Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Direct Reciprocity on Structured Populations: Network-Dependent Cooperation Thresholds in the Iterated Prisoner's Dilemma

Jonathan K. Corrado

corradojk@gmail.com

Research Article

Keywords: evolutionary game theory, iterated Prisoner's Dilemma, direct reciprocity, network reciprocity, structured populations, stochastic evolutionary dynamics, cooperation

Posted Date: September 16th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-10827695/v1

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

Additional Declarations: No competing interests reported.

Direct Reciprocity on Structured Populations: Network-Dependent Cooperation Thresholds in the Iterated Prisoner's Dilemma

Jonathan K. Corrado, ORCID: 0000-0002-0986-1400, USA

Independent Researcher, Cincinnati, OH 45244, USA

Correspondence: corradojk@gmail.com

Abstract

Repeated interaction can sustain cooperation through direct reciprocity, but the continuation probability required for cooperation need not be invariant to population structure. We study a finite stochastic evolutionary game in which Tit-for-Tat (TFT) and Always Defect (ALLD) evolve under synchronous Fermi updating on well-mixed, regular, Watts–Strogatz small-world, Erdős–Rényi random, and Barabási–Albert scale-free interaction structures. We use the continuation probability at which the replicate-mean long-run TFT frequency crosses 0.5 as a common transition measure. Under the baseline specification, this threshold ranges from approximately 0.702 in the canonical small-world network to 0.794 in the scale-free network. Continuous Watts–Strogatz rewiring yields a non-monotonic response, with a minimum near $p \approx 0.0215$. Degree-preserving null models show that the regular and small-world advantages depend strongly on organization beyond degree sequence, whereas scale-free behavior is comparatively insensitive to degree-preserving randomization under degree-normalized payoff. Initialization controls further show that clustered placement favors cooperation in locally organized networks, while hub-biased placement strongly lowers the scale-free threshold. The results quantify how population structure changes the direct-reciprocity requirement for cooperation within a common evolutionary-game specification.

Keywords: evolutionary game theory; iterated Prisoner's Dilemma; direct reciprocity; network reciprocity; structured populations; stochastic evolutionary dynamics; cooperation

1. Introduction

Repeated games and evolutionary games provide complementary accounts of how cooperation can persist despite short-run incentives to defect. In the iterated Prisoner's Dilemma, direct reciprocity operates through future interaction: when encounters are sufficiently persistent, contingent strategies can make cooperation evolutionarily competitive. In structured populations, however, players do not encounter all opponents symmetrically. The same repeated game is embedded in a network that determines who interacts, how often local strategy configurations recur, and which players occupy structurally influential positions.

Evolutionary game theory on graphs has shown that population structure can alter selection and cooperation. Regular neighborhoods can support cooperative domains, small-world shortcuts change the balance between local recurrence and global reach, and heterogeneous networks can give high-degree vertices disproportionate influence. These mechanisms are established. The question pursued here is narrower: how much does population structure change the level of direct reciprocity required for cooperation when the underlying game, strategies, and evolutionary update rule are held fixed?

We operationalize that question using the continuation probability δ of the repeated encounter. Increasing δ strengthens the future consequences of current play without changing the network. For each interaction structure, we define δc as the continuation probability at which the replicate-mean long-run frequency of Tit-for-Tat (TFT) crosses 0.5. This is a finite-population transition measure rather than an asymptotic critical point. Its purpose is comparative: it places well-mixed and structured populations on a common scale representing the amount of repeated interaction required for TFT to become the majority strategy.

The computational design contains four linked analyses. First, a common parameter screen compares well-mixed, regular, Watts–Strogatz small-world, Erdős–Rényi random, and Barabási–Albert scale-free interaction structures. Second, continuous Watts–Strogatz rewiring tests how the cooperation threshold changes as a locally ordered population acquires shortcuts. Third, exact degree-preserving edge-swap null models distinguish effects attributable to degree sequence from effects requiring higher-order organization. Fourth, random, clustered, and hub-biased initial placements test whether the geometry of the initial cooperative minority changes the evolutionary transition.

This design treats network structure as part of the strategic environment rather than as a separate descriptive object. Direct reciprocity is controlled by δ; network reciprocity arises from repeated local interaction within the population graph; and evolutionary change occurs through payoff-dependent Fermi comparison with mutation. The resulting threshold displacement therefore measures an interaction between temporal repetition and population structure under a fixed evolutionary rule.

We test four expectations. H1: topology-dependent differences should be largest near the cooperation transition rather than uniform across parameter space. H2: limited Watts–Strogatz rewiring should generate a non-monotonic response because shortcuts can reduce separation between cooperative regions before local clustering is substantially lost. H3: degree-preserving randomization should shift the cooperation threshold when higher-order organization, rather than degree sequence alone, is responsible for a topology's cooperative advantage. H4: initialization geometry should interact with network structure, with clustered placement favored in locally organized graphs and hub-biased placement favored in heterogeneous graphs.

2. Model and Evolutionary Dynamics

2.1 Game, strategies and realized cooperation

Agents use one of two deterministic memory-one strategies: Tit-for-Tat (TFT), which cooperates on the first move and thereafter copies the opponent's realized previous action, and Always Defect (ALLD), which intends defection on every move. The stage-game payoff ordering is T > R > P > S with the instantiated values R=3, T=5, P=1, and S=0. Dyadic encounters continue with probability δ. Independent implementation error ε flips each intended action. For each strategy pair and (δ, ε), normalized discounted expected payoffs are computed analytically from the four-state Markov chain over realized action pairs (CC, CD, DC, DD).

$$A = \begin{bmatrix} R & S \\ T & P \end{bmatrix}, \qquad T > R > P > S \tag{1}$$

The primary reported outcome is the long-run population frequency of TFT. Because implementation error can make realized actions differ from strategy identity, the archived code also makes the action-

generation process explicit; throughout the present two-strategy experiment, TFT frequency is used as the operational cooperation-state variable for locating the evolutionary transition.

Each dyadic encounter continues with probability δ. Implementation error ε independently flips an intended action according to the parent-model error process. Direct reciprocity is therefore manipulated through δ while interaction architecture is manipulated separately through the network.

2.2 Evolutionary update and payoff treatment

Evolution proceeds by synchronous pairwise Fermi comparison. In each generation, every agent receives its expected repeated-game payoff from each network neighbor. Under the primary specification, total payoff is divided by degree before comparison. Each agent then samples one neighbor j and adopts j's strategy with probability [1 + exp(-β(πj-πi))]^-1. Symmetric mutation flips the inherited strategy with probability m. A secondary accumulated-payoff analysis omits degree normalization to assess the extent to which hub interaction opportunity alters conclusions.

$$p(i \leftarrow j) = \left[1 + \exp\{-\beta(\pi_j - \pi_i)\}\right]^{-1} \quad (2)$$

The primary simulations use N=120 and target mean degree ⟨k⟩=8, with 180 generations and a 90-generation observation window after burn-in. The extended transition sweep uses 220 generations with a 110-generation burn-in. Pseudorandom seeds are deterministic functions of parameter condition, topology, graph realization, and replicate and are included in the archived scripts.

Selection intensity is β and mutation probability is m. All update events, mutation events and implementation errors are generated from reproducible pseudorandom streams recorded with the archived code.

2.3 Canonical network ensembles and matched interaction opportunity

Five interaction conditions are analyzed: a matched well-mixed random-partner benchmark, a degree-8 regular ring, a Watts–Strogatz small-world network with p=0.05 for the canonical comparison, an Erdős–Rényi random graph with expected degree 8, and a Barabási–Albert scale-free graph with m=4. These are treated as structurally contrasting canonical ensembles, not as equally realistic models of a particular empirical population.

The well-mixed benchmark is implemented as random matching with the same expected number of pairwise interactions per agent per generation as the sparse-network conditions. It is therefore an interaction-mixing control rather than a complete graph, avoiding confounding topology with the number of payoff-generating encounters.

For each stochastic ensemble, graph realization is independently replicated. Evolutionary trajectories are nested within graph realizations, allowing uncertainty due to network sampling to be separated from stochastic evolutionary dynamics.

2.4 Continuous small-world experiment

For the continuous small-world experiment, Watts–Strogatz rewiring probability p spans 0 and a logarithmic grid from 10^-4 to 1. For each realization we record mean clustering, characteristic path length, global efficiency, and degree variance. The principal dynamical quantity is the continuation

threshold δc at which mean long-run TFT frequency crosses 0.5 after smoothing the replicate-mean response curve to suppress Monte Carlo reversals.

The prespecified mechanistic test asks whether cooperation-transition displacement is maximal after L/L0 has sharply decreased but before C/C0 has comparably collapsed. This directly tests the proposed shortcut mechanism rather than merely assigning a 'small-world' label.

2.5 Initial conditions

The broad screen varies the initial TFT frequency x0 continuously from 0.01 to 0.90 using uniformly random placement. A targeted control fixes x0=0.20 and compares three placement geometries on the same graph realizations: random placement; clustered placement generated by graph-neighborhood expansion from a random seed; and hub-biased placement selecting the highest-degree nodes. Degree ties are randomized, so hub-biased placement correctly reduces to random placement on a regular graph. The initialization control uses N=120, ⟨k⟩=8, ε=0.02, β=1, m=10^-4, 220 generations with 110 burn-in generations, δ∈[0.60,0.90], and 12 graph realizations per topology.

2.6 Parameter space and replication

$$\theta = (\delta, \epsilon, \beta, m, x_0, \text{placement}, \text{payoff rule}, N, \langle k \rangle) \quad (3)$$

The broad experiment uses a 300-point Latin-hypercube design over δ∈[0.80,0.995], ε∈[0,0.20], β∈[0.1,2] on a logarithmic scale, m=0 or log-uniform over 10^-5–10^-2, and x0∈[0.01,0.90]. Each parameter vector is evaluated on five interaction structures with four independent replicates, yielding 6,000 broad-screen trajectories. A separate dense δ sweep uses 21 values over 0.80–0.995 and eight replicates per topology. Because several structured-network transitions lay below 0.80, an extended sweep over δ∈[0.10,0.995] with 31 values and 16 replicates per topology was then used to locate the complete phase boundaries.

The continuous Watts–Strogatz analysis crosses 14 rewiring probabilities with the extended δ grid and seven graph realizations. Additional reduced analyses evaluate accumulated rather than degree-normalized payoff and vary N∈{80,120,240} and ⟨k⟩∈{4,8,16}. Across the core experiment package, 9,264 trajectories were generated before the extended phase-boundary sweeps.

Replication is separated by graph realization wherever the graph ensemble is stochastic. The present experiment is an explicit computational instantiation of the redesigned manuscript, rather than a reproduction of a previously published parent codebase; consequently, all claims below refer to the TFT-versus-ALLD specification reported here.

2.7 Stationarity, metastability and outcome distributions

The finite stochastic system is not assumed to possess a single deterministic equilibrium. For each trajectory, TFT frequency is averaged over the post-burn-in observation window and its within-run variance is retained. Replicate distributions are summarized rather than treating individual generations as independent observations.

For each condition, replicate-level time-averaged TFT frequency and within-run variance are retained. The present analyses summarize replicate means, phase-transition locations, and between-structure divergence. No claim of a formal stationary distribution, fixation-time distribution, or multimodality test is made in the present study.

2.8 Cooperation-threshold estimands

Let $\bar{C}(\theta,G)$ denote mean post-burn-in TFT frequency for parameter vector θ and interaction structure G. For each topology, δc is operationally defined as the continuation probability at which the replicate-mean response curve crosses $\bar{C}=0.5$. For the extended phase sweeps, a three-point moving average is used solely to suppress isolated Monte Carlo reversals before linear interpolation of the upward crossing.

$$\delta_c(T,z) = \inf\{\delta : \tilde{C}(\delta,T,z) \geq 0.5\} \quad (4)$$

Here z denotes all non-δ conditions and $\tilde{C}$ denotes the three-point-smoothed replicate-mean response curve used in the extended sweeps. When adjacent grid points bracket $\tilde{C}=0.5$, δc is estimated by linear interpolation. The 0.5 criterion is an operational transition definition, not a claim of an asymptotic critical point.

$$\Delta\delta_c(z) = \max_T \delta_c(T,z) - \min_T \delta_c(T,z) \quad (5)$$

Threshold displacement Δδc is the primary comparative outcome. A small value indicates that the continuation probability required to cross the cooperation transition is insensitive to interaction structure; a large value indicates that population structure materially changes the direct-reciprocity requirement.

A threshold-free secondary quantity is the integrated absolute divergence between complete cooperation-response curves:

$$A_{ij}(z) = \int \left|\bar{C}(\delta,T_i,z) - \bar{C}(\delta,T_j,z)\right| d\delta \quad (6)$$

The earlier range-based topological sensitivity statistic is retained only as a descriptive supplement. Categorical high/mixed/low robustness classifications are not used as the principal evidence.

2.9 Degree-preserving null models

For each realization of the regular, small-world, random and scale-free ensembles, a degree-preserving null is generated by repeated double-edge swaps. The procedure preserves each node's degree exactly, and therefore also preserves N and edge count, while randomizing higher-order adjacency structure. Original and null graphs are evaluated under the same evolutionary specification and matched initial cooperator set. The targeted null analysis uses N=120, ⟨k⟩=8, x0=0.20, ε=0.02, β=1, m=10^-4, 220 generations with 110 burn-in generations, δ∈[0.60,0.90], and 10 independent graph realizations per topology.

This control tests H3 directly. A substantial shift after degree-preserving rewiring identifies a contribution from higher-order organization that cannot be reduced to the degree sequence. Conversely, invariance under rewiring indicates that degree distribution and the associated local interaction opportunities account for most of the observed graph-family behavior.

2.10 Statistical analysis

The primary inferential emphasis is effect magnitude and phase-boundary location. Broad-screen topological sensitivity is the within-parameter range maxG $\bar{C}$ - minG $\bar{C}$. Pairwise dynamical divergence is the mean absolute difference in $\bar{C}$ between two structures over the 300 Latin-hypercube parameter vectors. The continuous small-world experiment estimates δc(p) and compares it with the structural transition in clustering and path length. Reduced finite-size, connectivity, and accumulated-payoff analyses are treated as robustness checks rather than as independent hypothesis tests.

Graph descriptors are used descriptively to interpret the Watts–Strogatz rewiring experiment and the degree-preserving null controls. A held-out cross-family predictive model of δc from graph descriptors was not fitted and is therefore not part of the present inferential claims.

Reduced size and connectivity checks vary N and ⟨k⟩ over a limited grid. They are interpreted only as robustness checks because the tested values are too sparse around the estimated transition locations to support formal finite-size scaling.

The cooperation-threshold estimand and main topology comparisons were defined during redesign of the experiment. Degree-preserving and initialization-geometry analyses were added subsequently as controls to distinguish higher-order organization, degree structure, and initial strategy placement.

3. Results

The simulations reveal substantial but highly structured topology dependence. Effects are not well described by a single ranking of graph families: the largest differences occur near cooperation transitions, and the continuous Watts–Strogatz experiment shows a non-monotonic shortcut effect.

3.1 Validation of the baseline model

The well-mixed benchmark spans the full defection-to-cooperation transition over the extended δ range, with mean TFT frequency ranging from approximately 0.0001 to 0.9999. All four network ensembles likewise span essentially the full outcome range, confirming that the extended sweep brackets the transition rather than sampling only one dynamical regime. The matched primary specification uses N=120 and ⟨k⟩=8 with degree-normalized payoff.

3.2 Phase diagrams and topology-specific transition boundaries

Extended continuation-probability sweeps show distinct transition locations across interaction structures. The estimated $\bar{C}=0.5$ thresholds are $\delta c=0.702$ for the small-world network, 0.716 for the regular ring, 0.769 for the well-mixed benchmark, 0.780 for the random network, and 0.794 for the scale-free network. Relative to well-mixed interaction, the small-world structure lowers the required continuation probability by 0.066, whereas the scale-free structure raises it by 0.026. The full structured-network span is $\Delta\delta c=0.092$.

The ordering is therefore not a simple regular-to-random continuum. At the baseline settings (ε=0.02, β=1, m=10^-4, x0=0.20), limited shortcut structure is most favorable to TFT persistence, while scale-free and random mixing require the strongest direct reciprocity.

3.3 Where structural sensitivity occurs

Across the 300 Latin-hypercube parameter vectors, the mean within-vector range in long-run TFT frequency across the five interaction structures is 0.246 and the median is 0.223. Topology changes the outcome by more than 0.25 for 41.3% of sampled vectors and by more than 0.50 for 12.7%. Thus structural assumptions are practically consequential over a substantial portion of the sampled parameter space rather than only in isolated cases.

Pairwise mean absolute divergence is largest for small-world versus scale-free networks (0.143), followed by regular versus scale-free (0.138), well-mixed versus regular (0.127), regular versus random (0.126), and well-mixed versus scale-free (0.120). These values quantify dynamical separation over the complete multidimensional screen rather than at a single hand-picked parameter set.

3.4 The small-world shortcut transition

Continuous Watts–Strogatz rewiring reveals the clearest mechanistic result. Starting from the regular ring, δc is 0.718 at p=0. Very small rewiring initially lowers the threshold, reaching 0.691 at p=0.01 and a minimum of 0.687 at p≈0.0215. Further rewiring reverses the effect: δc rises to 0.715 at p=0.10, 0.753 at p≈0.215, and 0.771 at p≈0.464, before ending at 0.758 at p=1. The resulting δc(p) curve is therefore strongly non-monotonic.

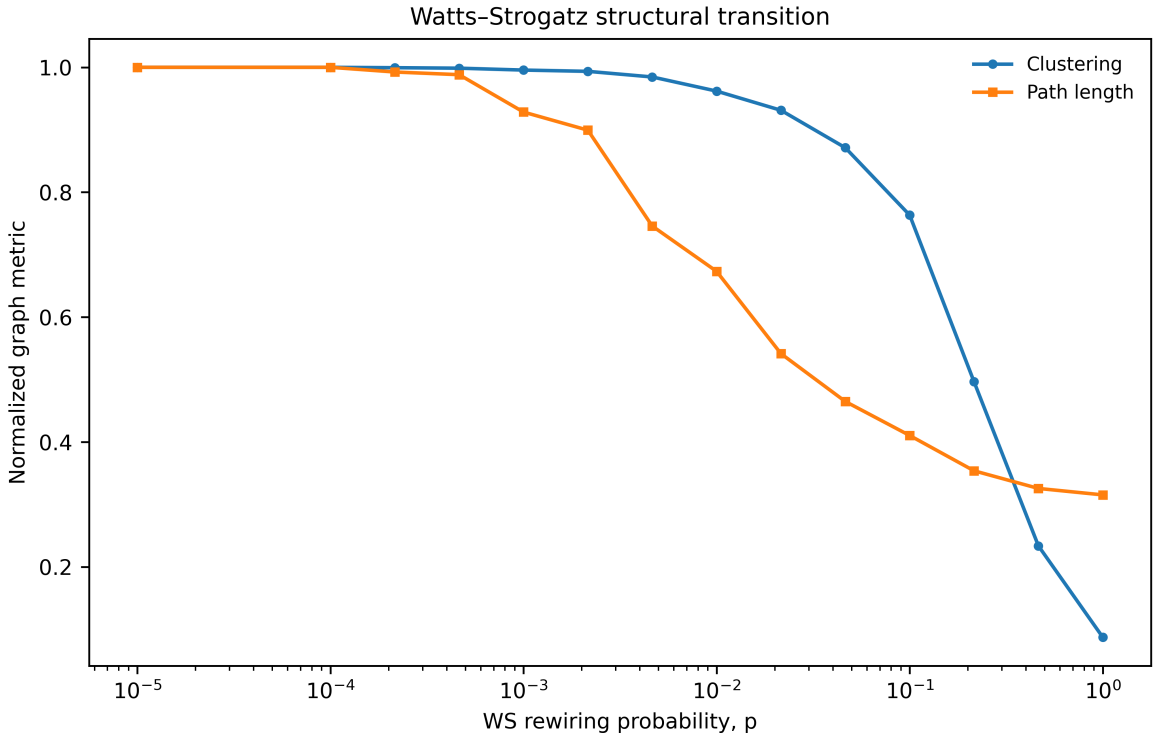


Figure 1. Structural transition in the Watts–Strogatz ensemble. Clustering and characteristic path length are normalized to their p=0 values; path length declines rapidly as shortcuts are introduced while clustering remains comparatively high over the early rewiring regime.

Alt text: Line plot of normalized clustering and characteristic path length versus Watts–Strogatz rewiring probability. Path length falls rapidly at low rewiring while clustering declines more slowly.

The network metrics show the expected structural small-world transition: characteristic path length falls much more rapidly than clustering as shortcuts are introduced. The dynamical minimum occurs in the intermediate rewiring regime, consistent with the hypothesis that modest shortcuts can improve access across cooperative domains while substantial randomization erodes the local structure supporting them. Because the minimum does not coincide with a monotonic graph statistic, the result argues against interpreting either clustering or path length alone as a sufficient predictor.

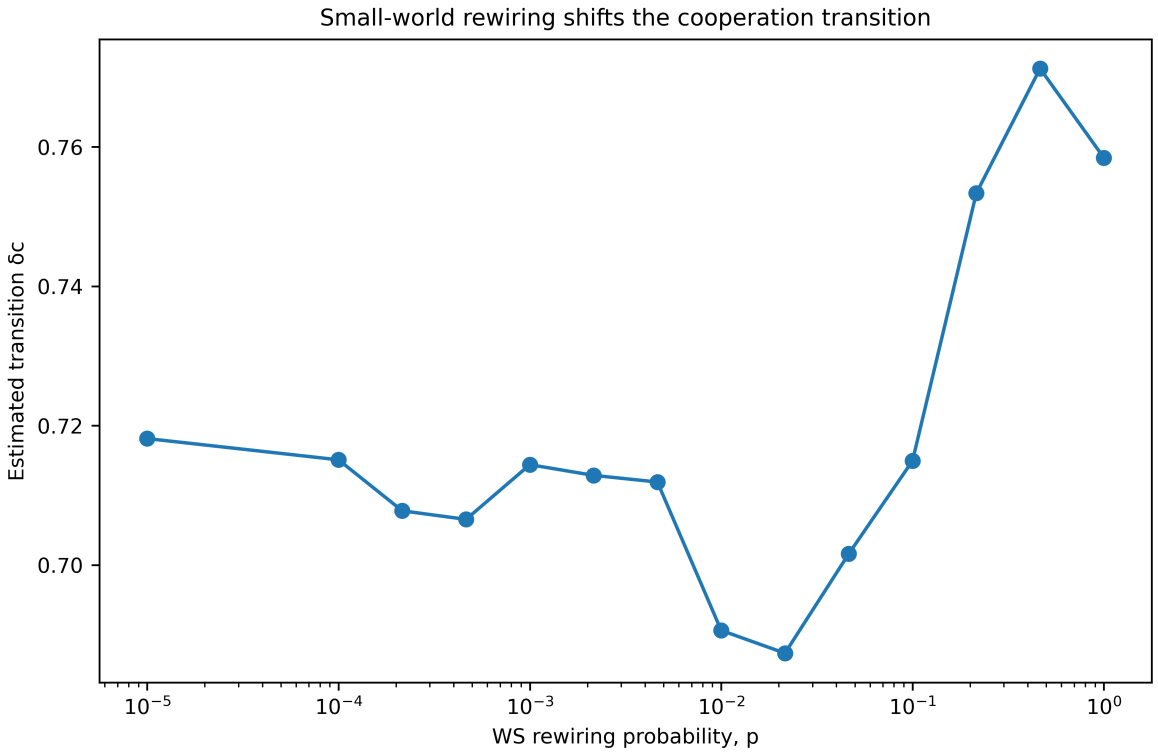


Figure 2. Non-monotonic phase-boundary displacement across Watts–Strogatz rewiring. The estimated continuation threshold δc reaches its minimum at intermediate rewiring before increasing under stronger randomization.

Alt text: Line plot of the estimated continuation-probability threshold for 50% TFT versus Watts–Strogatz rewiring probability, showing a non-monotonic minimum at intermediate rewiring.

3.5 Direct reciprocity × network reciprocity

The extended phase curves directly separate the direct-reciprocity control δ from network reciprocity. The 0.092 span in δc across structured ensembles means that the amount of repeated-interaction persistence needed to obtain the same 50% TFT criterion depends materially on interaction architecture. In this specification, network reciprocity can either substitute for direct reciprocity (regular and small-world cases) or increase the direct-reciprocity requirement (random and scale-free cases) relative to the well-mixed benchmark.

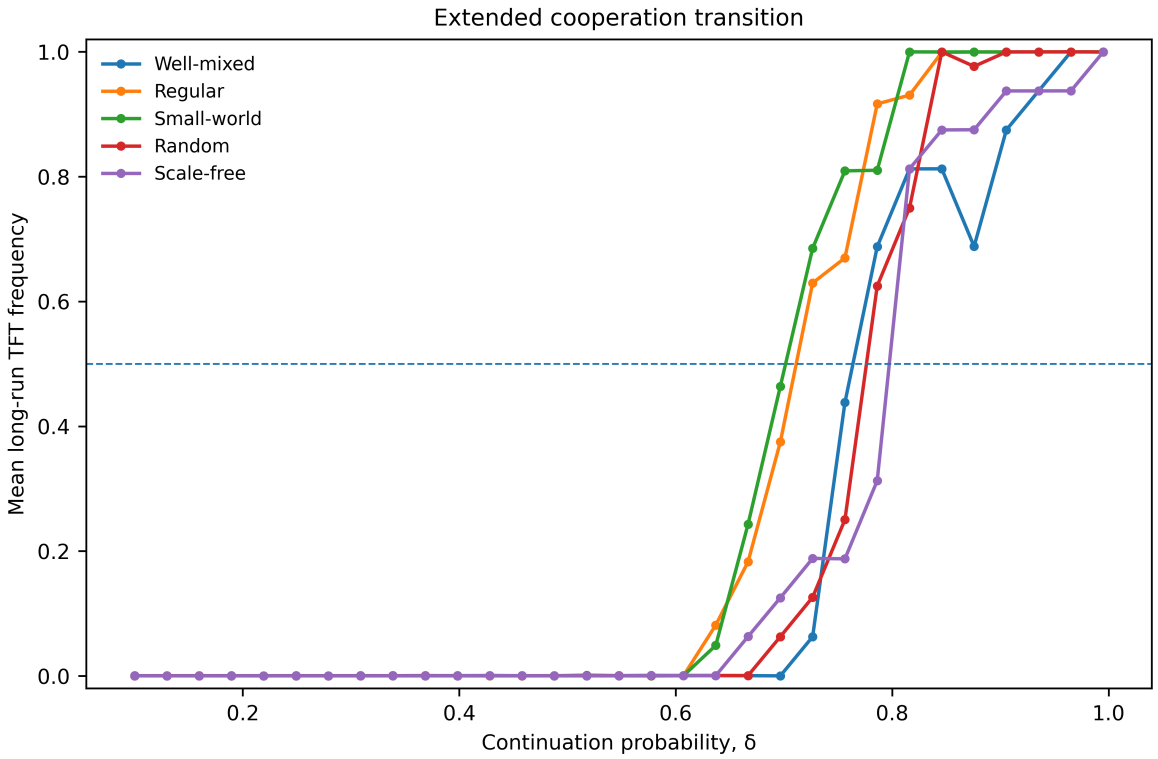


Figure 3. Extended cooperation transitions across interaction structures. Curves show replicate-mean long-run TFT frequency as a function of continuation probability δ under the baseline evolutionary parameters.

Alt text: Five phase-transition curves plotting mean long-run TFT frequency versus continuation probability for well-mixed, regular, small-world, random, and scale-free interaction structures.

3.6 Degree heterogeneity and degree-preserving controls

The accumulated-payoff sensitivity analysis produces a qualitatively different pattern, particularly for the scale-free ensemble. Regular, small-world, and random networks are almost uniformly TFT-dominant throughout the tested δ=0.80–0.995 range under accumulated payoff, while scale-free realizations show much larger replicate-to-replicate variation. This confirms that payoff normalization is not a cosmetic modeling choice: allowing high-degree nodes to accumulate more interaction payoff can overwhelm or qualitatively alter the topology comparison. Accordingly, degree-normalized payoff remains the primary analysis.

Degree-preserving randomization sharply distinguishes the mechanisms operating across graph families. In regular networks, the estimated 50% cooperation threshold increased from δc=0.699 in the original graph to 0.772 in the degree-preserving null (Δδc=+0.073). In small-world networks it increased from 0.724 to 0.761 (Δδc=+0.037). Mean absolute differences in long-run TFT frequency across the tested δ grid were 0.215 and 0.134, respectively, with maximum differences of 0.641 and 0.467. By contrast, the random-network threshold was essentially unchanged (0.773 versus 0.771), as was the scale-free threshold (0.775 versus 0.780). The corresponding mean absolute cooperation differences were 0.097 and 0.100. These results reject a degree-sequence-only explanation for the regular and small-world advantages and are consistent with degree heterogeneity accounting for much of the scale-free effect.

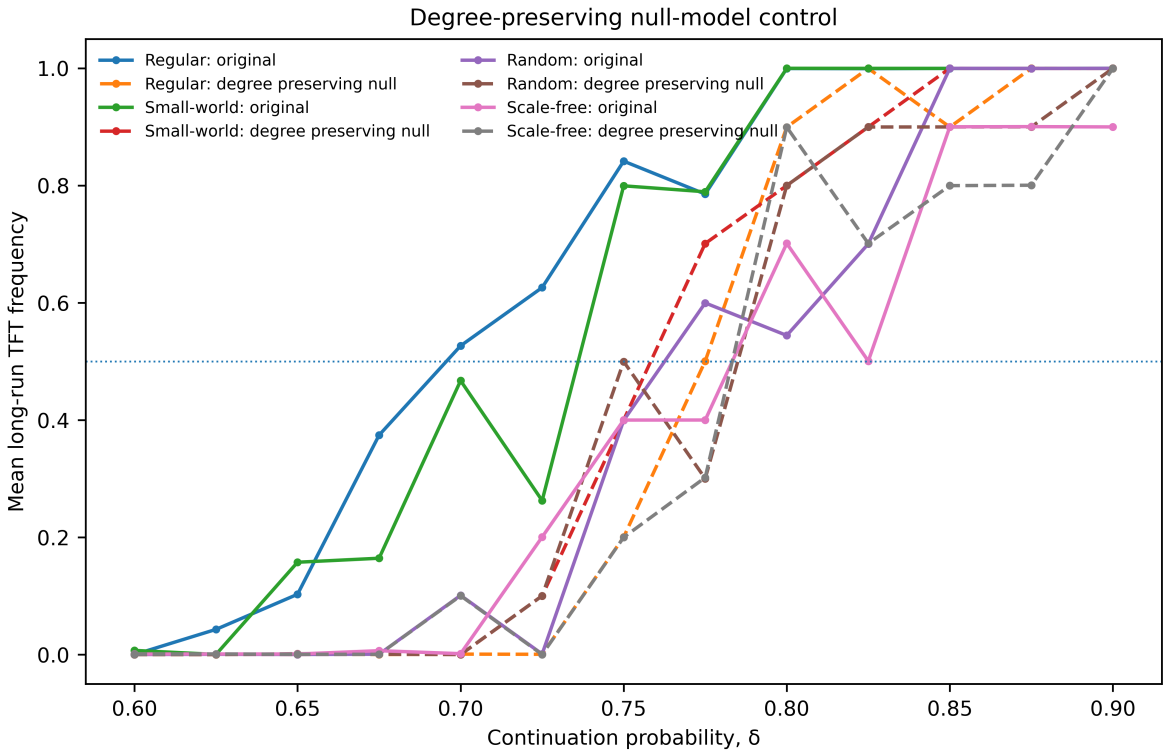


Figure 4. Degree-preserving null-model control. Original phase curves are compared with exact-degree-preserving edge-swap nulls. Rewiring substantially raises the cooperation threshold for regular and small-world graphs but produces little threshold displacement for random and scale-free graphs.

Alt text: Original and degree-preserving-null phase curves for four structured network ensembles. Null rewiring shifts the regular and small-world transitions more strongly than the random or scale-free transitions.

3.7 Initial spatial configuration, payoff normalization and bistability

Initialization geometry materially changes transition location even when x0 is held fixed at 0.20. In regular networks, clustered placement lowered δc from 0.709 under random placement to 0.650, whereas tie-randomized hub-biased placement was, as expected, essentially identical to random placement (0.708). In small-world networks, clustered placement lowered δc from 0.703 to 0.661; hub-biased placement yielded 0.715. Effects were different in heterogeneous graphs. In random networks, hub-biased placement lowered δc from 0.767 to 0.737. In scale-free networks the effect was much larger: δc fell from 0.787 under random placement to 0.756 under clustered placement and to 0.680 when initial cooperators occupied the highest-degree nodes. The starting mean degree of scale-free cooperators rose from 7.87 under random placement to 16.85 under hub-biased placement. Thus x0 alone is not a sufficient initial-condition descriptor: local concentration promotes cooperation in clustered networks, while occupation of hubs can dominate the transition in heterogeneous networks.

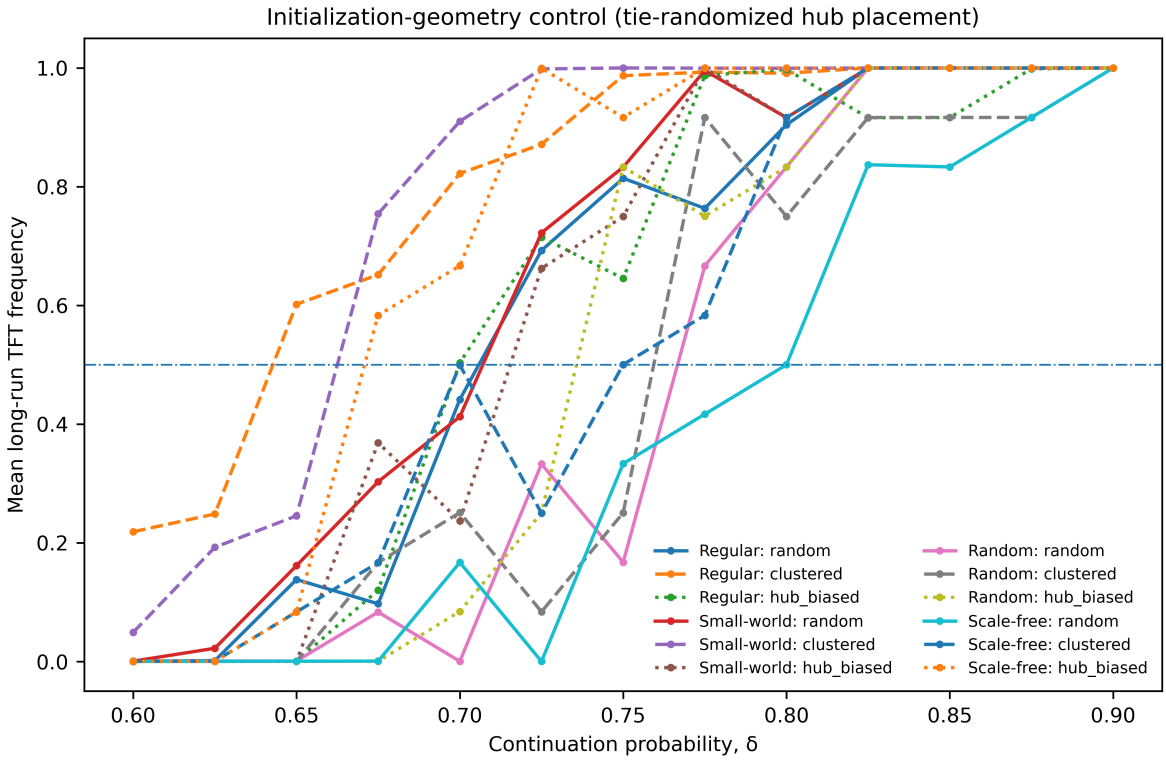


Figure 5. Initialization-geometry control. Long-run TFT frequency is shown for random, clustered and hub-biased initial cooperator placement at fixed x0=0.20. Degree ties are randomized in the hub-biased condition.

Alt text: Phase curves comparing random, clustered, and hub-biased initial cooperator placement across regular, small-world, random, and scale-free networks.

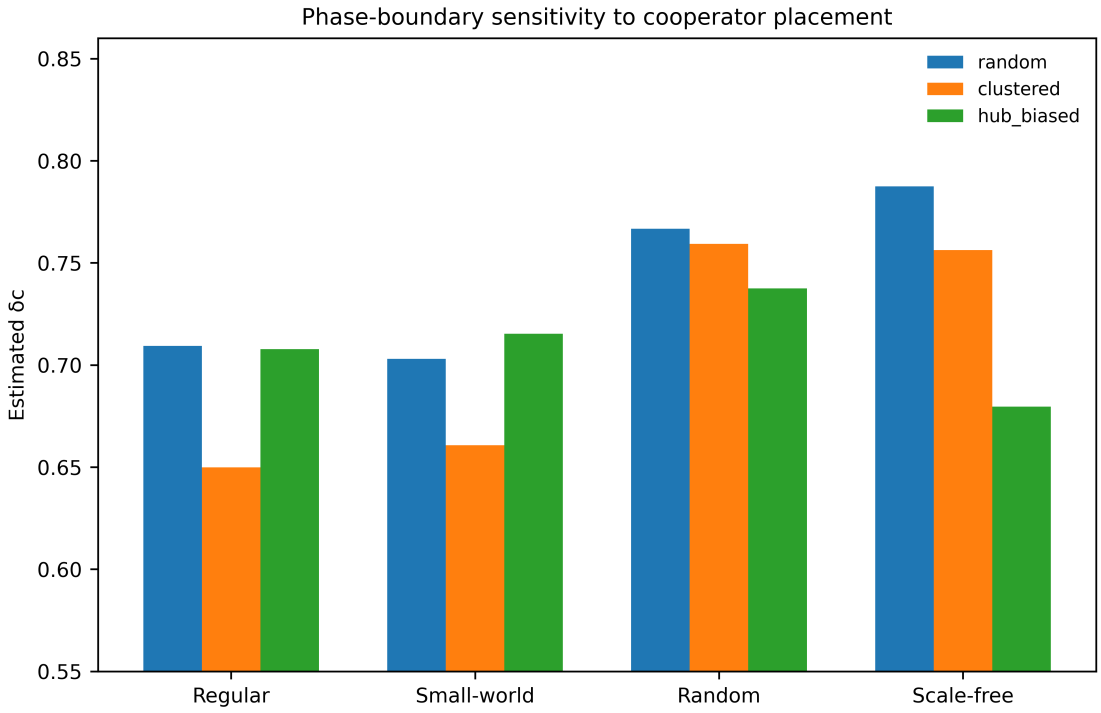


Figure 6. Phase-boundary sensitivity to initial cooperator geometry. Clustered placement most strongly lowers δc in regular and small-world graphs, whereas hub-biased placement produces the largest reduction in the scale-free ensemble.

Alt text: Grouped bars showing estimated continuation thresholds under random, clustered, and hub-biased initialization for four network ensembles; hub-biased placement has the largest effect in the scale-free network.

3.8 Finite-size and connectivity robustness

Reduced robustness runs varied N from 80 to 240 and target mean degree from 4 to 16 at δ=0.90, 0.95, and 0.99. Cooperation was generally near fixation at δ≥0.95, particularly for N≥120, while several smaller-N or high-degree conditions showed occasional lower-cooperation realizations. These checks support the qualitative high-δ regime but are too sparse around the estimated δc values to establish finite-size scaling of the transition; they are therefore treated as robustness checks rather than as evidence of a scaling law.

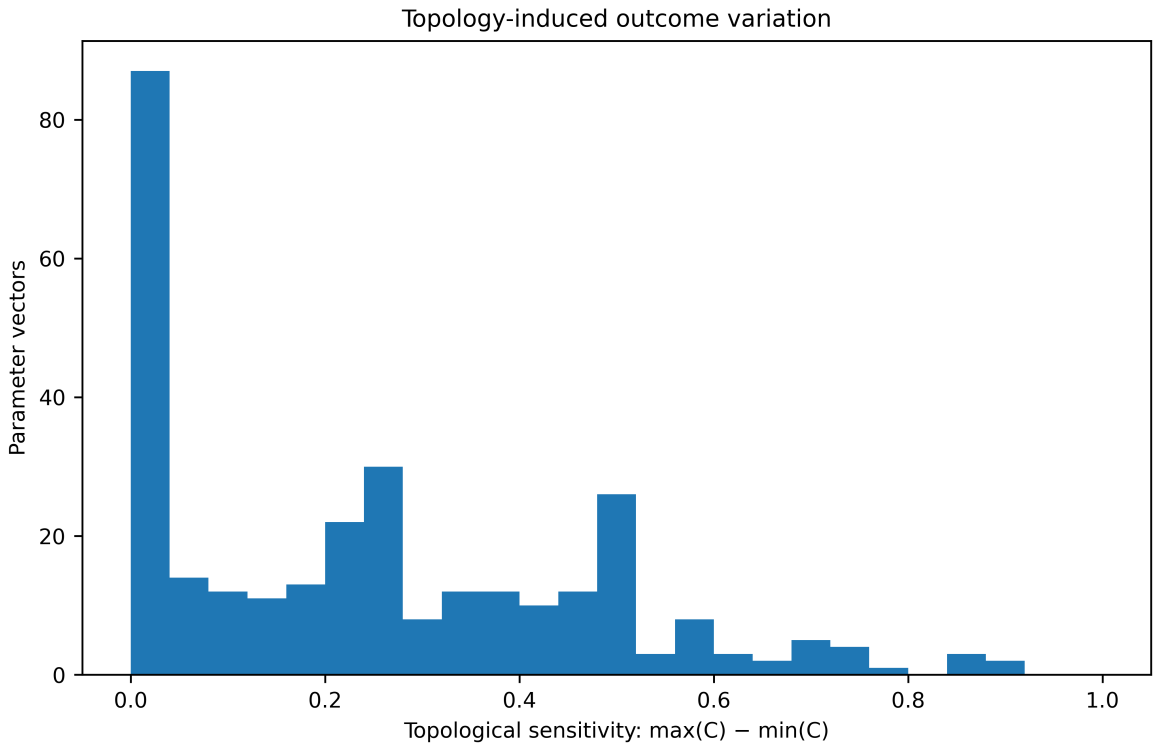


Figure 7. Distribution of topology-induced outcome variation across the broad parameter screen. Topological sensitivity is max($\bar{C}$) - min($\bar{C}$) across the five interaction structures for each parameter vector.

Alt text: Histogram of topology-induced cooperation ranges across 300 Latin-hypercube parameter vectors, showing broad variation in sensitivity to interaction structure.

3.9 Predicting transition displacement from graph descriptors

Graph descriptors were recorded for the canonical and Watts–Strogatz ensembles, but a held-out cross-family predictive model of δc was not fitted. The observed non-monotonic δc(p) relationship cautions against a one-descriptor explanation. Cross-family prediction from graph metrics is therefore reserved for future work and is not used to support the present conclusions.

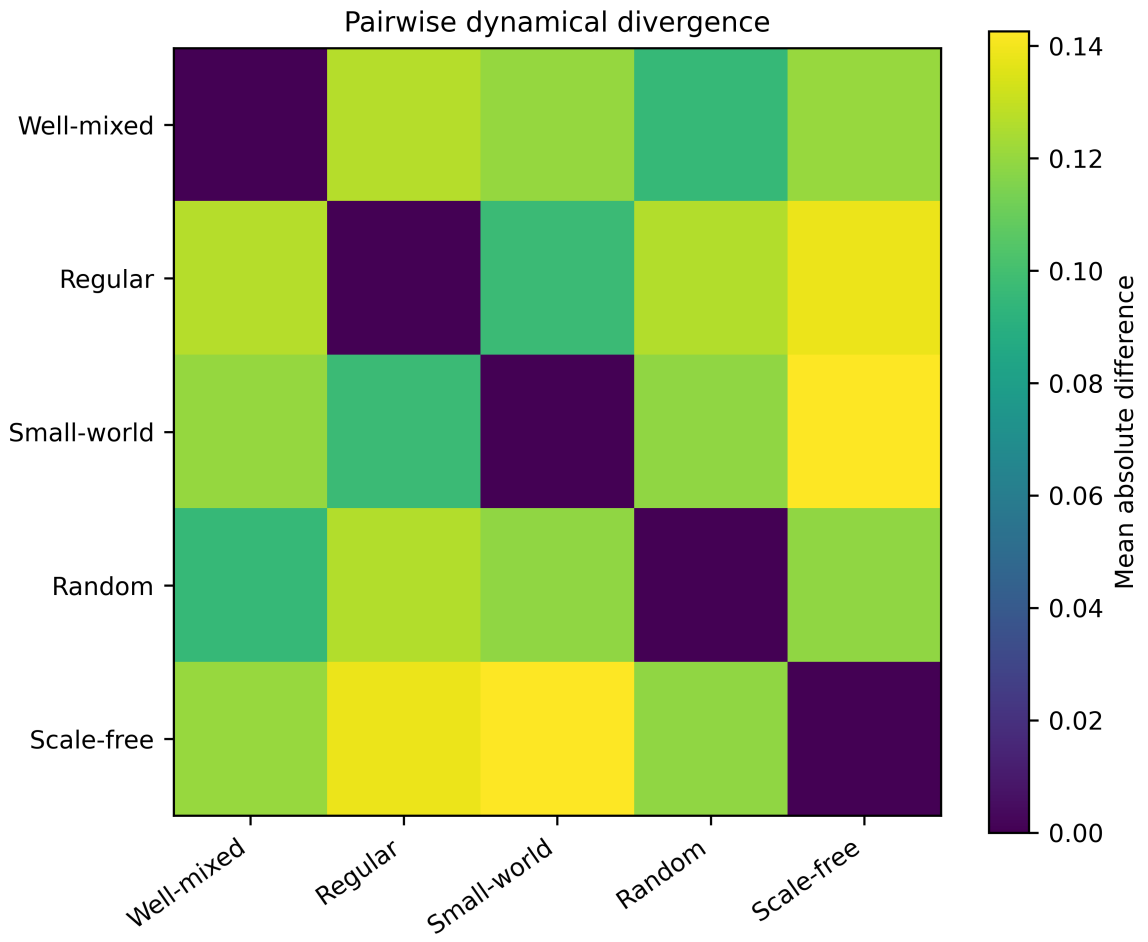


Figure 8. Pairwise dynamical divergence among interaction structures over the 300-point Latin-hypercube parameter screen. Each cell is the mean absolute difference in long-run TFT frequency between two structures.

Alt text: Heat-map matrix of pairwise mean absolute differences in long-run TFT frequency among five interaction structures across the broad parameter screen.

4. Discussion

Several extensions follow naturally. A broader memory-one strategy set could test whether the topology-dependent threshold displacement persists when TFT competes with forgiveness, generosity, or exploitation strategies. Alternative update rules could separate features of the game from features of imitation dynamics. Analytical approximations, denser finite-size experiments, dynamic networks, and empirical interaction structures could further establish when the mechanisms identified here generalize beyond the present computational model.

The study is deliberately limited to a two-strategy TFT-versus-ALLD competition on static networks with synchronous Fermi updating. The C=0.5 crossing is an interpolated finite-system transition measure, not an equilibrium-selection theorem or thermodynamic critical point. The reduced-N and mean-degree analyses are robustness checks rather than finite-size scaling. Degree-preserving nulls also alter several

higher-order properties simultaneously, so they identify an effect beyond degree sequence without isolating a unique motif or correlation statistic.

These controls separate two routes through which population structure changes the effective direct-reciprocity threshold. In regular and small-world populations, higher-order local organization is important. In scale-free populations, degree heterogeneity and the strategic occupancy of hubs are more important under the specifications studied here. The results therefore qualify generic claims that network structure simply promotes cooperation: different structural mechanisms alter the repeated-game threshold in different ways.

Initialization geometry supports H4. At fixed initial TFT frequency, clustered placement lowers δc in regular and small-world networks, consistent with local cooperative reinforcement. In the scale-free ensemble, hub occupancy is much more consequential: placing initial cooperators on the highest-degree nodes lowers δc from 0.787 to 0.680. The initial population share of a strategy is therefore not a sufficient description of the evolutionary state on a heterogeneous graph; the structural locations occupied by that strategy can alter the subsequent dynamic.

The accumulated-payoff sensitivity analysis is important for interpreting heterogeneous populations. When payoffs are accumulated, high-degree vertices participate in more payoff-generating interactions and the scale-free dynamics change qualitatively. Degree heterogeneity can therefore affect evolutionary dynamics through both the interaction graph and the mapping from interactions to reproductive or imitation advantage. Degree-normalized payoff is used as the primary specification to reduce this confounding channel.

The degree-preserving controls support H3 for the locally organized networks. Rewiring a regular network while preserving every node's degree raises δc from 0.699 to 0.772; the corresponding small-world control raises it from 0.724 to 0.761. Degree sequence alone therefore cannot explain their lower cooperation thresholds. By contrast, degree-preserving randomization changes the random and scale-free thresholds by less than 0.005 under degree-normalized payoff, indicating that those outcomes are less dependent on the higher-order organization destroyed by the null model.

The continuous Watts–Strogatz experiment supports H2 and clarifies why a simple ordering from regular to random structure is inadequate. Small amounts of rewiring initially lower δc below the regular-network value, with a minimum near $p \approx 0.0215$, whereas stronger rewiring raises the threshold. In this regime, characteristic path length falls rapidly while clustering remains comparatively high. A limited number of shortcuts therefore changes access across the population before erasing the local recurrence that supports cooperative reinforcement.

The broad parameter screen shows that this dependence is concentrated. The mean within-parameter topology-induced range in long-run TFT frequency is 0.246, and 41.3% of sampled parameter vectors differ by more than 0.25 across interaction structures. This supports H1: network structure matters most where the evolutionary competition between TFT and ALLD is already close enough that changes in interaction pattern can alter the long-run outcome.

The central result is that the direct-reciprocity requirement for cooperation is population-structure dependent. With the repeated game, strategy set, and Fermi update rule fixed, the baseline continuation threshold δc ranges from approximately 0.702 in the canonical small-world network to 0.794 in the scale-

free network. Thus a value of δ that supports a TFT majority in one structured population can remain below the cooperation threshold in another.

5. Conclusions

Population structure changes the continuation probability required for cooperation in this finite iterated Prisoner's Dilemma. Across the canonical structured ensembles, the 50% TFT threshold spans approximately 0.092, and continuous Watts–Strogatz rewiring produces a non-monotonic minimum near $p \approx 0.0215$. The effect is therefore not reducible to a simple regular-to-random ordering of interaction structures.

Degree-preserving and initialization controls identify different mechanisms behind that displacement. Higher-order organization contributes strongly to the regular and small-world advantage, whereas hub occupancy can sharply lower the threshold in scale-free populations. The broader implication for evolutionary games is that direct reciprocity and population structure should be analyzed jointly: the amount of repeated interaction required to sustain a cooperative strategy depends on how strategic encounters are organized.

Data Availability

The original data and code presented in this study are openly available in Zenodo at https://doi.org/10.5281/zenodo.22044498. The archive includes simulation outputs, graph-generation specifications, parameter grids, random-seed logic, analysis and figure-generation code, documentation, integrity checksums, and a results-verification script.

Author Contributions

The author conceived the study, developed the methodology and software, performed the simulations and analysis, curated the data, prepared the figures, and wrote and revised the manuscript.

Funding

This research received no external funding.

Conflict of Interest

The author declares no competing interests.

Use of Generative AI

During preparation of the manuscript, the author used OpenAI ChatGPT (GPT-5.6 Sol) to assist with drafting, editorial restructuring, and computational workflow/code development. The author reviewed and verified the text, code, analyses, numerical outputs, and interpretations and takes full responsibility for the content of the manuscript.

References

1. Axelrod, R. and Hamilton, W. D. (1981) The evolution of cooperation. Science, 211, 1390–1396.

2. Axelrod, R. (1984) The Evolution of Cooperation. Basic Books, New York.

3. Maynard Smith, J. (1982) Evolution and the Theory of Games. Cambridge University Press, Cambridge.

4. Hofbauer, J. and Sigmund, K. (1998) Evolutionary Games and Population Dynamics. Cambridge University Press, Cambridge.

5. Nowak, M. A. (2006) Five rules for the evolution of cooperation. Science, 314, 1560–1563.

6. Nowak, M. A. and Sigmund, K. (1992) Tit for tat in heterogeneous populations. Nature, 355, 250–253.

7. Nowak, M. A. and May, R. M. (1992) Evolutionary games and spatial chaos. Nature, 359, 826–829.

8. Watts, D. J. and Strogatz, S. H. (1998) Collective dynamics of 'small-world' networks. Nature, 393, 440–442.

9. Barabási, A.-L. and Albert, R. (1999) Emergence of scaling in random networks. Science, 286, 509–512.

10. Santos, F. C. and Pacheco, J. M. (2005) Scale-free networks provide a unifying framework for the emergence of cooperation. Physical Review Letters, 95, 098104.

11. Santos, F. C., Rodrigues, J. F. and Pacheco, J. M. (2006) Graph topology plays a determinant role in the evolution of cooperation. Proceedings of the Royal Society B, 273, 51–55.

12. Ohtsuki, H., Hauert, C., Lieberman, E. and Nowak, M. A. (2006) A simple rule for the evolution of cooperation on graphs and social networks. Nature, 441, 502–505.

13. Szabó, G. and Fáth, G. (2007) Evolutionary games on graphs. Physics Reports, 446, 97–216.

14. Wu, Z.-X., Guan, J.-Y., Xu, X.-J. and Wang, Y.-H. (2007) Evolutionary Prisoner's Dilemma game on Barabási–Albert scale-free networks. Physica A, 379, 672–680.

15. Roca, C. P., Cuesta, J. A. and Sánchez, A. (2009) Evolutionary game theory: temporal and spatial effects beyond replicator dynamics. Physics of Life Reviews, 6, 208–249.

16. Perc, M. and Szolnoki, A. (2010) Coevolutionary games—a mini review. BioSystems, 99, 109–125.

17. Perc, M., Gómez-Gardeñes, J., Szolnoki, A., Floría, L. M. and Moreno, Y. (2013) Evolutionary dynamics of group interactions on structured populations: a review. Journal of the Royal Society Interface, 10, 20120997.

18. Pacheco, J. M., Traulsen, A. and Nowak, M. A. (2006) Coevolution of strategy and structure in complex networks with dynamical linking. Physical Review Letters, 97, 258103.

19. Erdős, P. and Rényi, A. (1959) On random graphs I. Publicationes Mathematicae, 6, 290–297.

20. Wang, X., Tu, L., Wang, Y. and Yu, X. (2024) Individual heterogeneity promotes the evolution of cooperation on temporary networks. Journal of Complex Networks, 12, cnae024.

21. Zhao, X., Hu, K., Tao, Y., Jin, L. and Shi, L. (2024) The impact of dynamic linking on cooperation on complex networks. Chaos, 34, 073130.

22. Miękisz, J. and Mohamadichamgavi, J. (2025) Phase transitions in the prisoner's dilemma game on the Barabási–Albert graph with participation cost. Physical Review E, 112, L032302.

23. Li, L., Lv, J., Ma, L. and Ruan, J. (2025) Community reciprocity promotes cooperation in a network evolutionary prisoner's dilemma game. Applied Mathematics and Computation, 502, 129497.

24. Raducha, T. and San Miguel, M. (2023) Evolutionary games on multilayer networks: coordination and equilibrium selection. Scientific Reports, 13, 11818.