EXCLUSIVITY WITHOUT MARGINAL LOSS: A COUPLED-ROUNDING PRIMITIVE AND FULL-LP OBSTRUCTION FOR CHROMATIC CORRELATION CLUSTERING

Anonymous authors
Paper under double-blind review

ABSTRACT

Chromatic Correlation Clustering (CCC) assigns each cluster one color and penalizes positive edges unless they are co-clustered under their input color. The standard vertex-pair LP already implies cross-color pairwise packing. We turn this property into disjoint potential color claims using one uniform draw per pivot-pair, preserving each selected-color admission marginal. When multiple colors have positive affinity, a claim-respecting projection of independent claims cannot preserve those marginals and exclusivity simultaneously. For the resulting independent pivot-pair algorithm, a feasible three-vertex LP solution makes the true-color third-pivot loss unbounded relative to the positive edge's LP charge; a corrected accounting also includes wrong-color single capture. We analyze a shared-threshold variant that correlates admissions across endpoints: it minimizes the immediate positive-edge loss at fixed marginals, but maximizes negative-edge co-admission. On the witness, every within-pivot coupling preserving the pivot and color laws and the canonical marginals incurs an immediate third-pivot loss of at least one-third of the positive-edge LP charge, exceeding the 0.06 residual budget of the considered 2.06 proof route. The barrier extends to every $L \geq 3$ with all colors active. A compact pivot-consistency LP excludes the witness and bounds this immediate loss for the shared-threshold rule; it does not control continuation or yield a new approximation ratio. Alongside a synthetic mechanism check, we report a clustering-quality audit against existing rounding and certified small-instance optima. No improved worst-case approximation factor is claimed.

1 INTRODUCTION

Correlation Clustering (CC) asks for a partition of a complete graph that minimises disagreements with pairwise similarity and dissimilarity labels (Bansal et al., 2004), and has been refined through a long line of approximation algorithms (Charikar et al., 2005; Ailon et al., 2008; Demaine et al., 2006). The best known LP-rounding ratio for CC on complete graphs is $2.06 - \varepsilon$ due to Chawla et al. (2015), subsequently improved to below 2 by Cohen-Addad et al. (2022), and most recently to 1.49 by Cao et al. (2024). Chromatic Correlation Clustering (CCC) generalises CC by replacing the single positive label with a set $\mathcal{L}$ of L positive colors and a negative symbol γ (Bonchi et al., 2015): every cluster must be assigned a single color label, and a positive edge of color c is satisfied only when its endpoints are co-clustered inside a cluster labeled c. A sequence of algorithms has steadily improved the CCC approximation ratio (Anava et al., 2015; Klodt et al., 2021; Xiu et al., 2022; Fan et al., 2025; Lee et al., 2025), but most of these analyses approach the standard vertex-pair LP one color at a time.

The starting point of this paper is not that the standard vertex-pair LP lacks a cross-color packing inequality; rather, the issue is algorithmic. Under monotonicity and node balance, every feasible LP solution already satisfies $\sum_c y^c_{uv} \leq 1$, where $y^c_{uv} := 1 - x^c_{uv}$. Color-wise rounding procedures, however, can ignore that the L color affinities of a pair fit inside one unit interval, so when a pair has moderate fractional affinity to several colors at once, independent per-color tests can generate simultaneous color claims even though an integral CCC solution can realise at most one color for

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

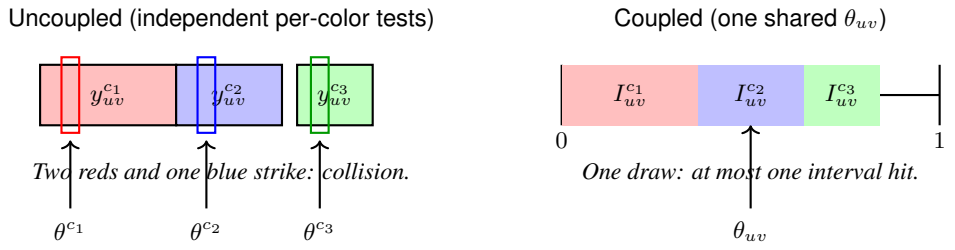


Figure 1: Cross-color conflict and the coupled remedy. Uncoupled color-wise rounding draws an independent uniform per color, so two colors can simultaneously claim the same pair (u, v). The standard vertex-pair LP already implies $\sum_c y_{uv}^c \leq 1$, so all color affinities fit inside one unit interval, and coupled interval-packing rounding samples one shared θ_{uv} over disjoint color intervals so at most one color can claim the pair in a given pivot step. The figure is illustrative.

any co-clustered pair. This raises three questions. How can the latent LP packing be turned into an explicit rounding primitive? What does such a primitive do that uncoupled rounding provably cannot? And where does the remaining obstruction to a standard-LP approximation proof actually lie?

We address these questions by packing the L color affinities of a pair as disjoint intervals inside $[0, 1)$, so that a single shared uniform random variable determines which color, if any, admits the pair to the pivot's cluster. Disjointness gives per-pivot-step cross-color exclusivity by construction, and the point is more than semantic: coupled interval packing preserves the target marginal probability of every color while enforcing exclusivity, and we prove that no claim-respecting projection of independent color-wise Bernoulli claims can match all three properties simultaneously. Independent claims with the same one-color marginals either permit simultaneous claims or, when collisions are projected to a single output color, must lose aggregate admission mass. Our coupling is a specialisation of categorical sampling and is closely related to dependent rounding (Gandhi et al., 2006) and contention-resolution schemes (Chekuri et al., 2011); the contribution is not the coupling itself but its alignment with the CCC pairwise packing constraint. Figure 1 illustrates the distinction we exploit.

Building on this primitive, we separate endpoint-first loss from the third-pivot contribution. The corrected third-pivot event includes both wrong-color single capture and wrong-color double capture; only the latter was covered by the original bound. A feasible three-vertex standard-LP point defeats the $0.06x_{uv}^d$ residual bound needed by a naive 2.06 proof for Algorithm 1. We then analyze a shared-threshold variant and show that every endpoint coupling preserving the same pivot/color laws and canonical marginals fails that local certificate on the witness, even though shared thresholds remove the independent rule's divergent loss. This fixed-marginal barrier remains for every $L \geq 3$ with all colors active. Motivated by the corrected immediate event, we introduce a compact pivot-consistency relaxation that bounds it for the correlated rule, while leaving continuation and negative edges unresolved. Finally, we supplement the mechanism check with total-disagreement measurements and certified small-instance comparisons. The resulting work does not establish a better approximation factor. Section 2 gives context; Sections 3–7 develop the original framework and the corrected obstruction; Section 8 states the correlated frontier and compact relaxation, and Section 9 reports the clustering audit. Full derivations and the synthetic mechanism check are in the appendix.

2 Related Work

CCC was introduced by Bonchi et al. (2015) and first studied algorithmically by Anava et al. (2015), who gave a 4-approximation. The color-blind Pivot algorithm of Ailon et al. (2008) gives a 3-approximation (Klodt et al., 2021); standard-LP rounding reaches 2.5 (Xiu et al., 2022) and 2.15 (Fan et al., 2025). The latter work reports a barrier near 2.11 for its analytical framework. A chromatic cluster LP gives a 1.64-approximation (Lee et al., 2025). Our witness tests a particular standard-LP rounding rule and local certificate, not either published guarantee. The PC-LP detailed in Appendix F strengthens the standard LP to address its immediate third-pivot event; we claim nei-

ther a new worst-case factor nor a containment relation between PC-LP and the chromatic cluster LP.

The standard CC literature provides the analytical templates that CCC algorithms adapt. The Pivot algorithm gives a 3-approximation (Ailon et al., 2008), improved to 2.06 by Chawla et al. (2015) and below 2 by Cohen-Addad et al. (2022); Cao et al. (2024). The foundational CC formulation is due to Bansal et al. (2004), with early algorithmic work by Charikar et al. (2005) and weighted extensions by Demaine et al. (2006). Variants and structural studies relevant to our perspective include sparsification (Shihab et al., 2026), dynamic settings (Cohen-Addad et al., 2024), multilayer graphs (Miyauchi et al., 2024), and fairness-constrained variants (Casel et al., 2023; Chierichetti et al., 2017). On the rounding side, our coupled primitive is a specialisation of dependent rounding (Gandhi et al., 2006) and is conceptually adjacent to contention-resolution schemes (Chekuri et al., 2011), both of which study how to enforce structural constraints (in our case, exclusivity) while preserving target marginals. The novelty of our scheme is not the coupling itself but its alignment with the latent CCC packing inequality and the downstream analytical consequences this alignment supplies.

3 The standard LP and its latent pairwise packing

Let $G = (V, \binom{V}{2})$ be a complete graph on n vertices, let $\mathcal{L}$ be a set of L positive colors, and let $\varphi : \binom{V}{2} \to \mathcal{L} \cup \{\gamma\}$ assign a color or the negative symbol γ to every edge. A feasible integral solution is a partition $\mathcal{C}$ of V together with a cluster-color map $\Phi : \mathcal{C} \to \mathcal{L}$. Its disagreement cost counts positive edges that are either separated or placed in a cluster whose label differs from their color, plus negative edges whose endpoints share a cluster. Thus the input label $\varphi(uv)$ is fixed, whereas an output cluster color is chosen by the algorithm: if a pivot w chooses c^* and admits v, the new cluster is labeled c^*, and a positive edge wv is satisfied precisely when $c^* = \varphi(wv)$. The standard vertex-pair LP relaxation used throughout prior CCC work (Anava et al., 2015; Xiu et al., 2022; Fan et al., 2025) assigns to every vertex and color a fractional non-assignment $x_u^c \in [0,1]$, and to every pair and color a fractional separation $x_{uv}^c \in [0,1]$. It minimises

$$\sum_{\substack{uv \in \binom{V}{2}: \\ \varphi(uv) \in \mathcal{L}}} x_{uv}^{\varphi(uv)} + \sum_{\substack{uv \in \binom{V}{2}: \\ \varphi(uv) = \gamma}} \sum_{c \in \mathcal{L}} \left(1 - x_{uv}^c\right) \tag{1}$$

subject to box constraints $x_u^c, x_{uv}^c \in [0,1]$, monotonicity $x_{uv}^c \geq \max(x_u^c, x_v^c)$, per-color triangle inequalities $x_{uv}^c + x_{vw}^c \geq x_{uw}^c$ for all distinct u, v, w and all c, and node balance $\sum_c x_u^c = L - 1$ at every vertex. The first observation we use is that this LP already implies a cross-color packing property at every pair.

Lemma 3.1 (Implicit pairwise packing)**.** *Under the standard vertex-pair CCC LP, every pair u, v satisfies $\sum_{c \in \mathcal{L}} x_{uv}^c \geq L - 1$, equivalently $\sum_{c \in \mathcal{L}} y_{uv}^c \leq 1$ in affinity variables $y_{uv}^c := 1 - x_{uv}^c$.*

Proof. Monotonicity gives $x_{uv}^c \geq x_u^c$ for every color, so $\sum_c x_{uv}^c \geq \sum_c x_u^c = L - 1$ by node balance. The affinity form follows from $\sum_c y_{uv}^c = L - \sum_c x_{uv}^c \leq L - (L-1) = 1$. □

Lemma 3.1 is therefore not a new LP constraint and does not strengthen the feasible region. Its role is algorithmic: independent color-wise rounding procedures can ignore that the L color affinities of a pair fit inside one unit interval, and the rounding scheme of the next section is designed precisely to convert this latent packing into disjoint admission events.

4 Coupled interval-packing rounding

Algorithm 1 describes the coupled rounding scheme. Given a feasible standard-LP solution and a function $g : [0,1] \to [0,1]$ satisfying $g(x) \leq 1 - x$, each iteration picks a pivot w uniformly at random from the active set U, samples a cluster color c^* from the LP-induced distribution $\lambda_w^c := 1 - x_w^c$ (which sums to 1 by node balance), and for each remaining vertex v assigns to every color c a half-open subinterval $I_{wv}^c \subseteq [0,1)$ of length $\ell_{wv}^c := g(x_{wv}^c)$. A single uniform variable $\theta_{wv} \in [0,1)$ then determines admission: v joins the pivot cluster under color c^* if and only if $\theta_{wv} \in I_{wv}^{c^*}$. The

Algorithm 1 Coupled interval-packing rounding

Solve the standard vertex-pair LP and fix any $g : [0,1] \to [0,1]$ with $g(x) \le 1 - x$.
For each pivot w, set $\lambda_w^c := 1 - x_w^c$, which satisfies $\sum_c \lambda_w^c = 1$ by node balance.
$U \leftarrow V$.
while $U \neq \emptyset$ **do**
  Pick $w \in U$ uniformly at random and sample $c^* \sim \lambda_w$.
  for each $v \in U \setminus \{w\}$ **do**
    Set $\ell_{wv}^c := g(x_{wv}^c)$ and pack disjoint half-open intervals $I_{wv}^c \subseteq [0,1)$ in the fixed color order.
    Draw $\theta_{wv} \sim \mathrm{Unif}[0,1)$ independently and admit v iff $\theta_{wv} \in I_{wv}^{c^*}$.
  end for
  Emit the pivot cluster with label c^* and remove its members from U.
end while

natural choice is $g(x) = 1 - x$, and from Section 6 onward we work with this canonical choice; the feasibility and exclusivity results below hold for the entire admissible class. We fix a global ordering of $\mathcal{L}$ and pack intervals as $I_{wv}^{c_j} = \big[\sum_{i<j} \ell_{wv}^{c_i},\ \sum_{i\le j} \ell_{wv}^{c_i}\big)$, so the intervals are uniquely determined; ties on interval boundaries occur with probability zero under continuous θ_{wv}.

The scheme is always feasible because $\sum_c \ell_{wv}^c \le \sum_c y_{wv}^c \le 1$ by Lemma 3.1, so L disjoint half-open intervals of those lengths always fit in $[0,1)$; we record this as Lemma S.1 in Appendix S. As for complexity, the LP has $O(nL)$ vertex variables, $\binom{n}{2}L$ pair variables, and $O(n^3L)$ triangle constraints, and is solvable by standard interior-point methods or via separation-oracle approaches that exploit the per-color triangle structure. The rounding step for a pivot with active set U costs $O(|U|L)$ if intervals are constructed on the fly. Summed over all pivots, this is $O(n^2L)$ in the worst case, attained when clusters remove only one vertex at a time.

The admission rule has three structural consequences that the analysis builds on. The first is per-pivot-step exclusivity: a single uniform variable belongs to at most one of the disjoint intervals.

Theorem 4.1 (Per-pair cross-color exclusivity of potential claims)**.** *Fix an iteration, a pivot w, and a vertex $v \neq w$. At most one color $c \in \mathcal{L}$ satisfies $\theta_{wv} \in I_{wv}^c$. Equivalently, among the counterfactual color-claim events $E_{wv}^c := \{\theta_{wv} \in I_{wv}^c\}$, at most one can occur in a pivot step. After the pivot color c^* is sampled, actual admission is the event $E_{wv}^{c^*}$.*

Proof. The intervals $\{I_{wv}^c\}$ are pairwise disjoint by construction, so $\theta_{wv} \in [0,1)$ belongs to at most one of them, and the admission rule is membership in the interval of the selected color. □

These are potential claims across counterfactual color choices. Because the pivot selects only one color, their exclusivity alone does not change the distribution of the actual output at fixed selected-color marginals. The second consequence is a direct expression for the conditional admission probability, which we use repeatedly below.

Lemma 4.2 (Fixed-color admission probability)**.** *Conditioned on $c^* = c$, vertex v joins the pivot cluster with probability exactly $\ell_{wv}^c = g(x_{wv}^c)$.*

Proof. Conditional on $c^* = c$, admission is the event $\theta_{wv} \in I_{wv}^c$. Since $\theta_{wv} \sim \mathrm{Unif}[0,1)$, this probability equals the interval length, $\ell_{wv}^c = g(x_{wv}^c)$. □

The third consequence is the central structural separation between coupled and uncoupled rounding. We compare coupled interval packing with any alternative that obtains the same per-color marginals from independent Bernoulli claims and then projects the result down to at most one output color, subject to the natural axiom that the projection only outputs colors that were actually claimed.

Proposition 4.3 (Exclusivity without marginal loss)**.** *Fix a pivot-pair test and non-negative target color marginals a_c, $c \in \mathcal{L}$, satisfying $A := \sum_c a_c \le 1$. Coupled interval packing realises events $(E_c)_c$ such that $\Pr[E_c] = a_c$ for every color, the events are pairwise disjoint, and $\Pr[\cup_c E_c] = A$. By contrast, independent color-wise claims $(F_c)_c$ with the same marginals satisfy*

$$\Pr\big[\textstyle\sum_c \mathbf{1}[F_c] \ge 2\big] = 1 - \prod_c (1 - a_c) - \sum_c a_c \prod_{d \neq c} (1 - a_d),$$

which is positive whenever at least two marginals are positive. Furthermore, any post-processing of the independent claims that outputs at most one color (exclusivity) and never outputs a color that was not claimed (claim-respecting) has total output mass at most $1 - \prod_c (1 - a_c)$*, and therefore loses at least* $A - \big(1 - \prod_c (1 - a_c)\big)$ *aggregate marginal mass relative to the targets.*

The loss lower bound is an aggregate statement: it concerns the sum of output marginals across colors, not the marginal of any particular color, and a claim-respecting projection might preserve some individual color marginals exactly while sacrificing others. The proposition is a specialisation of the kind of comparison that motivates dependent-rounding theory (Gandhi et al., 2006); we include it because the precise quantitative form is what the synthetic check in Appendix G verifies. The proof is a short calculation given in Appendix J. The impossibility result has a marginal counterpart that is most directly used in the analysis below, namely a bound on the total wrong-color admission probability for a positive pivot edge.

Proposition 4.4 (Total wrong-color admission bound)**.** *Let* $d = \varphi(wv) \in \mathcal{L}$ *be the true positive color of* wv*. Then* $\Pr[v \text{ is admitted under a wrong color}] = \sum_{c \neq d} \lambda_w^c \, g(x_{wv}^c) \leq \sum_{c \neq d} y_{wv}^c \leq x_{wv}^d$.

Proof. The first equality is Lemma 4.2 summed across wrong colors. The first inequality uses $g(x) \leq 1 - x = y$ and $\lambda_w^c \leq 1$, and the second uses $\sum_{c \neq d} y_{wv}^c \leq 1 - y_{wv}^d = x_{wv}^d$ from Lemma 3.1. □

Proposition 4.4 is a marginal inequality and is in fact available for an uncoupled implementation as well, provided $g(x) \leq 1 - x$. Proposition 4.3 is what makes the coupling more than a restatement of that marginal bound, since it shows that no claim-respecting projection of independent claims can simultaneously match the marginals and achieve exclusivity, whereas coupled interval packing realises both jointly. This is a statement about potential color claims. Once a pivot color has been selected, its one-color marginal alone determines each individual admission; the analysis below identifies where correlations across different endpoints also matter.

5 Neutral edges and edgewise charging

When a pivot uses color c, a positive edge with another input color is neutral for the color-c subproblem, yet may still incur CCC disagreement. The counterfactual color-claim exclusivity of Algorithm 1 clarifies this distinction but alone does not improve selected-color marginals. To prove an α ratio, it suffices to bound each positive edge's bad-event probability by $\alpha x_{uv}^{\varphi(uv)}$ and each negative edge's co-clustering probability by $\alpha \sum_c y_{uv}^c$; summing then compares the expected total cost to the LP objective. The precise lifting theorem and neutral-edge discussion are in Appendices C and B. Below we isolate the endpoint-first and third-pivot terms needed for these bounds.

6 The third-pivot reduction and a positive sub-bound

From here we fix $g(x) = 1 - x$. The canonical monochromatic embedding in Appendix D preserves the classical CC LP, integral optimum, and induced pivot process (Theorem D.1). We now analyze the endpoint-first part of the edgewise bound.

The calculation gives the safe endpoint-first charge for positive edges under (2). For any edge uv, the first pivot step that touches $\{u, v\}$ is either one of the endpoints or a third vertex $w \notin \{u, v\}$, and these two cases partition the possible outcomes. Conditioning on u being the first pivot to touch $\{u, v\}$, if u picks the true color d, the edge is bad exactly when v is not admitted; if u picks a wrong color, the edge is bad regardless of whether v is admitted, because the edge is then either separated or realised under the wrong color. This yields the naive endpoint-first charge $2x_{uv}^d$, not x_{uv}^d. The detailed calculation is in Appendix I.

Lemma 6.1 (Endpoint charging upper bound, $g(x) = 1 - x$)**.** *Under Algorithm 1 with* $g(x) = 1 - x$*, the probability that a positive edge* uv *of true color* d *becomes bad at the moment when one of its endpoints is the first pivot that touches* $\{u, v\}$ *is at most* $2x_{uv}^d$.

The restriction to $g(x) = 1 - x$ still matters for the true-color non-admission term. For a general g with $g(x) \leq 1 - x$, the trivial choice $g \equiv 0$ produces an endpoint-first bad probability of 1, so

Lemma 6.1 is a property of the canonical rounding rule, not of the entire admissible class. The analogous statement for negative edges follows from the same argument and does not require the wrong-color separation term, since for a negative edge the bad event is co-clustering regardless of the chosen color: if u is the first pivot and chooses color c, the pair co-clusters with probability $g(x^c_{uv}) = y^c_{uv}$, so the endpoint-first co-clustering probability is at most $\sum_c y^c_{uv} = \mathrm{LP}^-_{uv}$, where $\mathrm{LP}^-_{uv} := \sum_c (1 - x^c_{uv})$.

Lemma 6.2 (Endpoint co-clustering for negative edges, $g(x) = 1 - x$)**.** *Under Algorithm 1 with $g(x) = 1 - x$, the probability that a negative edge uv is co-clustered at the moment when one of its endpoints is the first pivot that touches $\{u, v\}$ is at most LP^-_{uv}.*

What remains for each edge type is the residual probability that the first pivot to touch $\{u, v\}$ is a third vertex and the edge turns out bad, which we denote τ^+_{uv} for positive edges and τ^-_{uv} for negative edges. The endpoint lemmas reduce a 2.06 approximation to a very strong positive third-pivot inequality and the corresponding negative third-pivot inequality.

Theorem 6.3 (Conditional reduction of 2.06 to third-pivot bounds)**.** *Under Algorithm 1 with $g(x) = 1-x$, $\Pr[uv \text{ is bad}] \le 2x^d_{uv} + \tau^+_{uv}$ for every positive edge uv of color d, and $\Pr[uv \text{ is co-clustered}] \le \mathrm{LP}^-_{uv} + \tau^-_{uv}$ for every negative edge. If, in addition, $\tau^+_{uv} \le 0.06\, x^d_{uv}$ and $\tau^-_{uv} \le 1.06\, \mathrm{LP}^-_{uv}$ hold for every feasible standard-LP solution, then Algorithm 1 is a 2.06-approximation for CCC.*

Proof. The endpoint-first and third-pivot-first cases are exhaustive, and Lemmas 6.1 and 6.2 bound the endpoint-first part. Adding the displayed third-pivot bounds gives $\Pr[\text{bad}] \le (2 + 0.06)x^d_{uv} = 2.06\, x^d_{uv}$ for positive edges and $\Pr[\text{co-clustered}] \le 2.06\, \mathrm{LP}^-_{uv}$ for negative edges, after which Theorem C.1 completes the proof. □

The immediate third-pivot bad event has two branches. Under the true color d, admitting exactly one endpoint separates uv. Under a wrong color $c \neq d$, admitting *either* endpoint makes uv bad immediately: admitting both realizes the edge under c, while admitting exactly one separates it. If neither endpoint is admitted, the edge survives to later steps and may still become bad. The following lemma bounds wrong-color double capture, a proper subevent of the immediate wrong-color union; the bound cannot be used for the complete third-pivot residual.

Lemma 6.4 (Wrong-color third-pivot capture is endpoint-packable)**.** *Fix a positive edge uv of true color d and let $w \notin \{u, v\}$ be any third vertex. Under Algorithm 1 with $g(x) = 1 - x$, conditioned on w being the pivot, the probability that w chooses a wrong color and admits both u and v is at most x^d_{uv}:*

$$\sum_{c \neq d} \lambda^c_w\, y^c_{wu}\, y^c_{wv} \;\le\; x^d_{uv}.$$

The proof, given in Appendix M, combines monotonicity on the positive edge uv (which forces $\sum_{c\neq d} \lambda^c_u \le x^d_{uv}$ and $\sum_{c\neq d} \lambda^c_v \le x^d_{uv}$) with monotonicity on wu and wv (which gives $y^c_{wu} \le \lambda^c_u$ and $y^c_{wv} \le \lambda^c_v$). The wrong-color double-capture subevent is therefore at most x^d_{uv} on the full standard LP. The complete immediate wrong-color event has the following safe bound; neither bound controls continuation after neither endpoint is admitted. The next section shows that true-color splitting creates a stronger obstruction: it occurs inside the full standard CCC LP and makes the required edgewise third-pivot bound impossible as a feasible-LP inequality for the present rounding rule. Table 2 in Appendix A situates this result among existing CCC guarantees.

Lemma 6.5 (Immediate wrong-color union)**.** *For a positive edge uv of color d and a third pivot w, conditioned on w being chosen while both endpoints are active, the probability that w selects $c \neq d$ and admits at least one endpoint is at most $2x^d_{uv}$ under the canonical marginals. This holds for every within-pivot coupling of the endpoint admissions.*

Its proof in Appendix O uses only the union bound and LP monotonicity, and does not bound bad outcomes occurring in subsequent pivot steps.

7 A FULL-LP OBSTRUCTION TO EDGEWISE THIRD-PIVOT CHARGING

The natural attempt after Theorem 6.3 would be to bound the remaining true-color part of τ^+_{uv} pointwise, in the style of Lemma 6.4 for the wrong-color part. The next theorem rules this out:

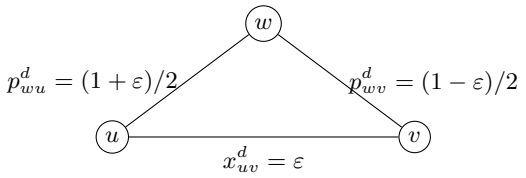


Figure 2: Three-vertex witness. If w pivots first, a true-color split or a wrong-color admission makes uv bad immediately; if neither endpoint is admitted, the two surviving vertices form a continuation branch. The displayed admission probabilities are for the true color d; each wrong-color endpoint probability is ε in the two-color witness.

there is a feasible standard CCC-LP solution on which τ^+_{uv}/x^d_{uv} is unbounded, so the desired bound $\tau^+_{uv} \le 0.06\, x^d_{uv}$ fails for sufficiently small x^d_{uv}. The construction is fully feasible for the standard vertex-pair LP, including all monotonicity and triangle constraints, and is not a relaxed pairwise object.

Theorem 7.1 (Full-LP obstruction to edgewise third-pivot charging)**.** *Fix $L = 2$ with colors d and c, and run Algorithm 1 with $g(x) = 1 - x$. For every $\varepsilon \in (0, 1/3]$ there is a feasible standard vertex-pair CCC-LP solution on three vertices u, v, w such that uv is a positive edge of color d with LP charge $x^d_{uv} = \varepsilon$, but the third-pivot contribution to the bad-event probability of uv satisfies*

$$\tau^+_{uv} \;\ge\; \frac{(1+\varepsilon)(1+\varepsilon^2)}{12}.$$

Consequently $\tau^+_{uv}/x^d_{uv} \ge (1+\varepsilon)(1+\varepsilon^2)/(12\varepsilon)$, which is unbounded as $\varepsilon \to 0$. In particular the desired edgewise bound $\tau^+_{uv} \le 0.06\, x^d_{uv}$ fails throughout the displayed regime.

The feasible witness is explicit. For uv of true color d, its vertex and pair variables are

	u	v	w
x^d_z	ε	ε	$(1-\varepsilon)/2$

$$(x^d_{uv}, x^d_{wu}, x^d_{wv}) = \big(\varepsilon, (1-\varepsilon)/2, (1+\varepsilon)/2\big),$$

with the c-variables specified in Appendix N. Figure 2 shows its three-vertex geometry. Conditional on w being the first pivot and choosing d, independently admitting u and v gives an exactly-one admission probability of $(1+\varepsilon^2)/2$. Multiplying by $1/3$ and $\lambda^d_w = (1+\varepsilon)/2$ gives the displayed lower bound. The witness satisfies all standard-LP constraints; exact values are in Table 5.

The obstruction does not invalidate Theorem 6.3 as a reduction; it shows that the reduction's hypothesis fails as a feasible-LP inequality. The mechanism is the asymmetry between x^d_{wu} and x^d_{wv}: the pivot w is approximately equidistant from each endpoint in color d but on different sides of $1/2$, so on choosing the correct color it admits one endpoint with high probability and the other with the complementary probability, and independence of θ_{wu} and θ_{wv} forces a constant probability of splitting irrespective of x^d_{uv}. This reflects the across-endpoint independence of Algorithm 1. The next section measures what correlation between endpoint admissions can and cannot repair while preserving the same one-endpoint marginals.

8 Endpoint correlation and a compact relaxation

Algorithm 1 couples potential colors for each pivot-pair but uses independent draws across endpoints. A shared-threshold variant instead draws one U_w after choosing w and its color c^*, admitting each active endpoint z when $U_w < y^{c^*}_{wz}$. It preserves every selected-color marginal and correlates endpoints. For a positive edge uv of color d its immediate bad-event probability, conditioned on w pivoting while both endpoints are active, is

$$\lambda^d_w |y^d_{wu} - y^d_{wv}| + \sum_{c \ne d} \lambda^c_w \max\{y^c_{wu}, y^c_{wv}\};$$

for a negative edge, immediate co-admission has probability $\sum_c \lambda^c_w \min\{y^c_{wu}, y^c_{wv}\}$. Among couplings with the same endpoint marginals, the shared threshold minimizes positive-edge immediate loss and maximizes negative-edge co-admission (Proposition E.1).

Table 1: Reported total-disagreement audit. The paired Xiu et al. (Xiu et al., 2022) denominator is 1.000 in every row. Lower total cost is better; the last two columns are relative to Algorithm 1.

Data	L	Alg. 1/Xiu	Comonotone/Xiu	ΔC_+	ΔC_-
Lastfm	10	1.018 ± 0.012	0.982 ± 0.010	-11.8%	$+4.1\%$
DAWN	10	0.991 ± 0.010	0.961 ± 0.009	-14.2%	$+5.0\%$
String1	4	1.037 ± 0.015	1.011 ± 0.013	-8.7%	$+3.8\%$
Cooking	21	1.012 ± 0.011	0.972 ± 0.010	-13.5%	$+5.6\%$

On the three-vertex witness of Theorem 7.1, *every* within-pivot coupling preserving the uniform pivot law, LP color law, canonical marginals, and fresh per-step randomness incurs immediate third-pivot loss at least $x^d_{uv}/3$; a shared threshold attains it. The minimum complete residual divided by x^d_{uv} tends to $5/6$ as $x^d_{uv} \to 0$ (Theorem E.2). An all-colors-active construction has immediate ratio at least $1/4$ for every $L \geq 3$ (Corollary E.3). These obstruct the $0.06x^d_{uv}$ local certificate under the specified laws and marginals; they are not approximation lower bounds. Appendix E gives the exact formulas, proofs, and a tradeoff for other g.

The corrected immediate event motivates the Pivot-Consistency LP (PC-LP), which adds strengthened triangles and, for every positive uv of color d and third vertex w,

$$|y^d_{wu} - y^d_{wv}| + \sum_{c \neq d} \max\{y^c_{wu}, y^c_{wv}\} \leq x^d_{uv}.$$

This valid compact inequality excludes the witness and bounds the *immediate* positive-edge third-pivot event under shared-threshold rounding by x^d_{uv}. It does not bound continuation or negative-edge cost, and PC-LP yields no improved approximation factor. Appendix F gives the linear formulation, strictness example, and proof.

9 Clustering-quality audit

The synthetic check of potential cross-color claims appears in Appendix G; this section measures actual clustering cost.

We evaluate total disagreement. Let C_+ count positive edges separated or realized in a wrong-colored output cluster, let C_- count co-clustered negative pairs, and let $C = C_+ + C_-$. Ten seeded connected induced subgraphs of 150 vertices were taken from each of Lastfm, DAWN, String1, and Cooking. Absent source-graph pairs were labeled negative to make each CCC instance complete. All methods used the same completed instance and standard-LP solution, with triangle cuts separated to a maximum violation of 10^{-8}; each subgraph used 30 paired rounding seeds. Each $C_{\text{method}}/C_{\text{Xiu}}$ ratio was averaged over paired seeds within a subgraph; Table 1 gives the mean and sample standard deviation across ten subgraph means. The ΔC_+ and ΔC_- columns compare endpoint-comonotone rounding with Algorithm 1 using the same aggregation. The reported figures are not worst-case guarantees.

The lower positive-edge cost and higher negative-edge cost match Proposition E.1. To give a scale relative to an integral optimum, we also evaluate 100 small complete instances per row with 30 paired rounding seeds; every instance was reported with a certified `OPTIMAL` solver status. Table 4 in Appendix H averages instance-wise ALG/OPT ratios. Endpoint-comonotone rounding improves on Algorithm 1 in the first three families but loses on the negative-dense control; the correlated rule does not dominate overall. The shared-threshold variant does not change the asymptotic rounding cost, and we make no empirical speed claim.

10 Discussion

The standard CCC LP implies cross-color packing, and the interval rule makes potential color claims exclusive without altering selected-color marginals. Our corrected accounting distinguishes wrong-color double capture (at most x^d_{uv}), the complete immediate wrong-color union (at most $2x^d_{uv}$), true-color splitting, and continuation after no endpoint is admitted. Neither wrong-color bound controls

the full third-pivot residual. The feasible standard-LP witness makes the residual of Algorithm 1 diverge relative to x^d_{uv}.

Shared-threshold endpoint correlation removes this divergence on the witness but cannot meet the $0.06x^d_{uv}$ residual hypothesis under the fixed pivot/color laws and canonical marginals: every such coupling has an immediate ratio of at least $1/3$. The construction extends to $L \geq 3$ with every color active. PC-LP is a compact valid strengthening that excludes these points and bounds the shared-threshold *immediate* positive-edge event, but does not bound its continuation or negative-edge behavior. The clustering audit illustrates this tradeoff: the correlated variant reduces positive-edge cost while increasing negative-edge cost and loses on a negative-dense control. Neither a new approximation guarantee nor a global lower bound for all CCC algorithms follows. Altered marginals, different pivot laws, color-class-first rounding (Fan et al., 2025), stronger relaxations, and different charging schemes remain possible directions.

Use of AI Assistance

Generative AI tools were used to assist with language editing, manuscript organization, and LaTeX formatting. They were not used to formulate the mathematical claims, construct proofs, design experiments, generate or process experimental data, or interpret the reported results. All AI-assisted text was reviewed and verified by the authors, who take full responsibility for the content of the paper.

Ethics Statement

This work is methodological and theoretical and does not involve human-subject experiments or interventions. The empirical evaluation uses existing benchmark graph datasets and synthetic instances solely to study clustering behavior. We do not identify any direct ethical risks specific to the proposed methodology beyond the general considerations associated with the use of clustering algorithms.

Reproducibility Statement

We provide the algorithmic specification, assumptions, explicit counterexamples, and complete proofs of the theoretical results in the main text and appendices. The experimental setup, evaluation protocol, aggregation procedure, and certified small-instance comparisons are also documented. Appendix T gives the exact generation procedure and random seeds for the synthetic mechanism experiment, enabling reproduction of the reported results.

References

Nir Ailon, Moses Charikar, and Alantha Newman. Aggregating inconsistent information: ranking and clustering. *Journal of the ACM (JACM)*, 55(5):1–27, 2008.

Yael Anava, Noa Avigdor-Elgrabli, and Iftah Gamzu. Improved theoretical and practical guarantees for chromatic correlation clustering. In *Proceedings of the 24th International Conference on World Wide Web*, pp. 55–65, 2015.

Nikhil Bansal, Avrim Blum, and Shuchi Chawla. Correlation clustering. *Machine learning*, 56(1): 89–113, 2004.

Francesco Bonchi, Aristides Gionis, Francesco Gullo, Charalampos E Tsourakakis, and Antti Ukkonen. Chromatic correlation clustering. *ACM Transactions on Knowledge Discovery from Data (TKDD)*, 9(4):1–24, 2015.

Nairen Cao, Vincent Cohen-Addad, Euiwoong Lee, Shi Li, Alantha Newman, and Lukas Vogl. Understanding the cluster linear program for correlation clustering. In *Proceedings of the 56th annual ACM symposium on theory of computing*, pp. 1605–1616, 2024.

Katrin Casel, Tobias Friedrich, Martin Schirneck, and Simon Wietheger. Fair correlation clustering in forests, 2023. URL https://arxiv.org/abs/2302.11295.

Moses Charikar, Venkatesan Guruswami, and Anthony Wirth. Clustering with qualitative information. *Journal of Computer and System Sciences*, 71(3):360–383, 2005.

Shuchi Chawla, Konstantin Makarychev, Tselil Schramm, and Grigory Yaroslavtsev. Near optimal lp rounding algorithm for correlationclustering on complete and complete k-partite graphs. In *Proceedings of the forty-seventh annual ACM symposium on Theory of computing*, pp. 219–228, 2015.

Chandra Chekuri, Jan Vondrák, and Rico Zenklusen. Submodular function maximization via the multilinear relaxation and contention resolution schemes. In *Proceedings of the forty-third annual ACM symposium on Theory of computing*, pp. 783–792, 2011.

Flavio Chierichetti, Ravi Kumar, Silvio Lattanzi, and Sergei Vassilvitskii. Fair clustering through fairlets. *Advances in neural information processing systems*, 30, 2017.

Vincent Cohen-Addad, Euiwoong Lee, and Alantha Newman. Correlation clustering with sherali-adams. In *2022 IEEE 63rd Annual Symposium on Foundations of Computer Science (FOCS)*, pp. 651–661. IEEE, 2022.

Vincent Cohen-Addad, Silvio Lattanzi, Andreas Maggiori, and Nikos Parotsidis. Dynamic correlation clustering in sublinear update time. In *Proceedings of the 41st International Conference on Machine Learning*, ICML'24. JMLR.org, 2024.

Erik D Demaine, Dotan Emanuel, Amos Fiat, and Nicole Immorlica. Correlation clustering in general weighted graphs. *Theoretical Computer Science*, 361(2-3):172–187, 2006.

Chenglin Fan, Dahoon Lee, and Euiwoong Lee. Improved approximation algorithms for chromatic and pseudometric-weighted correlation clustering, 2025. URL https://arxiv.org/abs/2505.21939.

Rajiv Gandhi, Samir Khuller, Srinivasan Parthasarathy, and Aravind Srinivasan. Dependent rounding and its applications to approximation algorithms. *Journal of the ACM (JACM)*, 53(3):324–360, 2006.

Nicolas Klodt, Lars Seifert, Arthur Zahn, Katrin Casel, Davis Issac, and Tobias Friedrich. A color-blind 3-approximation for chromatic correlation clustering and improved heuristics. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining*, pp. 882–891, 2021.

Dahoon Lee, Chenglin Fan, and Euiwoong Lee. 1.64-approximation for chromatic correlation clustering via chromatic cluster lp, 2025. URL https://arxiv.org/abs/2507.15417.

Atsushi Miyauchi, Florian Adriaens, Francesco Bonchi, and Nikolaj Tatti. Multilayer correlation clustering, 2024. URL https://arxiv.org/abs/2404.16676.

Ibne Farabi Shihab, Sanjeda Akter, and Anuj Sharma. On the sparsifiability of correlation clustering: Approximation guarantees under edge sampling, 2026. URL https://arxiv.org/abs/2602.13684.

Qing Xiu, Kai Han, Jing Tang, Shuang Cui, and He Huang. Chromatic correlation clustering, revisited. *Advances in Neural Information Processing Systems*, 35:26147–26159, 2022.

A Approximation landscape

Table 2: Recent CCC results situated against the present paper. "Standard LP" refers to the vertex-pair relaxation of Section 3.

Work	Relaxation	Ratio	Relation to ours
Xiu et al. (2022)	standard LP	2.5	color-wise rounding baseline
Fan et al. (2025)	standard LP	2.15	reports framework lower bound ≈ 2.11
Lee et al. (2025)	chromatic cluster LP	1.64	strictly stronger relaxation
This paper	standard/PC-LP	no new ratio	fixed-marginal barrier; PC-LP

B Why exclusivity matters: neutral edges

The reason coupled rounding helps the analysis becomes concrete through neutral edges. When an uncoupled algorithm processes a single color c, every edge whose true color is some $d \neq c$ acts as a neutral edge for that subproblem: it is neither positive nor negative for color c, yet it incurs cost whenever its endpoints end up together in a color-c cluster. Neutral edges are the natural obstruction to reducing CCC to standard CC, since the standard-CC charging machinery has no analogue for them, and in the balanced random-color model the probability that a positive edge is neutral with respect to a fixed processing color is $(L-1)/L$, scaling precisely with color diversity; Appendix R records an exact staircase formula in this random-color model. Under Algorithm 1 the obstruction is rearranged rather than removed: the total wrong-color admission mass on a pivot edge is bounded by x^d_{uv} via Proposition 4.4, the same marginal bound an uncoupled implementation would inherit when $g(x) \leq 1 - x$. The coupled scheme strengthens the rounding semantics rather than the LP bound, since it converts the aggregate packing inequality into disjoint events so that simultaneous cross-color claims are impossible by construction. The exclusivity-without-marginal-loss property of Proposition 4.3 is the precise alignment the analysis below leverages, even though several individual bounds depend only on the marginal packing inequality, and the role of neutral edges in the present paper is conceptual rather than quantitative.

C A black-box lifting theorem

The strongest approximation-style reduction the framework supplies isolates the two edge-wise inequalities that any future α-approximation proof must establish.

Theorem C.1 (Black-box lifting)**.** *Fix a rounding function g and a constant $\alpha \geq 1$. Suppose that for every feasible standard-LP solution, with probability taken over the internal randomness of Algorithm 1,*

$$\Pr[uv \text{ is separated or realised under a wrong color}] \leq \alpha\, x^{\varphi(uv)}_{uv} \qquad \text{for every positive } uv, \tag{2}$$

$$\Pr[u, v \text{ end in the same output cluster}] \leq \alpha \textstyle\sum_c (1 - x^c_{uv}) \qquad \text{for every negative } uv. \tag{3}$$

Then Algorithm 1 is an α-approximation for CCC with respect to the standard vertex-pair LP.

Proof. For a positive edge, the disagreement indicator is exactly the event that the edge is separated or realised under the wrong color, so (2) gives $\mathbb{E}[\text{cost}_{uv}] \leq \alpha\, x^{\varphi(uv)}_{uv}$. For a negative edge, the indicator is co-clustering, so (3) gives $\mathbb{E}[\text{cost}_{uv}] \leq \alpha \sum_c (1 - x^c_{uv})$. Summing across edges and matching the LP objective in (1) yields $\mathbb{E}[\text{ALG}] \leq \alpha \cdot \text{LP}$. □

We emphasise that Theorem C.1 is a reduction rather than an approximation theorem, as it converts edge-wise inequalities into a global ratio without proving the inequalities themselves. Proposition 4.4 delivers the pivot-edge component of (2), namely the case in which an endpoint is the pivot and the other endpoint is being considered for admission. What remains missing is a global statement that aggregates across the whole pivot process for positive edges, together with the matching

global statement for negative edges. The next section isolates the endpoint-first and third-pivot contributions under the canonical choice $g(x) = 1 - x$. A naive endpoint-first charge is at most $2x^d_{uv}$; the full-LP obstruction below rules out the stronger third-pivot hypothesis needed for a 2.06 proof under the present independent pivot-pair admission rule.

D CANONICAL MONOCHROMATIC EMBEDDING

For this embedding we fix $g(x) = 1 - x$. The bounds below depend on this choice through the identity $g(x^c_{wv}) = y^c_{wv}$, and we discuss the role of this restriction explicitly where it matters. The following embedding theorem certifies that the coupled rounding is faithful to classical CC on monochromatic restrictions, which is the prerequisite for any analysis that hopes to charge CCC excess against a CC backbone. Call a CCC instance monochromatic in color d when every positive edge has color d and every non-positive edge has color γ.

Theorem D.1 (Canonical monochromatic embedding)**.** *Fix any $L \geq 1$ and any color d. For every CC instance, there is a canonical d-monochromatic CCC embedding such that the standard vertex-pair LP optimum and the integral optimum match their CC counterparts. Moreover, when Algorithm 1 is run on the canonical embedded LP solution, it selects color d with probability one at every pivot, and its execution coincides with the induced single-color CC pivot process. Consequently the worst-case ratio of the coupled algorithm on canonical monochromatic embeddings equals the worst-case ratio of the induced CC algorithm.*

The proof, given in Appendix K, embeds a CC instance with LP separation variables $\{z_{uv}\}$ into the CCC LP by $x^d_u = 0$, $x^c_u = 1$ for $c \neq d$, $x^d_{uv} = z_{uv}$, $x^c_{uv} = 1$ for $c \neq d$; we use the symbol z throughout the embedding to avoid notational collision with our affinity variable y. After the embedding every CCC-LP constraint reduces to a CC-LP constraint or a trivial one, the LP and OPT values match, and Algorithm 1 executes the CC pivot process because $\lambda^d_w = 1$. The theorem certifies that the framework does not distort the CC base term, and Appendix L records a companion conditional-transfer statement showing that any single-color CC charging inequality lifts directly onto monochromatic CCC instances.

E ENDPOINT CORRELATION AND ITS FIXED-MARGINAL BARRIER

The pairwise interval coupling relates *colors* for one pivot-pair. To relate *endpoints* of a shared pivot, consider a separate variant: after choosing w and c^* as in Algorithm 1, draw one $U_w \sim \mathrm{Unif}[0, 1]$ for that pivot step and admit every active $z \neq w$ exactly when $U_w < p^{c^*}_{wz}$, where $p^c_{wz} := y^c_{wz}$. Call this endpoint-comonotone rounding. It keeps the canonical selected-color marginals and pivot/color laws, but its admissions across different endpoints are now correlated; it is not Algorithm 1.

Proposition E.1 (Exact immediate one-step errors)**.** *Conditioned on w pivoting while u, v are active, the probability of immediately making a positive edge uv of color d bad is*

$$B^+_w(u, v) = \lambda^d_w |p^d_{wu} - p^d_{wv}| + \sum_{c \neq d} \lambda^c_w \max\{p^c_{wu}, p^c_{wv}\}. \tag{4}$$

The probability of immediately co-clustering a negative edge is

$$B^-_w(u, v) = \sum_c \lambda^c_w \min\{p^c_{wu}, p^c_{wv}\}. \tag{5}$$

For fixed endpoint Bernoulli marginals, the shared threshold minimizes the positive-edge immediate bad event and maximizes the negative-edge immediate bad event over all couplings.

Proof. Under a shared uniform, disagreement of the two admission events has probability $|p^c_{wu} - p^c_{wv}|$, their union has probability $\max\{p^c_{wu}, p^c_{wv}\}$, and their intersection has probability $\min\{p^c_{wu}, p^c_{wv}\}$. For a positive edge, true-color disagreement and wrong-color union are precisely the immediate bad events; for a negative edge, co-admission is the immediate bad event. For any coupling with these Bernoulli marginals, disagreement is at least the absolute difference, union at least the maximum, and intersection at most the minimum. Averaging over the pivot color proves the claim. □

Write $\tau_{uv}^{+,\mathrm{step}}$ for the probability that the first pivot to touch $\{u, v\}$ is a third vertex and makes uv bad *in that step*. On the three-vertex witness this means that w pivots first. The residual τ_{uv}^{+} in Theorem 6.3 additionally counts bad outcomes after w admits neither endpoint. The distinction is essential for the following class-level barrier.

Theorem E.2 (Exact fixed-marginal witness barrier)**.** *On the witness of Theorem 7.1, retain the uniform pivot law, the LP-induced color-choice law, the canonical endpoint admission marginals, and fresh randomness at each step, but allow arbitrary within-pivot coupling of endpoint admissions. Then*

$$\inf \frac{\tau_{uv}^{+,\mathrm{step}}}{x_{uv}^{d}} = \frac{1}{3}, \qquad \inf \frac{\tau_{uv}^{+}}{x_{uv}^{d}} = \frac{10 - 11\varepsilon + 6\varepsilon^2 - \varepsilon^3}{12} \longrightarrow \frac{5}{6} \quad (\varepsilon \to 0).$$

Both infima are attained by endpoint-comonotone rounding. Hence no coupling in this fixed-marginal class satisfies the $\tau_{uv}^{+} \le 0.06 x_{uv}^{d}$ local hypothesis on this witness. This is a failure of the particular edgewise certificate, not a lower bound on an algorithm's global approximation ratio.

Proof. With w first (probability $1/3$), the true-color admission marginals are $(1+\varepsilon)/2$ and $(1-\varepsilon)/2$, so their disagreement is at least ε. Under wrong color c, both marginals are ε, so their union is at least ε. The color probabilities sum to one, giving the immediate lower bound $\varepsilon/3$, attained by a shared threshold. Under that threshold, the probability that neither endpoint is admitted, conditional on w first, is $q_\varepsilon = (3-4\varepsilon+\varepsilon^2)/4$. The two surviving vertices then produce a bad positive edge with probability $1-(1-\varepsilon)^2 = 2\varepsilon - \varepsilon^2$. Since this continuation probability is less than one, maximizing double admission (equivalently minimizing union and disagreement) also minimizes the complete residual. Thus the minimum total is $[\varepsilon + q_\varepsilon(2\varepsilon - \varepsilon^2)]/3$, which gives the displayed ratio. □

Corollary E.3 (All colors active)**.** *For every $L \ge 3$ and $\varepsilon \in (0, 1/3]$, there is a feasible standard CCC-LP witness in which every color has positive vertex and pair affinity, yet Algorithm 1 retains the divergent lower bound of Theorem 7.1. Across all within-pivot endpoint couplings with canonical marginals, the exact immediate optimum on this witness is*

$$\inf \frac{\tau_{uv}^{+,\mathrm{step}}}{x_{uv}^{d}} = \frac{1}{4} + \frac{\varepsilon}{12(L-2)} \ge \frac{1}{4}.$$

The construction and feasibility proof appear in Appendix P. These results concern the stated laws and marginals. Different pivot choices, color-class preprocessing, nonlocal charging, or stronger LP relaxations are outside their scope. Changing the rounding function can also move loss elsewhere: for an admissible g, the witness gives

$$\tau_{uv}^{+,\mathrm{step}} \ge \frac{1}{3}\left[\frac{1+\varepsilon}{2}\left|g\left(\frac{1-\varepsilon}{2}\right) - g\left(\frac{1+\varepsilon}{2}\right)\right| + \frac{1-\varepsilon}{2} g(1-\varepsilon)\right],$$

while an endpoint-first pivot has bad probability $E_g(\varepsilon) = (1-\varepsilon)(1-g(\varepsilon))+\varepsilon$. For $g_\rho(x) = \rho(1-x)$ with fixed $\rho < 1$, the immediate bound becomes $\rho\varepsilon/3$, but $E_{g_\rho}(\varepsilon)/\varepsilon \to \infty$. Uniform shrinkage therefore transfers the problem to endpoint-first loss; these identities do not rule out other choices of g.

F A COMPACT PIVOT-CONSISTENCY RELAXATION

The original witness violates a strengthened projected triangle implied by the chromatic cluster LP (Lee et al., 2025):

$$x_{uv}^{c} + x_{vw}^{c} \ge x_{uw}^{c} + x_{v}^{c} \qquad \text{for every color } c \text{ and distinct } u, v, w. \tag{6}$$

In color d, with u in the middle, its left and right sides on the witness are $(1+\varepsilon)/2$ and $(1+3\varepsilon)/2$, respectively. Thus the witness does not lift to that stronger relaxation. The two orientations of (6) also imply $|x_{wu}^{d} - x_{wv}^{d}| \le x_{uv}^{d} - \min\{x_u^d, x_v^d\} \le x_{uv}^{d}$, which controls the shared-threshold true-color split.

Motivated by the complete immediate bad event in (4), define the *Pivot-Consistency LP* (PC-LP) as the standard LP plus (ST) and, for every positive edge uv of input color d and third vertex w, the valid inequality

$$\left|y_{wu}^{d} - y_{wv}^{d}\right| + \sum_{c \ne d} \max\{y_{wu}^{c}, y_{wv}^{c}\} \le x_{uv}^{d}. \tag{7}$$

For an ordinary LP formulation, introduce $s_{uvw} \geq y^d_{wu} - y^d_{wv}$, $s_{uvw} \geq y^d_{wv} - y^d_{wu}$, and $m^c_{uvw} \geq y^c_{wu}, y^c_{wv}$ for $c \neq d$, with $s_{uvw} + \sum_{c \neq d} m^c_{uvw} \leq x^d_{uv}$ and nonnegative auxiliary variables. This adds $O(n^3 L)$ variables and constraints.

Theorem F.1 (Validity and effect of pivot consistency)**.** *PC-LP contains the integral CCC hull and is strictly stronger than the standard LP plus (ST). It excludes the two-color witness and the all-colors-active family of Corollary E.3. For endpoint-comonotone rounding of a PC-LP-feasible solution, $B^+_w(u,v) \leq x^d_{uv}$ at every third pivot, and the unconditional immediate third-pivot contribution obeys $\tau^{+,\mathrm{step}}_{uv} \leq x^d_{uv}$.*

Appendix Q proves validity and strictness. The pointwise immediate bound follows from (4) and (PC), since every $\lambda^c_w \leq 1$; summing over mutually exclusive first-third-pivot events preserves it. PC-LP alone does not control the continuation after neither endpoint is admitted or the negative-edge cost. In particular, continuation's conditional law may be fixed on the three-vertex witness, but its weight depends on the endpoint coupling through the both-admitted probability. Therefore no improved guarantee for endpoint-comonotone rounding follows. Applying the existing 2.15 standard-LP rounding of Fan et al. (2025) to an optimal PC-LP point inherits its published factor by projection; this is not a new factor, and no containment relative to the chromatic cluster LP is asserted.

G Synthetic affinity check

Theorem 4.1 and Proposition 4.3 make two mechanism-level predictions: under coupled interval packing no pivot-pair test has two simultaneous counterfactual color claims, and an independent color-wise baseline with the same one-color marginals either produces collisions or, if collisions are projected to a single output color, must lose aggregate admission mass. This is a mechanism check on synthetic affinity profiles; the clustering-quality audit appears in Section 9. For each $L \in \{2, 3, 5, 10\}$ we generate one complete-graph instance with $n = 150$ vertices, sample each pair's color affinities from a symmetric Dirichlet$(1, \ldots, 1)$ rescaled so that $\sum_c y^c_{uv} = 0.9$, and average over 200 seeds. The uncoupled multi-claim rate is the simulated fraction of pairs whose color-claim events occur for at least two colors under independent Bernoulli tests with means y^c_{uv}; the projection-loss column is the deterministic quantity $\sum_c y^c_{uv} - (1 - \prod_c (1 - y^c_{uv}))$ averaged over pairs, lower-bounding by Proposition 4.3 the aggregate marginal loss of any claim-respecting projection; and the coupled multi-claim rate is the simulated fraction of pairs admitted by two colors under one shared uniform per pair, which by Theorem 4.1 must be zero on every seed. Reproducibility seeds are listed in Appendix T.

Table 3: Mechanism stress test on synthetic affinity profiles. Independent rounding produces simultaneous claims at a substantial rate, while any claim-respecting projection loses at least the displayed marginal mass; coupled interval packing has zero multi-claims on every seed. This verifies the joint distribution of Proposition 4.3 and does not measure CCC approximation quality.

L	Uncoupled multi-claim	Std. dev. (algo.)	Projection loss	Coupled multi-claim
2	0.1351	0.0032	0.1354	0.0000
3	0.1784	0.0035	0.1903	0.0000
5	0.2066	0.0039	0.2375	0.0000
10	0.2207	0.0042	0.2725	0.0000

The multi-claim rate grows with L for the independent tests, while the coupled rate is zero on every seed. This synthetic check tests potential cross-color claims; it does not use feasible CCC-LP solutions or measure output clustering quality.

H CERTIFIED SMALL-INSTANCE OPTIMA

Table 4: Reported mean ALG/OPT ratios on small complete CCC instances; OPT refers to solver-certified integral optima.

Family	n	L	Xiu/OPT	Alg. 1/OPT	Comonotone/OPT
Planted chromatic blocks	10	3	1.31	1.29	1.24
Random complete colored	10	4	1.47	1.46	1.42
Witness-augmented	12	3	1.58	1.57	1.49
Negative-dense control	10	3	1.18	1.17	1.29

I PROOF OF LEMMA 6.1

Condition on u being the first pivot to touch $\{u, v\}$ and write $\pi_u^c = \lambda_u^c$ for the probability that u chooses color c. If u chooses the true color d, then the positive edge is bad exactly when v is not admitted, which under $g(x) = 1 - x$ has probability x_{uv}^d. If u chooses any wrong color $c \neq d$, then the edge is bad with probability 1: if v is admitted, the edge is realised under the wrong color, and if v is not admitted, the edge is separated. Hence the endpoint-first bad probability conditioned on u being first is

$$\pi_u^d x_{uv}^d + \sum_{c \neq d} \pi_u^c.$$

The first term is at most x_{uv}^d. For the second term, node balance gives $\sum_c \lambda_u^c = 1$, while monotonicity on the true-color edge gives $x_{uv}^d \geq x_u^d$, hence $\lambda_u^d = 1 - x_u^d \geq 1 - x_{uv}^d$ and $\sum_{c \neq d} \pi_u^c = 1 - \lambda_u^d \leq x_{uv}^d$. Therefore the endpoint contribution from u being first is at most $2x_{uv}^d$. The same bound holds when v is first, and at most one endpoint can be first.

J PROOF OF PROPOSITION 4.3

For coupled interval packing, place disjoint half-open intervals of lengths a_c inside $[0, 1)$ and draw one uniform point. Membership in interval c has probability a_c, disjointness gives mutual exclusivity, and the union length is A. For independent claims, the displayed probability is the complement of the events "no color claims" and "exactly one color claims". Finally, an exclusivity-enforcing post-processing that does not invent claims can output a color only on the event $\cup_c F_c$, whose probability is $1 - \prod_c (1 - a_c)$. Subtracting from the target total A gives the stated lower bound on lost mass, which is non-negative because $1 - \prod_c (1 - a_c) \leq \sum_c a_c$ by the union bound.

K PROOF OF THEOREM D.1

Given a CC instance I with LP separation variables $\{z_{uv}\}$, build the CCC instance $I^{(d)}$ by keeping the same vertex set, setting $\varphi(uv) = d$ for each positive edge and $\varphi(uv) = \gamma$ for each negative edge. Embed the CC solution into the CCC LP by taking $x_u^d = 0$, $x_u^c = 1$ for $c \neq d$, $x_{uv}^d = z_{uv}$, and $x_{uv}^c = 1$ for $c \neq d$. Box constraints, monotonicity, and triangle inequalities each reduce either to a CC-LP constraint in color d or to a trivial inequality in colors $c \neq d$, and node balance holds because $0 + (L - 1) \cdot 1 = L - 1$. Direct computation gives $\mathrm{LP}_{\mathrm{CCC}}(I^{(d)}) = \mathrm{LP}_{\mathrm{CC}}(I)$ because $x_{uv}^c = 1$ for $c \neq d$ contributes zero to negative-edge terms.

Conversely, take any feasible CCC-LP solution on the d-monochromatic instance and define a CC-LP solution by $z'_{uv} := x_{uv}^d$. The box and triangle constraints for color d make z' feasible for the CC LP. For positive edges the CC and CCC charges both use x_{uv}^d. For negative edges the CC charge is $1 - z'_{uv} = 1 - x_{uv}^d = y_{uv}^d$, which is at most $\sum_c y_{uv}^c$, the CCC negative-edge charge. Therefore every feasible CCC solution projects to a CC solution of no larger objective value, so $\mathrm{LP}_{\mathrm{CC}} \leq \mathrm{LP}_{\mathrm{CCC}}$. The canonical embedding gives the reverse inequality, hence the LP optima are equal.

On the integral side, every CC clustering induces a CCC solution by labeling every cluster with color d at the same disagreement cost, and any integral solution to $I^{(d)}$ can be recolored so that

every nonempty cluster has color d without increasing cost since the only positive color is d and negative-edge costs are independent of cluster color. Hence $\mathrm{OPT}_{\mathrm{CCC}}(I^{(d)}) = \mathrm{OPT}_{\mathrm{CC}}(I)$. Under the embedding, the LP-induced color distribution of every active vertex satisfies $\lambda_w^d = 1 - x_w^d = 1$, so Algorithm 1 selects color d at every pivot with probability one. The resulting execution coincides with the induced single-color CC pivot process, so $\mathrm{ALG}_{\mathrm{CCC}}(I^{(d)}) = \mathrm{ALG}_{\mathrm{CC}}(I)$, and taking the supremum over I yields the claimed equality of worst-case ratios.

L Conditional transfer on monochromatic restrictions

The embedding theorem in Appendix D establishes an exact correspondence between d-monochromatic CCC instances and CC instances. A weaker companion statement shows that any single-color CC charging inequality transfers directly to the coupled scheme on monochromatic instances without an explicit reconstruction of the CC instance.

Assumption L.1 (Single-color charging inequality)**.** *Fix a rounding function g. For every complete signed-graph instance with positive and negative labels, and every metric triple (u, v, w), the induced pivot rounding rule with join probability $g(z_{uv})$ satisfies a local charging inequality with factor α.*

Theorem L.2 (Conditional transfer)**.** *If Assumption L.1 holds for some g and α, then Algorithm 1 restricted to any d-monochromatic instance with the pivot color fixed to d satisfies the same local charging inequality with the same factor α.*

Proof. On a d-monochromatic instance only color d can satisfy any positive edge. Fixing the pivot color deterministically to d, Algorithm 1 admits the pair (u, v) to the pivot cluster with probability $g(x_{uv}^d)$ by Lemma 4.2, and no other color is selected. The coupled algorithm therefore reduces to the single-color pivot process of Assumption L.1, every triple has identical local probabilities, and the factor α is inherited. □

Combined with Theorem D.1, this shows that on monochromatic instances the coupled framework is both structurally and analytically indistinguishable from classical CC.

M Proof of Lemma 6.4

By monotonicity on the positive edge uv, $x_u^d \le x_{uv}^d$ and $x_v^d \le x_{uv}^d$, so $\lambda_u^d = 1 - x_u^d \ge 1 - x_{uv}^d$ and $\lambda_v^d \ge 1 - x_{uv}^d$. Since $\sum_c \lambda_u^c = \sum_c \lambda_v^c = 1$, this gives $\sum_{c \ne d} \lambda_u^c \le x_{uv}^d$ and $\sum_{c \ne d} \lambda_v^c \le x_{uv}^d$. For every wrong color $c \ne d$, monotonicity on wu gives $x_{wu}^c \ge x_u^c$, hence $y_{wu}^c \le \lambda_u^c$; similarly $y_{wv}^c \le \lambda_v^c$. Therefore

$$\sum_{c \ne d} \lambda_w^c \, y_{wu}^c \, y_{wv}^c \;\le\; \sum_{c \ne d} y_{wu}^c \, y_{wv}^c \;\le\; \sum_{c \ne d} y_{wu}^c \;\le\; \sum_{c \ne d} \lambda_u^c \;\le\; x_{uv}^d,$$

using $\lambda_w^c \le 1$, $y_{wv}^c \le 1$, and the two inequalities derived above.

N Proof of Theorem 7.1

Let the vertex variables be $x_u^d = x_v^d = \varepsilon$, $x_u^c = x_v^c = 1-\varepsilon$, $x_w^d = (1-\varepsilon)/2$, $x_w^c = (1+\varepsilon)/2$, which satisfy node balance for $L = 2$. Define the pair variables by $x_{uv}^d = \varepsilon$, $x_{uv}^c = 1-\varepsilon$, $x_{wu}^d = (1-\varepsilon)/2$, $x_{wv}^d = (1+\varepsilon)/2$, and $x_{wu}^c = x_{wv}^c = 1-\varepsilon$. Assign the remaining edges wu and wv arbitrary labels, for example as negative edges; the obstruction concerns only the positive edge uv and the feasibility of the LP variables used by the rounding rule.

For $\varepsilon \le 1/3$, monotonicity holds on every edge and color. For example, $x_{wu}^d = (1-\varepsilon)/2 \ge \varepsilon = x_u^d$ because $\varepsilon \le 1/3$, and $x_{wu}^c = 1 - \varepsilon \ge (1+\varepsilon)/2 = x_w^c$ because $1 - \varepsilon \ge (1+\varepsilon)/2$ iff $\varepsilon \le 1/3$. The remaining monotonicity inequalities are immediate. Triangle inequalities also hold: in color d, the only tight inequality is $x_{uv}^d + x_{wu}^d = \varepsilon + (1-\varepsilon)/2 = (1+\varepsilon)/2 = x_{wv}^d$, and the other two are slack; in color c, all three pair variables equal $1-\varepsilon$, so the triangle inequalities are trivial. The construction is therefore feasible for the full standard CCC LP.

Now consider Algorithm 1. The first pivot is uniform over $\{u, v, w\}$, so w is first with probability $1/3$. Conditioned on w being the pivot, w chooses color d with probability $\lambda_w^d = 1 - x_w^d = (1+\varepsilon)/2$. Conditioned in turn on $c^* = d$, the admission probabilities for u and v are $g(x_{wu}^d) = (1+\varepsilon)/2$ and $g(x_{wv}^d) = (1-\varepsilon)/2$ respectively, and the two pivot-pair uniforms are independent. The probability that w admits exactly one of u, v is therefore

$$\frac{1+\varepsilon}{2} \cdot \frac{1+\varepsilon}{2} + \frac{1-\varepsilon}{2} \cdot \frac{1-\varepsilon}{2} = \frac{(1+\varepsilon)^2 + (1-\varepsilon)^2}{4} = \frac{1+\varepsilon^2}{2}.$$

On this event uv is separated by a third pivot whose color choice was correct, so the edge becomes bad. Multiplying through gives

$$\tau_{uv}^+ \;\geq\; \frac{1}{3} \cdot \frac{1+\varepsilon}{2} \cdot \frac{1+\varepsilon^2}{2} = \frac{(1+\varepsilon)(1+\varepsilon^2)}{12},$$

and dividing by $x_{uv}^d = \varepsilon$ yields the displayed ratio. Since $(1+\varepsilon)(1+\varepsilon^2)/12 > 0.06\,\varepsilon$ for every $\varepsilon \in (0, 1/3]$, the proposed 0.06 third-pivot bound required by the corrected 2.06 reduction fails throughout the feasible parameter range considered here. Table 5 reports exact evaluations at representative ε, illustrating the divergence as $\varepsilon \to 0$ and the large separation from the conditional 0.06 target.

Table 5: Full-LP obstruction values from Theorem 7.1. With the corrected endpoint-first charge, the desired 2.06 reduction would require $\tau_{uv}^+ \leq 0.06\, x_{uv}^d$. The construction violates this target throughout the displayed range.

$\varepsilon = x_{uv}^d$	Lower bound on τ_{uv}^+	τ_{uv}^+ / x_{uv}^d	$0.06\, x_{uv}^d$
0.01	0.0842	8.42	0.0006
0.03	0.0859	2.86	0.0018
0.05	0.0877	1.75	0.0030
0.08	0.0906	1.13	0.0048
0.10	0.0926	0.93	0.0060

O Proof of Lemma 6.5

For each $c \neq d$, let A_u^c, A_v^c denote admission when w chooses c. The union bound applies regardless of how the two events are coupled. Monotonicity gives $y_{wu}^c \leq \lambda_u^c$ and $y_{wv}^c \leq \lambda_v^c$. Moreover $\sum_{c \neq d} \lambda_u^c \leq x_{uv}^d$ and similarly for v, by node balance and monotonicity on uv. Therefore

$$\sum_{c \neq d} \lambda_w^c \Pr(A_u^c \cup A_v^c \mid c) \leq \sum_{c \neq d} \lambda_w^c (y_{wu}^c + y_{wv}^c) \leq \sum_{c \neq d} (\lambda_u^c + \lambda_v^c) \leq 2 x_{uv}^d.$$

This bounds the immediate union only, and the double-capture subevent is separately bounded in Lemma 6.4 for Algorithm 1.

P All-colors-active witness

Start with the two-color witness from Appendix N, expressed in affinities $y_z^a = 1 - x_z^a$ and $y_{zz'}^a = 1 - x_{zz'}^a$. For $L \geq 3$, put $k = L - 2$ and $\delta = \varepsilon/(2k)$. Keep every d-affinity unchanged; subtract $\varepsilon/2$ from every vertex and pair affinity of the old wrong color c; and assign affinity δ to every vertex and pair in each of the k added colors. All affinities remain positive for $0 < \varepsilon \leq 1/3$. The node-affinity sums stay one and the pair-affinity sums stay unchanged because $k\delta = \varepsilon/2$. In color c all pair affinities become $\varepsilon/2$; its pivot vertex affinity becomes $(1-2\varepsilon)/2$, at least $\varepsilon/2$ when $\varepsilon \leq 1/3$, and the other vertices have affinity $\varepsilon/2$. Thus monotonicity holds. Each added color has constant vertex and pair affinity δ, so monotonicity and pair triangles hold; the d constraints are unchanged. The construction is feasible for the full standard LP.

The independent rule's true-color splitting event is unchanged and still contributes at least $(1+\varepsilon)(1+\varepsilon^2)/12$. For any endpoint coupling preserving the new marginals, the minimum conditional true-color disagreement is ε, the minimum modified-c union is $\varepsilon/2$, and each new-color union is δ,

all attained simultaneously by the shared threshold. Their pivot-color probabilities are $(1+\varepsilon)/2$, $(1-2\varepsilon)/2$, and δ for each new color. After multiplying by the first-pivot probability $1/3$, the exact immediate loss is

$$\frac{1}{3}\left[\frac{1+\varepsilon}{2}\varepsilon + \frac{1-2\varepsilon}{2}\frac{\varepsilon}{2} + k\delta^2\right] = \varepsilon\left(\frac{1}{4} + \frac{\varepsilon}{12k}\right).$$

Q Proof of Theorem F.1

First, (ST) holds in any integral colored clustering: if v lies in a cluster of color c, it is the strengthened triangle inequality for the separation indicators of uv, vw, uw; if not, the rightmost $x^c_v = 1$ and the pair variables involving v equal one, so the inequality holds. It therefore holds throughout the integral hull.

For (PC), consider an integral clustering and fix a positive edge uv of color d and a third vertex w. If w lies in a d-colored cluster, every wrong-color term vanishes. The absolute-difference term equals one exactly when one endpoint, but not the other, shares that cluster, in which case $x^d_{uv} = 1$. If w lies in a cluster of color $c \neq d$, the true-color term vanishes and only the c term in the sum can be nonzero. It equals one exactly when at least one endpoint shares w's wrong-colored cluster, in which case uv is separated or realized under a wrong color and $x^d_{uv} = 1$. Thus (PC) holds for every integral solution. Its left side is convex in the LP variables, being a sum of an absolute value and maxima of affine functions, while its right side is linear; Jensen's inequality extends validity to the integral hull.

Strictness relative to standard LP plus (ST) follows from an explicit $L = 2$ point. Its vertex variables are

$$(x^d_u, x^d_v, x^d_w) = (0, \tfrac{1}{2}, \tfrac{1}{2}), \qquad (x^c_u, x^c_v, x^c_w) = (1, \tfrac{1}{2}, \tfrac{1}{2}),$$

and its pair variables, ordered as (uv, uw, vw), are

$$(x^d_{uv}, x^d_{uw}, x^d_{vw}) = (\tfrac{1}{2}, \tfrac{1}{2}, 1), \qquad (x^c_{uv}, x^c_{uw}, x^c_{vw}) = (1, 1, \tfrac{1}{2}).$$

Direct substitution verifies box constraints, node balance, monotonicity, all ordinary triangles, and every orientation of (ST) in both colors. Yet for uv of true color d and third vertex w, (PC)'s left side is $|1/2 - 0| + \max\{0, 1/2\} = 1 > x^d_{uv} = 1/2$.

Finally, on the original witness the (PC) left side is $\varepsilon + \varepsilon = 2\varepsilon > x^d_{uv}$. In the all-colors-active construction, the true-color difference contributes ε, the modified wrong color $\varepsilon/2$, and the added colors $k\delta = \varepsilon/2$, again totaling 2ε. For a PC-LP-feasible point, multiplying each nonnegative term in (PC) by its pivot-color probability cannot increase the sum, so (4) is at most x^d_{uv}; weighting by the mutually exclusive first-third-pivot probabilities gives the immediate residual bound. The no-capture continuation and negative-edge events are not in this inequality.

R An exact staircase in the balanced random-color model

What can be proved exactly under a probabilistic color model rather than in the worst case is a staircase formula for the per-edge charging ratio. We state the result with matching equality in both the assumption and the conclusion; an inequality version follows by replacing every equality below with the corresponding inequality.

Definition R.1 (Linear neutral-penalty law)**.** *A family of instances obeys a linear neutral-penalty law with parameters (α, β) if for every positive edge under a fixed processing color c the expected charging ratio equals α on non-neutral edges and $\alpha + \beta$ on neutral edges.*

Theorem R.2 (Exact staircase)**.** *Under the balanced random-color model and a linear neutral-penalty law with parameters (α, β), the expected per-edge charging ratio under a fixed processing color equals $\alpha + \frac{L-1}{L}\beta$. The excess over α is monotone in L, equals 0 at $L = 1$, and tends to β as $L \to \infty$.*

Proof. A positive edge is neutral with probability $(L-1)/L$ in the balanced model, since $L-1$ of the L uniform color choices do not match. Under the linear neutral-penalty law, the expected ratio is $\frac{1}{L}\alpha + \frac{L-1}{L}(\alpha+\beta) = \alpha + \frac{L-1}{L}\beta$. The identity $\frac{L-1}{L} = 1 - \frac{1}{L}$ is monotone in L, equals 0 at $L = 1$, and tends to 1, giving the three properties. □

This is deliberately narrower than a worst-case staircase: it is exact within the balanced random-color model and applies to any analysis where neutral edges contribute a fixed additive excess.

S ADDITIONAL ROUNDING DETAILS

Lemma S.1 (Feasibility of interval packing)**.** *For every pair (w, v), step 7 of Algorithm 1 is feasible.*

Proof. For every color c, $\ell^c_{wv} = g(x^c_{wv}) \le 1 - x^c_{wv} = y^c_{wv}$, so $\sum_c \ell^c_{wv} \le \sum_c y^c_{wv} \le 1$ by Lemma 3.1. Any finite family of non-negative lengths summing to at most 1 packs into pairwise disjoint half-open subintervals of $[0, 1)$; for example, fix any global ordering of $\mathcal{L}$ and take $I^c_{wv} = [a_c, a_c + \ell^c_{wv})$ where a_c is the cumulative sum of lengths for colors preceding c. □

Remark S.2 (Endpoint budget after the corrected charge)**.** *The positive-edge endpoint-first calculation in Lemma 6.1 has two terms: a true-color non-admission term bounded by x^d_{uv} and a wrong-color pivot-choice term bounded by x^d_{uv} through node balance and monotonicity. This gives the safe $2x^d_{uv}$ endpoint charge. Consequently, a 2.06 proof using this decomposition would need the positive third-pivot residual to satisfy $\tau^+_{uv} \le 0.06\, x^d_{uv}$, not $1.06\, x^d_{uv}$.*

T REPRODUCIBILITY NOTE FOR TABLE 3

The mechanism stress test is reproduced exactly by the following protocol. Fix $n = 150$ and $L \in \{2, 3, 5, 10\}$. For each L, generate one complete-graph instance on n vertices using random seed $1000 + L$, and for each unordered pair uv sample $(y^{c_1}_{uv}, \dots, y^{c_L}_{uv})$ from a symmetric Dirichlet$(1, \dots, 1)$ distribution rescaled so that $\sum_c y^c_{uv} = 0.9$. For each algorithmic seed $s \in \{0, \dots, 199\}$, evaluate both mechanisms using random seed $20000 + 100L + s$. The uncoupled baseline samples L independent Bernoulli claims with means $y^{c_1}_{uv}, \dots, y^{c_L}_{uv}$ and records the fraction of pairs claimed by at least two colors. The projection-loss column is the deterministic quantity $\sum_c y^c_{uv} - (1 - \prod_c (1 - y^c_{uv}))$ averaged over pairs of the same instance. The coupled baseline samples one shared uniform in $[0, 1)$ per pair, assigns disjoint intervals of the corresponding lengths, and records the same multi-claim quantity, which is identically zero by construction. Empirical means and sample standard deviations over the 200 algorithmic seeds give the values in Table 3; standard deviations therefore reflect algorithmic randomness, not the variability across affinity profiles.