

Degree-Corrected Energy Relaxation Bounds for Competitive Recoloring on Complex Networks

Ranjan Veerabhadraswamy

VIT-AP University

Ajith Jubilson Emerson

ajith.jubilson@vitap.ac.in

VIT-AP University

Article

Keywords: complex networks, stochastic dynamics, discrete Lyapunov function, finite-time relaxation, potential games, graph recoloring

Posted Date: October 15th, 2026

DOI: <https://doi.org/10.21203/rs.3.rs-10874007/v1>

License:   This work is licensed under a Creative Commons Attribution 4.0 International License.

[Read Full License](#)

Additional Declarations: No competing interests reported.

Degree-Corrected Energy Relaxation Bounds for Competitive Recoloring on Complex Networks

Ranjan Veerabhadraswamy¹ and Ajith Jubilson Emerson
(corresponding author)^{1*}

¹School of Computer Science and Engineering, Vellore Institute of Technology, Andhra Pradesh (VIT-AP), Amaravati, Andhra Pradesh, 522241, India.

*Corresponding author(s). E-mail(s): ajith.jubilson@vitap.ac.in;
Contributing authors: ranjanvswamyjnv2005@gmail.com;

Abstract

Asynchronous competitive recoloring is a simple model of decentralized conflict reduction on a network: at each activation, one vertex changes state only when a new color strictly reduces its number of same-colored neighbors. Exact-potential convergence of this process is known, but that fact alone does not explain how network structure constrains relaxation. We establish a degree-corrected inequality that links excess conflict energy to the weighted population of vertices that can improve. Under independent uniform activation, the inequality yields a topology-sensitive harmonic upper certificate for the expected time to reach an equilibrium-compatible energy level. The same mechanism extends to fixed heterogeneous activation rates and gives an explicit finite-confidence relaxation certificate. When the number of colors exceeds the maximum degree, the energy level is zero and the argument bounds full expected absorption. Numerical experiments on 9,600 Erdős–Rényi, Barabási–Albert, and Watts–Strogatz trajectories support the expectation-level certificate across size, density, topology, and palette size. Barabási–Albert networks produce substantially larger certificates than the other ensembles without comparably slower observed relaxation, indicating that the worst-case structural factor can be highly conservative in hub-dominated topologies.

Keywords: complex networks, stochastic dynamics, discrete Lyapunov function, finite-time relaxation, potential games, graph recoloring

1 Introduction

Many networked systems organize through repeated local adjustments rather than centralized optimization. An agent observes its neighborhood, changes state according to a local objective, and thereby alters the environment of adjacent agents. The resulting non-equilibrium trajectory depends on both the update rule and the interaction graph, raising a basic question: what network structure can be converted into a quantitative relaxation guarantee?

Competitive graph recoloring provides a transparent setting. Vertices are agents, colors are discrete states, and edges represent pairwise interactions. A vertex minimizes the number of same-colored neighbors and changes only under strict improvement, defining an anti-coordination game. Related models arise in distributed graph coloring, local allocation, and graphical games [1–3]. The abstraction separates interaction topology, deterministic local response, and stochastic asynchronous activation.

Network structure influences relaxation through size, density, clustering, and degree heterogeneity. Hubs can produce large local effects while making worst-case estimates depend on only a few vertices. Homogeneous Erdős–Rényi (ER), heterogeneous Barabási–Albert (BA), and clustered Watts–Strogatz (WS) networks represent these features differently [4–6]. We ask how such structure constrains asynchronous competitive recoloring.

The classical foundation is the theory of potential games. Rosenthal introduced an early potential construction for congestion games [7], and Monderer and Shapley formalized exact and ordinal potential games [8]. In competitive recoloring, the number of monochromatic edges is an exact potential: a strict local improvement changes the global potential by exactly the same amount. This identity immediately rules out improvement cycles and guarantees finite convergence to a pure Nash equilibrium. MAX- k -CUT and graph-coloring games provide the closest algorithmic context [9–13]; related work studies approximation and convergence in potential games [14, 15], distributed selfish MAX-CUT dynamics [16], and locality-sensitive convergence time in graphical games [17]. Networked anti-coordination dynamics have also been studied through graphical dynamical systems, stochastic decision models, and mixed coordination games [18–21], while playing-sequence effects and stochastic strategy revision provide broader context for scheduler-dependent dynamics [22–24]. Thus neither the exact-potential identity nor the classical $|E|$ accepted-move bound is claimed as new here. The new contribution is the state-dependent, degree-corrected inequality in Theorem 1, which converts excess energy into a weighted supply of improving vertices and thereby yields an explicit expected relaxation certificate under uniform asynchronous activation.

A potential argument establishes energy dissipation but does not quantify the probability of improvement under random activation. We supply that connection through a degree-corrected instability inequality: above an equilibrium-compatible energy ceiling, excess energy forces a weighted population of improving vertices. Uniform activation then gives a harmonic expected hitting-time certificate; fixed heterogeneous activation gives expectation and finite-confidence extensions; and when the palette exceeds the maximum degree, the ceiling vanishes and the bound controls absorption. A balanced replay of 9,600 trajectories across three network families, four

sizes, five palettes, and four target mean degrees evaluates where the worst-case degree factor loses information. 46
47

2 Results 48

2.1 Degree-corrected instability mechanism 49

Let $G = (V, E)$ be a finite simple undirected graph, $n = |V|$, $m = |E|$, and $C = \{1, \dots, k\}$ with $k \geq 2$. Write $d_i = \deg(i)$. For a coloring $c \in C^V$, define the local conflict count of vertex i under candidate color γ and the global conflict energy by 50
51
52

$$x_i(\gamma; c) = \sum_{j \in N(i)} \mathbf{1}\{\gamma = c_j\}, \quad \Phi(c) = \sum_{\{u, v\} \in E} \mathbf{1}\{c_u = c_v\}.$$

When activated, vertex i chooses a minimizer of $x_i(\gamma; c)$ and changes only if that minimum is strictly below $x_i(c_i; c)$. Ties among improving colors are arbitrary. Let $I(c)$ be the set of vertices with a strict improving recoloring. Under the stochastic scheduler, vertices are selected independently and uniformly from V ; rejected activations leave the configuration unchanged. 53
54
55
56
57

The following classical foundation is included for self-containment and attributed to the exact-potential MAX- k -CUT literature [8–10]. 58
59

Proposition 1 (Known potential foundation) *If vertex i changes from α to β by strict improvement, then* 60
61

$$\Phi(c^+) - \Phi(c) = x_i(\beta; c) - x_i(\alpha; c) \leq -1.$$

Hence every strict-improvement path has at most $\Phi(c_0) \leq m$ accepted moves. Under independent uniform activation, an equilibrium is reached almost surely. 62
63

For each vertex define 64

$$b_i = \left\lfloor \frac{d_i}{k} \right\rfloor, \quad w_i = d_i - b_i, \quad S = \sum_{i \in V} b_i, \quad H = \left\lfloor \frac{S}{2} \right\rfloor, \quad W = \max_{i \in V} w_i.$$

Here b_i is the largest local conflict count that can be guaranteed by choosing a least frequent color in the neighborhood, and w_i is the remaining degree-dependent weight. The quantity H is a degree-corrected upper bound on equilibrium energy. It is not asserted to be tight: it is a topology-dependent ceiling derived from the degree sequence and palette size. 65
66
67
68
69

Theorem 1 (Degree-corrected instability and relaxation) *Let $m > 0$ and let $I(c)$ be as above. Every configuration satisfies* 70
71

$$\sum_{i \in I(c)} w_i \geq 2\Phi(c) - S. \tag{1}$$

72 Consequently, every equilibrium c^* satisfies $\Phi(c^*) \leq H$. Under independent uniform
 73 activation, set

$$\tau_H = \inf\{t \geq 0 : \Phi(c(t)) \leq H\}, \quad X_0 = (\Phi(c_0) - H)_+,$$

74 and $\varepsilon = S - 2H \in \{0, 1\}$. Then $\tau_H < \infty$ almost surely and

$$\mathbb{E}[\tau_H \mid c_0] \leq B_H := nW \sum_{r=1}^{X_0} \frac{1}{2r - \varepsilon}. \quad (2)$$

75 The empty sum is zero.

76 *Proof* Put $x_i = x_i(c_i; c)$. Each monochromatic edge is counted at both endpoints, so $2\Phi(c) =$
 77 $\sum_i x_i$. If $i \notin I(c)$, its current color is among the least frequent colors in its neighborhood.
 78 Since d_i neighbors are distributed among k colors, $x_i \leq \lfloor d_i/k \rfloor = b_i$. For $i \in I(c)$, trivially
 79 $x_i \leq d_i = b_i + w_i$. Therefore

$$2\Phi(c) \leq \sum_{i \notin I(c)} b_i + \sum_{i \in I(c)} (b_i + w_i) = S + \sum_{i \in I(c)} w_i,$$

80 which proves (1). At equilibrium $I(c^*) = \emptyset$; hence $2\Phi(c^*) \leq S$ and integrality gives $\Phi(c^*) \leq$
 81 H .

82 Suppose $\Phi(c) = H + r$ with $r \geq 1$. Since $S = 2H + \varepsilon$, (1) and $w_i \leq W$ imply

$$|I(c)| \geq \frac{2r - \varepsilon}{W}.$$

83 A uniform activation is therefore improving with probability at least $(2r - \varepsilon)/(nW)$. Rejected
 84 activations do not change the state, while an accepted activation strictly lowers the integer-
 85 valued potential and leaves level $H + r$. The expected residence time at that level is at most
 86 $nW/(2r - \varepsilon)$. Since the potential never increases, each positive excess level is visited at most
 87 once, although some may be skipped. Summing these residence-time bounds for $r = 1, \dots, X_0$
 88 proves (2) and almost-sure threshold attainment. $\square$

89 Inequality (1) states that a configuration cannot remain far above the equilibrium-
 90 compatible energy ceiling while containing only a negligible amount of locally improv-
 91 ing structure. The right-hand side measures excess energy, whereas the left-hand side
 92 measures instability through degree-corrected vertex weights. This is stronger than
 93 knowing only that an accepted update decreases Φ : it supplies a lower bound on the
 94 population from which an improving activation may be drawn.

95 As an independent finite-instance sanity check, the structural inequality and equi-
 96 librium ceiling were exhaustively verified for every coloring with $k \in \{2, 3\}$ on all
 97 connected non-isomorphic simple graphs with two to five vertices. The check covered
 98 6,440 configurations and found no violation (Supplementary Note S1).

99 2.2 Expected relaxation-time certificate

100 The hitting-time estimate in equation (2) separates three contributions. The factor
 101 n is the inverse scale of uniform vertex sampling. The quantity X_0 is the initial
 102 energy excess above H , so it records how many positive excess levels may need to be
 103 crossed. The factor W is a worst-case local structural weight determined by the largest
 104 degree-corrected vertex contribution. Finally, the harmonic sum reflects the progres-
 105 sive reduction in the guaranteed improving population as the process approaches the

ceiling. Relaxation can therefore be certified quickly when many weighted improving vertices are forced by the current energy, but the guarantee becomes more conservative near H .

The bound is conditional in expectation on the initial state c_0 . It is not a trajectory-wise ceiling and is not intended as a point prediction of the observed time. Individual realizations may exceed their corresponding numerical value of B_H without contradicting the theorem. The appropriate empirical comparison is consequently between conditional or controlled-group averages and their certificates. This distinction is central when the bound is evaluated below.

2.3 High-palette absorption

Corollary 1 (High-palette absorption) *If $k > \Delta(G)$, then $H = 0$, $W = \Delta(G)$, and the equilibrium absorption time satisfies*

$$\mathbb{E}[\tau_{\text{NE}} \mid c_0] \leq \frac{n\Delta(G)}{2} H_{\Phi(c_0)}, \quad H_q = \sum_{r=1}^q \frac{1}{r}.$$

Proof Now $b_i = 0$ for every i , so $S = H = \varepsilon = 0$. Theorem 1 shows that every equilibrium has $\Phi = 0$; conversely, $\Phi = 0$ admits no strict improvement. Thus $\tau_H = \tau_{\text{NE}}$, and (2) reduces to the displayed harmonic bound. $\square$

The condition $k > \Delta(G)$ makes every b_i vanish because each neighborhood contains fewer vertices than there are available colors. The equilibrium-energy ceiling is then zero, and reaching the ceiling is equivalent to reaching an absorbing conflict-free configuration. This gives a full expected absorption-time estimate rather than only a threshold-hitting estimate. The regime is structurally informative but represents only a small subset of the tested parameter grid, so it is analyzed separately.

2.4 Heterogeneous activation and finite-confidence relaxation

The structural inequality is not tied to uniform sampling. Let activations be independent draws from a fixed distribution $p = (p_i)_{i \in V}$ with $p_i > 0$ and $\sum_i p_i = 1$, and define

$$\lambda = \min_{i: d_i > 0} \frac{p_i}{w_i}.$$

Every improving vertex has positive degree and hence positive w_i , so this quantity is well defined whenever $m > 0$.

Corollary 2 (Heterogeneous activation) *Under the fixed activation distribution p ,*

$$\mathbb{E}[\tau_H \mid c_0] \leq B_{H,p} := \frac{1}{\lambda} \sum_{r=1}^{X_0} \frac{1}{2r - \varepsilon}. \quad (3)$$

For uniform activation, $\lambda = 1/(nW)$ and $B_{H,p} = B_H$.

135 *Proof* At level $\Phi(c) = H + r$, Theorem 1 gives $\sum_{i \in I(c)} w_i \geq 2r - \varepsilon$. Hence

$$\Pr(\text{improvement} \mid c) = \sum_{i \in I(c)} p_i \geq \lambda \sum_{i \in I(c)} w_i \geq \lambda(2r - \varepsilon).$$

136 The expected residence time at this level is therefore at most $1/[\lambda(2r - \varepsilon)]$. Summing over
 137 positive excess levels proves (3). Under uniform activation, $p_i/w_i = 1/(nw_i)$ on non-isolated
 138 vertices, whose minimum is $1/(nW)$. $\square$

139 **Corollary 3** (High-probability relaxation) *If $X_0 \geq 1$ and $\delta \in (0, 1)$, then under the fixed*
 140 *activation distribution p , with probability at least $1 - \delta$,*

$$\tau_H \leq B_{H,p} \log \frac{X_0}{\delta} + X_0. \quad (4)$$

141 *If $X_0 = 0$, then $\tau_H = 0$. In particular, uniform activation gives the same statement with*
 142 *$B_{H,p} = B_H$.*

143 *Proof* Conditional on visiting level $H + r$, its residence time is stochastically dominated by a
 144 geometric waiting time with success probability $q_r = \lambda(2r - \varepsilon)$. For every visited level, $q_r \leq 1$,
 145 because q_r is a lower bound on the actual improvement probability $\sum_{i \in I(c)} p_i \leq 1$. Applying
 146 the geometric tail bound at each of at most X_0 visited positive excess levels, followed by
 147 a union bound and $\lceil x \rceil \leq x + 1$, gives (4); details and exhaustive finite-instance checks are
 148 provided in Supplementary Note S2. $\square$

149 2.5 Numerical validation across complex-network families

150 The frozen experiment grid contains Erdős–Rényi (ER), Barabási–Albert (BA), and
 151 Watts–Strogatz (WS) networks with $n \in \{100, 200, 400, 800\}$, $k \in \{2, 3, 4, 5, 8\}$, target
 152 mean degree $\bar{d} \in \{4, 8, 12, 16\}$, and 40 independent seeded trials in every controlled
 153 cell. The balanced design contains $3 \times 4 \times 5 \times 4 \times 40 = 9,600$ trajectories and 240 controlled
 154 cells. A deterministic replay reproduced the stored graph statistics, initial and
 155 final energies, accepted moves, activation counts, and stopping states before adding
 156 the theorem-specific quantities H , W , X_0 , τ_H , and B_H . Complete controlled-cell
 157 summaries are provided in Supplementary Table S1.

158 Because equation (2) is conditional in expectation, Fig. 1a compares the mean
 159 τ_H with the mean B_H within each controlled 40-run cell; it does not treat B_H as
 160 a trajectory-wise ceiling. All 240 cell means lie below equality, and the largest ratio
 161 $\bar{\tau}_H/\bar{B}_H$ is 0.1646. Fig. 1b reports mean τ_H/B_H with descriptive bootstrap 95% confidence
 162 intervals for each topology and palette size; runs with $B_H = \tau_H = 0$ contribute
 163 zero.

164 Across all trajectories, mean τ_H is 94.60 (bootstrap 95% CI 90.83–98.57), mean
 165 B_H is 31,755.18 (bootstrap 95% CI 30,665.42–32,883.90), and mean τ_H/B_H is 0.00957
 166 (bootstrap 95% CI 0.00911–0.01005). Mean τ_H is 84.96, 86.01, and 112.82 for ER,
 167 BA, and WS, respectively. In contrast, mean B_H is 16,032, 68,934, and 10,299. BA
 168 therefore has the largest certificate by a wide margin without a correspondingly large
 169 threshold time. The high-palette subset contains 160 runs satisfying $k > \Delta(G)$; all
 170 have $\tau_H = \tau_{NE}$, with mean absorption time 1,003.35 and mean corollary certificate
 171 6,427.92. Of these runs, 159 are WS and one is ER; none is BA, so this subset does
 172 not support a topology comparison.




Fig. 1 Degree-corrected relaxation validation. **a**, Controlled-cell means of observed threshold time and theoretical certificate; the dashed line is equality. Each point represents 40 seeded trajectories. **b**, Mean normalized threshold time versus palette size for ER, BA, and WS networks. Error bars are two-sided nonparametric percentile-bootstrap 95% confidence intervals from 10,000 trajectory-level resamples within each topology-palette group. The dashed horizontal line marks equality.

2.6 Structural effects of topology, palette size, density, and size

Balanced marginal summaries reveal how the observed threshold time changes across the tested design (Fig. 2). Increasing network size raises the mean number of activations required to reach the degree-corrected ceiling in all three families. From $n = 100$ to $n = 800$, the pooled marginal mean rises from 23.96 to 179.83 for ER, from 23.89 to 181.66 for BA, and from 26.31 to 250.25 for WS. These finite-size trends are close to proportional over the tested range, but four sizes are insufficient to establish an asymptotic scaling law.

Palette size changes both the initial conflict landscape and the target ceiling. The mean threshold time increases from $k = 2$ to $k = 8$ in every family when averaged over size and density: 47.58 to 172.45 for ER, 38.84 to 143.45 for BA, and 37.44 to 296.38 for WS. The WS sequence is not monotone at intermediate palettes, illustrating why the result should not be compressed into a universal palette law. Larger k generally lowers H , so the dynamics are asked to reach a more stringent energy level even though more colors can reduce local conflicts.

The density trend is reversed over the tested grid. As target mean degree rises from 4 to 16, pooled mean τ_H falls from 175.02 to 43.22 for ER, from 164.04 to 41.66 for BA, and from 295.81 to 37.43 for WS. This is a descriptive property of the present threshold and parameterization rather than evidence that density universally accelerates recoloring. In particular, H increases with the degree sequence, while the mean initial excess X_0 remains of similar order across the tested densities; both features affect the measured hitting time.

Degree heterogeneity is most visible in the structural factors. Across all runs, BA has mean degree variance 73.95 and mean $W = 49.33$, compared with 9.36 and 14.37 for ER, and 0.94 and 9.78 for WS. The controlled-cell scatter in Fig. 2d shows that




Fig. 2 Structural determinants of observed relaxation. Marginal mean threshold time τ_H versus **a**, network size; **b**, palette size; and **c**, target mean degree, shown separately for ER, BA, and WS networks. The balanced design gives 800 trajectories per topology–size point, 640 per topology–palette point, and 800 per topology–degree point. Error bars are two-sided nonparametric percentile-bootstrap 95% confidence intervals from 10,000 resamples. **d**, Controlled-cell mean structural weight W versus the certificate-to-time ratio $\bar{B}_H/\bar{\tau}_H$; each point represents 40 trajectories and the vertical axis is logarithmic. All panels use only the frozen 9,600-run replay dataset.

large W is associated with very conservative certificate-to-time ratios, with BA occupying most of the high- W region. This comparison is descriptive and combines several controlled factors; it does not identify degree heterogeneity as a causal determinant of the observed trajectory time. It does, however, locate a clear source of looseness in the theorem: replacing the entire weight profile by its maximum substantially enlarges B_H on hub-dominated networks.

3 Discussion

The main result converts a qualitative Lyapunov argument into a state-dependent relaxation certificate. Exact-potential descent already guarantees that accepted recolorings cannot cycle and that only finitely many accepted moves occur. The degree-corrected inequality adds information that the potential alone does not provide: when conflict energy lies above an equilibrium-compatible ceiling, a weighted population

of vertices must admit a strict local improvement. Uniform asynchronous sampling then turns that population bound into a lower bound on the probability of an energy-decreasing activation. Summing the expected residence times across excess-energy levels yields the harmonic certificate. In this way, global energy, local instability, and network degree structure enter one quantitative argument.

The scheduler extensions clarify which part of this mechanism is structural and which part is probabilistic. For fixed heterogeneous activation, λ couples update rates to degree-corrected weights: a very small activation probability assigned to a heavily weighted vertex can weaken the guarantee even though convergence remains almost sure. The finite-confidence corollary then turns the same per-level progress estimate into a trajectory-wise statement at a chosen failure probability. Its union-bound form is explicit and robust, but it is not claimed to be tight.

The certificate bounds a conditional expectation; it is not a trajectory-wise kinetic prediction. It replaces improving weights by the maximum W , uses only aggregate excess energy, and bounds every positive level although accepted recolorings may skip levels. These robust simplifications explain why observed ratios can be much smaller than one and allow the simulations to diagnose where information is lost.

BA networks illustrate this conservatism: their maximum degree, degree variance, and W greatly exceed those of ER and WS, yet their mean threshold time is close to ER and below WS. A hub strongly affects the worst-case factor without necessarily remaining a bottleneck. Bounds based on the distribution or evolving sum of the w_i values may retain information discarded by the maximum, although the present data do not establish a sufficient degree statistic.

Across the tested range, threshold time grows with network size while normalized certificate ratios decline; four sizes do not establish an asymptotic exponent. Larger palettes generally increase the time to reach the lower ceiling H , with a non-monotone WS response at intermediate k . Higher target mean degree lowers the measured hitting time in this grid, but also changes the target itself. Because ER, BA, and WS differ in several statistics simultaneously, these comparisons are descriptive rather than causal; isolating degree heterogeneity or clustering would require matched ensembles or targeted rewiring.

The scope is finite, simple, static, undirected, and unweighted graphs; temporal interactions [25], state-dependent schedulers, and nonstationary activation are not treated. Updates are strict best responses, experiments use synthetic networks at four sizes, and the theorem controls threshold hitting rather than general absorption outside $k > \Delta(G)$. Degree-profile certificates, directed or weighted interactions, temporal networks, and sharper concentration are natural extensions. They should preserve the central separation used here: network structure identifies available local descent, and the scheduler converts it into a relaxation estimate.

4 Methods

4.1 Network generation

All graphs were generated with NetworkX using deterministic integer seeds. For ER graphs, each possible edge was included independently with probability $p =$

253 $\min\{1, \bar{d}/(n-1)\}$, where $\bar{d}$ is the target mean degree. For BA graphs, the preferential-
 254 attachment parameter was $q = \max\{1, \min(n-1, \text{round}(\bar{d}/2))\}$ edges per arriving
 255 vertex. The resulting mean degree is approximately $2q$ away from the finite initial con-
 256 dition. For WS graphs, vertices began on a ring lattice with the even neighborhood
 257 parameter $q = 2 \text{round}(\bar{d}/2)$, clipped to the valid range and reduced by one if neces-
 258 sary to remain even. Each lattice edge was rewired independently with probability 0.1.
 259 These procedures are the exact implementations used to generate the frozen dataset.

260 4.2 Competitive recoloring dynamics

261 Each vertex was initialized independently and uniformly with one of k integer-labeled
 262 colors. At an activation, neighbor colors were counted and the local conflict count
 263 was evaluated for every color. The current color was retained unless another color
 264 had strictly fewer same-colored neighbors. If several alternative colors attained pro-
 265 gressively lower conflict counts, iteration over the ordered color labels selected the
 266 lowest-index minimizer. This deterministic tie rule affects reproducibility but not the
 267 Lyapunov argument. A successful change was recorded as an accepted move; sampling
 268 a vertex with no strict improvement was recorded as a rejected activation. The process
 269 stopped when the incrementally maintained set of improving vertices became empty,
 270 which is equivalent to a pure Nash equilibrium for the strict best-response rule.

271 4.3 Scheduler and reproducibility

272 At each step, one vertex was sampled independently and uniformly. A master seed
 273 of 2026 was combined with a family code, n , k , target mean degree, and trial index
 274 through NumPy's `SeedSequence`; successive generated integers seeded graph construc-
 275 tion and the recoloring trajectory. The replay stopped on any disagreement in graph
 276 statistics, energies, accepted moves, activation count, or convergence status.

277 4.4 Parameter grid

278 The balanced grid comprised three graph families; $n \in \{100, 200, 400, 800\}$; $k \in$
 279 $\{2, 3, 4, 5, 8\}$; target mean degree $\bar{d} \in \{4, 8, 12, 16\}$; and 40 trials per controlled cell.
 280 The total number of trajectories was therefore $3 \times 4 \times 5 \times 4 \times 40 = 9,600$. No stochas-
 281 tic simulation was rerun for the expanded analysis. Figure 2 and all new marginal
 282 summaries were computed directly from the frozen replay table.

283 4.5 Measured quantities

284 For every run, the stored quantities included the edge count, realized mean degree,
 285 degree-square sum, initial and final potential, accepted recolorings, total vertex acti-
 286 vations, and final convergence status. The replay added maximum degree $\Delta(G)$, H ,
 287 W , initial excess X_0 , parity ε , threshold time τ_H , certificate B_H , and the indica-
 288 tor $k > \Delta(G)$. Degree variance was derived without reconstructing new graphs as
 289 $n^{-1} \sum_i d_i^2 - (2m/n)^2$. The squared coefficient of degree variation was this variance
 290 divided by $(2m/n)^2$. Runs with $B_H = \tau_H = 0$ were assigned normalized time zero,

matching the frozen validation analysis. Excluding runs for which both B_H and τ_H were zero did not materially alter the normalized-time results.

4.6 Statistical analysis

The analysis was descriptive; no null-hypothesis significance tests or causal models were used. All displayed confidence intervals and global uncertainty summaries are two-sided nonparametric percentile-bootstrap 95% intervals computed from 10,000 resamples with seed 20260827. Figure 1b resamples trajectory-level ratios within each topology–palette group. The global summaries use all 9,600 trajectories, while expectation-level validation uses 240 controlled cells of 40 trials each. Figure 2 reports balanced marginal means: each topology–size and topology–degree point contains 800 runs, and each topology–palette point contains 640 runs. Relative to corresponding normal intervals, bootstrap interval widths changed by at most 2.39% overall and by at most 1.54% for Fig. 1b; no qualitative interpretation changed (Supplementary Table S2). The distributions of τ_H , B_H , and their ratio were right-skewed; the intervals are descriptive uncertainty summaries rather than evidence for parametric group tests. Figure 2d uses cell means and a logarithmic vertical axis; it is an exploratory structural comparison and does not support causal attribution.

The replay was rechecked with Python 3.12, NetworkX 3.6.1, NumPy 1.26.4, pandas 3.0.5, and Matplotlib 3.11.1. Original CSV files were read without modification, and derived outputs were written separately.

4.7 Use of language-model tools

LLM and Gen AI tools were used only for Editing the document (e.g., grammar, spelling, word choice)

4.8 Code availability

The simulation, replay, bootstrap analysis, scheduler verification, and figure-generation code is publicly available in the GitHub repository at tag `v1.1-scientific-reports-upgrade`: <https://github.com/Ranjan-V/Competitive-Graph-Recoloring/tree/v1.1-scientific-reports-upgrade>.

Data availability

The simulation dataset generated and analyzed during this study, together with the code used to generate and analyze it, is publicly available at <https://github.com/Ranjan-V/Competitive-Graph-Recoloring/tree/v1.1-scientific-reports-upgrade>. Supplementary Tables S1 and S2 are provided with the Supplementary Information. For enquiries about data access, contact the corresponding author, Ajith Jubilson Emerson (ajith.jubilson@vitap.ac.in).

References

- [1] Kuhn, F. & Wattenhofer, R. On the complexity of distributed graph coloring. *Proc. ACM Symp. Principles Distrib. Comput.* 7–15 (2006).
- [2] Panagopoulou, P. N. & Spirakis, P. G. A game theoretic approach for efficient graph coloring. *Lect. Notes Comput. Sci.* **5369**, 183–195 (2008).
- [3] Szabó, G. & Fáth, G. Evolutionary games on graphs. *Phys. Rep.* **446**, 97–216 (2007).
- [4] Erdős, P. & Rényi, A. On random graphs I. *Publ. Math. Debrecen* **6**, 290–297 (1959).
- [5] Barabási, A.-L. & Albert, R. Emergence of scaling in random networks. *Science* **286**, 509–512 (1999).
- [6] Watts, D. J. & Strogatz, S. H. Collective dynamics of small-world networks. *Nature* **393**, 440–442 (1998).
- [7] Rosenthal, R. W. A class of games possessing pure-strategy Nash equilibria. *Int. J. Game Theory* **2**, 65–67 (1973).
- [8] Monderer, D. & Shapley, L. S. Potential games. *Games Econ. Behav.* **14**, 124–143 (1996).
- [9] Gourvès, L. & Monnot, J. The Max k -Cut game and its strong equilibria. *Lect. Notes Comput. Sci.* **6108**, 234–246 (2010).
- [10] Carosi, R. & Monaco, G. Generalized graph k -coloring games. *Theory Comput. Syst.* **64**, 1028–1041 (2020).
- [11] Frieze, A. M. & Jerrum, M. Improved approximation algorithms for MAX k -CUT and MAX BISECTION. *Algorithmica* **18**, 67–81 (1997).
- [12] Garuglieri, G., Madeo, D., Mocenni, C., Palma, G. & Rinaldi, S. Optimal coloring strategies for the Max k -Cut game. *Mathematics* **12**, 604 (2024).
- [13] D’Ascenzo, A., D’Emidio, M., Flammini, M. & Monaco, G. Digraph k -coloring games: From theory to practice. *Leibniz Int. Proc. Inform.* **233**, 20:1–20:18 (2022).
- [14] Christodoulou, G., Mirrokni, V. S. & Sidiropoulos, A. Convergence and approximation in potential games. *Theor. Comput. Sci.* **438**, 13–27 (2012).
- [15] Swenson, B., Murray, R. & Kar, S. On best-response dynamics in potential games. *SIAM J. Control Optim.* **56**, 2734–2767 (2018).

- [16] Auger, D., Cohen, J., Coucheney, P. & Rodier, L. Distributed selfish algorithms for the Max-Cut game. *Lect. Notes Electr. Eng.* **264**, 45–54 (2013).
- [17] Hirvonen, J., Schmid, L., Chatterjee, K. & Schmid, S. On the convergence time in graphical games: A locality-sensitive approach. *Leibniz Int. Proc. Inform.* **286**, 11:1–11:24 (2024).
- [18] Qiu, X. *et al.* Networked anti-coordination games meet graphical dynamical systems: Equilibria and convergence. *Proc. AAAI Conf. Artif. Intell.* **37**, 11663–11671 (2023).
- [19] Zhu, Y. & Xia, C.-Y. Asynchronous best-response dynamics of networked anti-coordination game with payoff incentives. *Chaos Solitons Fractals* **172**, 113503 (2023).
- [20] Jones, M. I., Pauls, S. D. & Fu, F. The dual problems of coordination and anti-coordination on random bipartite graphs. *New J. Phys.* **23**, 113018 (2021).
- [21] Arditti, L., Como, G., Fagnani, F. & Vanelli, M. Equilibria and learning dynamics in mixed network coordination/anti-coordination games. *Proc. 60th IEEE Conf. Decis. Control* 4982–4987 (2021).
- [22] Heinrich, T. *et al.* Best-response dynamics, playing sequences, and convergence to equilibrium in random games. *Int. J. Game Theory* **52**, 703–735 (2023).
- [23] Blume, L. E. The statistical mechanics of strategic interaction. *Games Econ. Behav.* **5**, 387–424 (1993).
- [24] Young, H. P. The evolution of conventions. *Econometrica* **61**, 57–84 (1993).
- [25] Holme, P. & Saramäki, J. Temporal networks. *Phys. Rep.* **519**, 97–125 (2012).

Acknowledgements

The authors thank their mentors and colleagues for helpful feedback. This research received no external funding.

Author contributions

R.V. contributed conceptualization, methodology, software, formal analysis, and writing of the original draft. A.J.E. contributed supervision, validation, and review and editing of the manuscript. Both authors reviewed and approved the final manuscript.

Competing interests

The authors declare no competing interests.

Supplementary Files

This is a list of supplementary files associated with this preprint. Click to download.

- [SupplementaryInformation1.pdf](#)