{c_j} (the unique child whose face contains v as a boundary endpoint of e^{**}). So the parent's in-spoke color at v = child's face-boundary-edge color at e^{**}.

For the projection $A(T_{c_j})$ to constrain parent's in spokes, we'd need to know how child's spoke colors determine cycle edge colors at specific positions. In the spoke-only case (Prop 1.13), each spoke color at u equals the "third color" = the unique color not appearing on the two cycle edges incident at u. So spoke colors determine the constraint set on cycle colors at each vertex.

(3) Roots and the cut. $T_1^{(\cdot)}$ are roots. Their out spokes are the pendant edges = the cut configuration σ_i. The achievable cut configurations:

$$\mathcal{R}_i := \bigcup_{\text{roots } T_1^{(f)}} A(T_1^{(f)}) \quad \text{(or restricted, depending on how root constraints compose).}$$

(4) Cross-cut. G' is properly 3-edge-colourable iff $\mathcal{R}_0 \cap \mathcal{R}_1 \neq \emptyset$ (under the bijection between the two sides' cut edges).

What's preserved through the tree DP

S_3-closure: preserved

Lemma (S_3-equivariance of tree DP)**.** *If every $A(T_{c_j})$ is S_3-closed (under diagonal action on colours), then $A(T_p)$ is S_3-closed.*

Proof. The proper-edge-coloring constraint at every vertex is preserved by S_3 acting on colours uniformly. Applying $\pi \in S_3$ to a valid χ for T_p gives another valid χ. Compatibility with children: parent's in spokes are uniformly π-shifted, and each child's $A(T_{c_j})$ is S_3-closed by hypothesis, so the shifted parent in spokes still hit $A(T_{c_j})$. Hence $\pi(\sigma) \in A(T_p)$ for any $\sigma \in A(T_p)$. □

By induction from leaves (where $A = \pi$ is S_3-closed by per-tire half), every $A(T)$ is S_3-closed. This is the easy half.

Non-emptiness: open, but constrained

Conjecture (Non-emptiness preservation)**.** *If every $A(T_{c_j})$ contains a full S_3-orbit of size 6, then $A(T_p)$ also contains a full S_3-orbit.*

This is the genuine open piece of the chain half.

Why it isn't trivial. Two S_3-closed subsets of $\{1,2,3\}^k$ can have empty intersection even if both contain S_3-orbits. Example: orbit of $(1,2,3)$ vs orbit of $(1,1,2)$ in $\{1,2,3\}^3$ are disjoint.

So the conjecture would require a structural reason that parent + children combined always have at least one common assignment with all 3 colours present (= a full S_3-orbit).

Why it's plausible. Empirical data from the partial- tire-dual chain pigeonhole (`tire_fiber_step2.tex`): 23/23 pairwise compatibility tests succeeded, with the intersections containing S_3-orbits and structured by "rainbow" or similar canonical orbits (`orbit_decomposition.tex`). These results suggest a structural reason for non-emptiness; we just don't have a clean proof.

What would close the proof. Show that for spoke-only cut tires, the per-tire projection $\pi(T)$ has the property: *for any specified colours on the in spokes that come from S_3-orbits, there is a compatible parent coloring.* Specifically:

Conjecture (Strong per-tire extendibility)**.** *Let T be a spoke-only cut tire with face boundary a simple cycle C_n ($n \geq 3$). For any $\sigma_{\text{in}} \in \{1,2,3\}^{n_{\text{in}}}$ such that σ_{in} lies in a non-trivial S_3-orbit (i.e. uses ≥ 2 colours and is in $\pi(T)$'s S_3-symmetric support), there exists a proper edge 3-coloring χ of T with $\chi|_{\text{in}-\text{spokes}} = \sigma_{\text{in}}$ and $\chi|_{\text{out}-\text{spokes}}$ a non-trivial S_3-orbit on the out-spoke side.*

If this conjecture holds, the chain DP preserves non-emptiness: for each non-trivial parent σ_{in} (= child's face boundary edges) provided by children, parent has a coloring with out spokes in a non-trivial S_3-orbit, hence $A(T_p)$ contains a full S_3-orbit.

Empirical next step

Cut-tire tree DP empirical test:

1. For each test graph (HM #0 through #5, dodecahedron, BuckyBall), build the cut tire forest on each side.
2. For each leaf, compute $A(T_{\text{leaf}})$.
3. Bottom-up propagate $A(\cdot)$ to roots.
4. Compare $\mathcal{R}_0 \cap \mathcal{R}_1$ at the cut.

This is the analogue of `tire_fiber_step2.tex` for the cut-tire setting. If empirically $\mathcal{R}_0 \cap \mathcal{R}_1 \neq \emptyset$ universally, the chain half is on firm empirical ground; the proof would still need Conjecture or a structural shortcut.

Caveat discovered empirically: cut tires are not spoke-only

A first-pass empirical test (`experiments/chain_dp_test.py`) on the dodecahedron and Holton–McKay #0 revealed a structural complication: *cut tires are not in general spoke-only.*

What goes wrong. H_d may have vertices of degree 3 (all three incident edges have depth d). In a cubic ambient graph G'_i, this happens when three depth-d edges meet at a single vertex. At such a vertex:

- There is no third edge available to be a spoke.
- The face boundary walk of H_d visits the vertex *twice* (as a branch point).

For example, in the dodecahedron, a 6-edge cut produces H_1 with 1 face whose boundary has length 20 but only 11 distinct vertices — 9 branch-point visits.

Consequence for Prop 1.13. The per-tire half (proven via Prop 1.13) covers *spoke-only* cut tires (face boundary = simple cycle C_n, $n \geq 3$). It does *not* cover branched cut tires.

Two ways forward.

1. *Restrict.* Identify which graphs G have the property that every cut tire of every 6-edge cut is spoke-only. This is a genuine *a priori* restriction.
2. *Generalise.* Extend the per-tire half to branched cut tires. Proper edge 3-coloring on such a structure is well-defined and probably has a non-empty S_3-closed projection by similar arguments, but the explicit count $2^n + 2(-1)^n$ no longer applies.

What this changes in the chain DP. The chain DP is still well-formed: enumerate proper 3-edge-colorings of each cut tire (now allowing branches), project to out spokes, restrict via children. The per-tire half just needs the generalized form.

Second issue: out-spoke projection loses $\mathbf{S}_3$ orbit. The per-tire half guarantees a full S_3 orbit on the *joint* in+out spoke projection $\pi(T)$. After restricting to OUT spokes only (which is what the parent uses), the projection $A(T)$ may contain fewer than 6 elements — e.g. all out spokes might be forced to a constant tuple by some structural symmetry, giving $|A(T)| = 3$ ($= \{(0,0,\dots),(1,1,\dots),(2,2,\dots)\}$). Initial empirical runs on the dodecahedron and HM #0 see exactly this happen at $\sim 20\%$ of cut tires. This is *not* a bug in the per-tire half; it is a genuine limitation of the OUT-only projection.

The correct chain DP formulation should track the joint (in + out) projection, not just OUT. This is the analogue of `tire_fiber_step2.tex`'s joint-support tracking.

Third issue: heuristic parent-finding. The current `cut_tire_tree.find_parent_face` uses a vertex-overlap heuristic (smallest face among overlapping candidates). Per the high-side proposition (`cut_tire_tree_structure.tex`), the geometrically-correct parent is unique, but the empirical script does not enforce this — it picks by smallest-face heuristic, which can mis-attribute children to wrong parents. For a rigorous empirical test, parent assignment should use the planar embedding's face-in-face containment, not vertex overlap.

Net status of the loose conjecture

component	status	note
Per-tire half (spoke-only $n \geq 3$)	proven	Prop 1.13
Per-tire half (branched)	open, needed for generality	
Tree structure (forest, high-side)	proven	`cut_tire_tree_structure.tex`
Chain DP S_3-equivariance	proven	this note, Lemma
Joint vs. OUT projection issue	flagged	need joint-support tracking
Chain DP non-emptiness preservation	open	Conj.
Bottom-line $\mathcal{R}_0 \cap \mathcal{R}_1 \neq \emptyset$	open, G-colorability gives it	

The chain half reduces to multiple structural claims:

- per-tire half for branched cut tires;
- joint-support DP (track $\pi(T)$, not just OUT projection);
- non-emptiness preservation (Conj. or Strong per-tire extendibility).

None are in hand yet. The S_3-equivariance and forest structure are. The full chain half is genuinely open and requires more work.