Simulating the Emergence of Cooperation through Spatially Constrained Iterative Games

By Eric Rothman, Fabricio Rua-Sanchez, Daniel Zhan

Applied Game Theory - Stage Final Paper

Introduction

Cooperative behavior poses a central paradox in evolutionary biology; if natural selection favors traits that maximize individual fitness, how can cooperative strategies that are often costly to individuals persist across ecological systems? Evolutionary game theory (EGT) provides a mathematical framework to study this phenomenon by modeling strategic interactions among self-interested agents. While classical EGT offers foundational insights, many models oversimplify the dynamic and spatially structured interactions that characterize real-world systems.

This project uses computational simulations to investigate how cooperation emerges and stabilizes under diverse evolutionary pressures. By simulating spatially constrained, iterated games, we analyze how strategic behaviors evolve under varying interaction rules, learning dynamics, and environmental constraints. Our modular simulation framework will enable systematic experimentation with game parameters, initial strategy distributions, and adaptation mechanisms. This approach bridges gaps between theoretical predictions and empirical observations, offering insights into the conditions that promote cooperative equilibria in spatial ecological systems.

Domain/Problem

EGT examines how strategies propagate through ecological populations based on their relative success in competitive interactions. Key concepts include the following:

- Evolutionary Stable Strategies (ESS) are strategies resistant to invasion by alternative strategies under selection pressures.
- Learning dynamics describe how agents update their strategies over time in response to feedback. In particular, Replicator Dynamics describe a mathematical model pertaining to how the frequency of strategies changes proportionally to their fitness.
- Network Reciprocity describes spatial or social structures that enable cooperative clusters to form and resist exploitation by defectors.

Classic games like the Prisoner’s Dilemma and Hawk-Dove illustrate tensions between individual and collective incentives. For example, in the Prisoner’s Dilemma, mutual cooperation yields higher collective payoffs than mutual defection, but defection is individually rational in one-off games. In contrast, the Hawk-Dove game has a different number of Nash equilibria depending on the cost of Hawk fighting, highlighting the need for strategies that balance aggression and cooperation. While these games provide theoretical benchmarks, their real-world applicability is limited by assumptions of perfect rationality and strictly defined costs for fighting and cooperating.

A major challenge in modeling cooperative evolution is balancing model complexity and interpretability. Overly simplistic models fail to capture spatial constraints, adaptive networks, or stochastic learning mechanisms. Conversely, overly complex models obscure causal relationships between variables. Other complicating factors further complicate outcomes and interpretability:

- Parameter Sensitivity: Outcomes in EGT often hinge on hyperparameters (e.g., payoff ratios, interaction rules, learning mechanisms). Small changes can shift equilibria, complicating generalization.
- Empirical Validation: Biological systems involve confounding variables (e.g., kin selection, environmental noise) that are challenging to replicate. Even for parameters that are included in our simulation, it can be difficult to identify what parameter values are used in empirical studies (e.g., what is the true average cost of two hawks fighting? Is using an average too reductive?). This challenge exacerbates the aforementioned parameter sensitivity of such simulations.
- Dynamic Adaptation: Real-world agents learn from past interactions, but many models assume fixed strategies or idealized update rules.

Prior work highlights these limitations. For instance, Axelrod’s Tit-for-Tat strategy succeeds in idealized iterated Prisoner’s Dilemma scenarios but falters in stochastic environments. Our work aims to address some of the limitations present in existing literature by offering comprehensive hyperparameter selection, allowing the user to experiment with a huge number of combinations of games, learning dynamics, payoff values, and interaction radii.

Design of Program

Our Python-based framework simulates evolutionary dynamics across customizable games, interaction topologies, and learning mechanisms. The design prioritizes modularity to support:

- Multiple Game Types: Predefined payoff matrices for Prisoner's Dilemma, Hawk-Dove, Stag Hunt, Rock-Paper-Scissors, and Battle of the Sexes, with support for custom games.
- Spatial Constraints: Agents interact within toroidal grids, balancing edge effects while enabling localized strategy diffusion.
- Adaptive Learning: Agents update strategies via replicator dynamics, Fermi rules, and other probabilistic update mechanisms.

Additionally, our simulation platform supports the following features to allow customized simulation runs and enhance interpretability and speed of research discovery:

- Visualization: A step-by-step grid update with color coded strategies displays how the strategy distribution evolves over time.
- Data output: Metrics and strategy distributions are saved automatically, and visualization software enables visualization of past runs.
- User-defined initial strategy distribution: An important hyperparameter that can be altered to see its effects on outcomes.
- User-defined opponent selection: Another important hyperparameter that can be seen as an extension of individual strategy.
- Object-oriented Agent class design: Encapsulates strategy, strategy history, score, interaction history, and position for each agent.

The simulation runs the following process until stopping criteria are satisfied:

1. Initialization: Agents are assigned strategies and positions.
2. Begin simulation loop:
    1. Interaction Phase: Agents play games with neighbors, accumulating payoffs.
    2. Strategy Update: Agents adopt new strategies based on learning dynamics.
    3. Visualization is updated.
3. Data is saved.

Expected Results from Literature

Through research into the different games and their natural occurrences we have discovered previous results that we will attempt to replicate with our simulation. Each of the games has a unique role as either a proof of concept or a game with relevance to processes found in biological ecosystems. Since the simulation is spatially constrained, we expect it is unlikely that a strategy included at the start is fully eliminated, but we can still measure dominance or ratios of success by quantifying the frequency of strategies after populations stabilize.

The first game that we examined is the Prisoner's Dilemma. In 1981, Alexrod famously ran a competition of different strategies for the Prisoner's Dilemma and concluded that Tit-for-Tat was the best strategy of all strategies entered. Thus we expect that if that strategy is included it will emerge as the dominant strategy. That means at the end of the simulation, a plurality of agents will have the Tit-for-Tat strategy and the others will be a mix of the remaining strategies. Although this game does not have clear implications in actual ecosystems, scientists have simulated the game with animals to see how they act in captivity. In defiance of the simulation's findings, Tit-for-Tat was never the dominant strategy, and pure betrayal was the most likely with the base payouts in bluebirds and rats. We thus conclude that, under specific conditions, betrayal may be a dominant strategy instead.

The Stag Hunt game highlights the tension between risk and reward. While this game is normally described in terms of human behavior, this game has clear implications in ecosystems beyond human activity. The best example of this is within Orcas' feeding habits; orcas will hunt herrings in a similar mapping from the Stag Hunt problem. They will work together to secure a school of herrings instead of breaking off and chasing a lesser amount. If only simple strategies are used, we are expecting hare hunting (individualistic behavior) to be the dominant strategy, but there could still be sizable clusters of stag hunting (cooperative behavior).

The Hawk-Dove game is a classic EGT game used to model competition over resources. The idea for this game originates from examining competitive behavior between different members of a species for a resource vs the cost of injury. This game occurs across a wide swath of animals and they each handle it a different way based on the likeliness of injury. Because of the similarity in our simulation, we are expecting a similar result to what was found by Fu, Nowak, and Hauert in "Invasion and expansion of cooperators in lattice populations: Prisoner's dilemma vs. snowdrift games." When the payoff difference between betraying and cooperating is small, cooperation becomes the dominant strategy, and betrayal becomes more likely as the difference grows. Except for huge payoff differences, both strategies will always be represented and there will always be clusters of cooperation even though they may be smaller.

Dawkin's Battle of the Sexes highlights how EGT can produce stable strategies that that not Nash Equilibriums and has a clear expected outcome. The game was created to show how female coyness might be a viable strategy for reproduction, as seen in animals like pigeons. As such we are expecting the Coy (Uncooperative) strategy to be the dominant strategy among female agents. Since we are representing reproduction, agents will be divided between male and female agents and they will have different strategies. The strategy that is expected to be dominant amongst males is the helpful (Cooperative) strategy but only after the unhelpful strategy took over most of the grid. This is expected since the male being unhelpful is only beneficial when the female is non-coy and we are expecting most females to be coy.

The final game we are experimenting with and the only three option game is Rock-Paper-Scissors. There are places in nature where this game occurs; one of which is within strains of E. Coli. This occurs since each of the strains grows at a different rate allowing one strain to out compete the others, but coincidentally that strain is killed by a toxin produced by the slowest strain. This forms a Rock-Paper-Scissors where each strain outcompetes a strain and loses to the other one. In "Local dispersal promotes biodiversity in a real-life game of rock–paper–scissors" Kerr and Co. looked to see what happens when these strains compete in different environments. When each strain only interacts locally, the end result was the stable Nash Equilibrium of about an equal amount of each strain. Thus, we are expecting that the final result will be that a third of the grid is dedicated to each type strategy in Rock-Paper-Scissors and a stability is formed.

Results from Simulation Study

Our simulated runs of the Prisoner's Dilemma show promising, repeatable results that not only replicate existing research results in literature but also serve as a strong proof of concept for the robustness and reliability of the underlying code and architecture of our simulation platform.

Our most important finding with Prisoner's Dilemma is that <u>spatial reciprocity dominates</u>. Using nearest-neighbor play (radius=1), a matrix size of 50x50 with toroidal edge effects, and the standard payoff matrix (3/3, 0/5, 5/0/ 1/1), replicator dynamics caused defection to appear strongly, stabilizing at around 60% of the grid on average, as seen in Figure 1. However, locally constrained clusters of cooperation and Tit-for-Tat strategies appear that persist indefinitely, seemingly immune to takeover from defecting agents that border these clusters. This pattern, first reported by Nowak, was reproduced in every trial using this combination of parameters and thus constitutes a reliably repeatable outcome that signifies the potential emergence of cooperation under clusters of adjacent agents that have developed “trust” in one another, even when surrounded by more cynical agents. We also found that increasing the interaction range gradually erodes cluster protection; at radius values of 2 and 3, if an initial cooperative pocket survives the first 20 iterations it is likely to persist, but defection often dominates the entire board. At a radius value of 5 every run ends up with total defection domination. We attribute this finding to the concept of the “shadow of the future”; when agents expect to interact with each other again, cooperation is encouraged, but as the radius of interaction is increased, it becomes increasingly unlikely that repeated interaction occurs, causing defection to become the rational strategy.

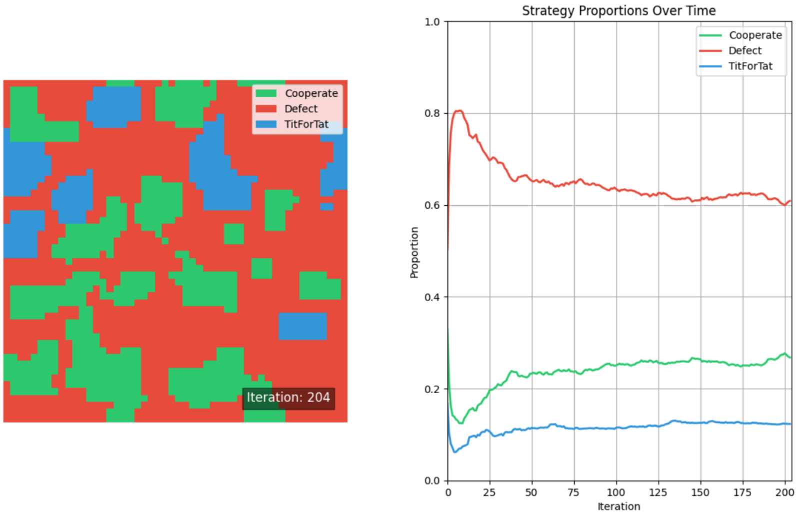

Figure 1: Prisoner’s Dilemma game played with default parameters (radius=1, replicator dynamics, default matrix).

Furthermore, our research involving the Stag Hunt game reveals that payoff sensitivity maps the basin of attraction. Using default hyperparameters (e.g., interaction radius of 1, payoff matrix values of 5/5, 0/3, 3/0, 3/3, replicator dynamics), the “always hunt stag” strategy consistently dominates the entire grid, albeit slowly (sometimes requiring upwards of 1000 iterations). Increasing the payoff of hunting a stag together only serves to hasten this convergence. However, reducing the payoff of hunting a stag to 4/4 introduces an interesting phenomenon; although mutual stag hunting is still a nash equilibrium, the reduction in payoff sometimes causes the “always hunt stag” strategy to stagnate at ~50% of the total population, and occasionally even go extinct entirely. In further pursuit of this phenomenon, we experimented with a custom matrix of 21/21, 0/20, 20/0, 20/20 and observed that the “always hunt stag” strategy consistently went extinct within 30 iterations. We thus conclude that, although the Nash equilibrium structure of the games we tried in this section are all identical, the payoff ratios are important proxy variables to consider when predicting the dominance of certain strategies.

The results of the Hawk-Dove games reinforce the emergence of cooperation as a dominant strategy within spatial environments. The program allowed us to test with and without the spatial component affecting the agent’s learning dynamics. This flexibility allowed us to validate the claims that we found in our research. In his paper, “Spatial Evolutionary Game Theory: Hawks and Doves Revisited”, Killinback concludes that the emergence of Dove as a dominant strategy is directly linked to the “general consequence of spatial pattern formation” (Killinback 1143). Thus in our tests we compared results of three proportions within the population (50/50 , 25/75, 75/25) in both a spatial learning dynamic and a random learning dynamic that ignored neighbor relationships.

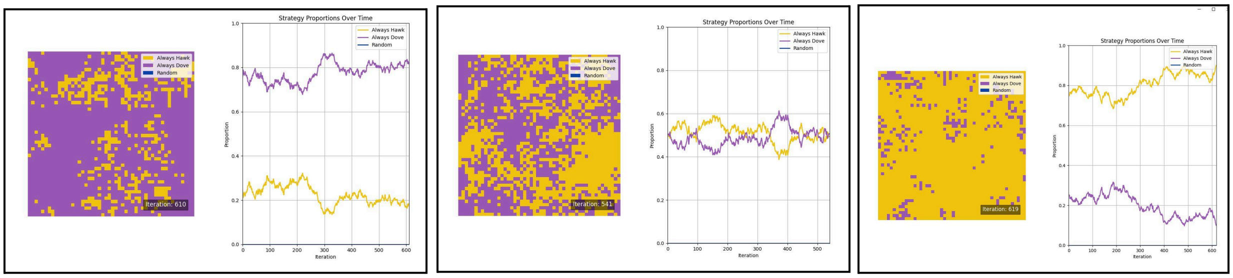

Figure 2: Results of running Hawk-Dove with non-spatial learning dynamics. The initial strategy distribution is as follows: 2a has 25% Dove 75% Hawk, 2b 50% Dove and 50% Hawk, 2c is 25% Dove and 75% Hawk.

Across all three non-spatial simulations, the population proportions remained relatively stable throughout the iterations. In these simulations there is no mechanism for local reinforcement or strategy clustering since each agent interacts randomly with others in the population, meaning that neither Hawks or Doves can consistently gain an advantage. The corresponding graphs show that the initial proportions tend to hold with minor fluctuations but no significant variation from the original values.

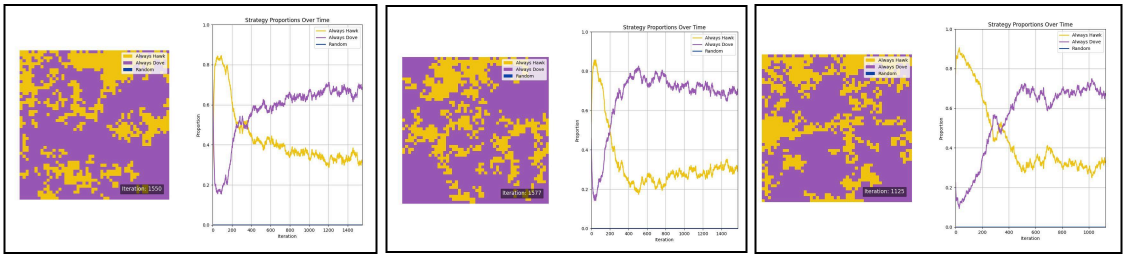

Figure 3: Results of running Hawk-Dove with spatial learning dynamics. The initial strategy distribution is as follows: 3a has 25% Dove 75% Hawk, 3b 50% Dove and 50% Hawk, 3c is 25% Dove and 75% Hawk.

In contrast, the second set of simulations introduced spatiality by placing agents on a two-dimensional toroidal grid, where interactions were limited to local neighbors. Using the same three initial conditions as before, we see that regardless of initial proportions in the population, the Dove strategy consistently emerged as dominant over time. Even when Hawks began as the majority 75/25, Doves were eventually able to form spatial clusters, which allowed them to reinforce each other's payoffs and resist invasion by aggressive strategies. These findings align strongly with Killinback's conclusion, which argues that spatial structure enables cooperative strategies to flourish through local reinforcement. The graphs generated from these spatial simulations show clear upward trends in Dove population share, with Hawk frequencies steadily declining across the board. These results illustrate how spatial interactions can fundamentally reshape evolutionary dynamics by promoting cooperation that would not emerge in well-mixed populations.

Our research into replicating the results of Dawkin's Battle of the Sexes matched our overall expected results but we were not able to directly replicate the findings from Postma's spatial simulation of the game. That is because that study's simulation replicated based on the current rounds' score while our simulation is based on the sum total score of each agent. The reason for this decision on our part is because for all other games we studied, it makes sense to see how the strategies evolve over a lifespan, while the BoS is looking at mating rituals which happen a significantly smaller amount of times in a lifespan. Despite this misalignment, we were able to arrive at both helpfulness in males and coyness in females being viable strategies. In fact they quickly tended to become the only viable strategies. In most experimental runs, unhelpfulness in both genders quickly becomes the dominant strategy, completely taking over the map except for a few instances of male helpfulness as seen in figure 4a. Afterwards, the male strategies slowly start to shift as the payout for being helpful with a coy female outshines the null payout of being unhelpful. This eventually converges with the only male strategy being helpfulness and the only female strategy being coyness as seen in figure 4b.

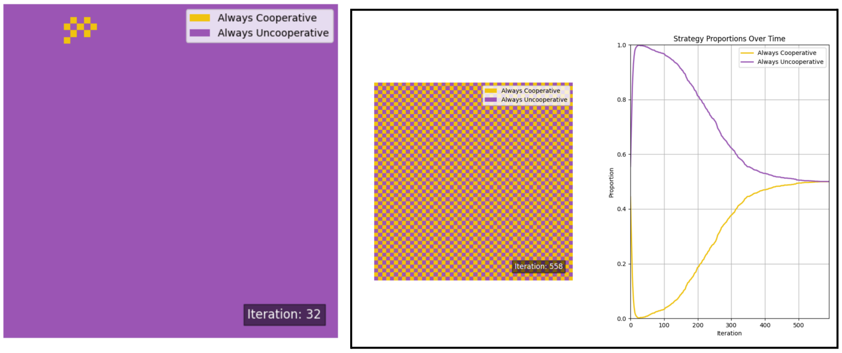

Figure 4: Results of running BoS with the replicator strategy. 4a is the initial result while 4b is the final one.

When BoS is run on a deterministic learning dynamic it quickly arrives at the check pattern as a final result. This occurs with the replicator, Fermi, and aspiration learning dynamics. The moran learning dynamic will not converge to any final result, however. As can be seen in figure 5 it will vary especially if the starting conditions favor cooperative behavior. Although it still tends towards the equilibrium mentioned previously, it requires more iterations to converge and is subject to the stochasticity of the learning dynamic. Occasionally, it will even form patches of complete cooperation which does not occur with any of the other learning dynamics.

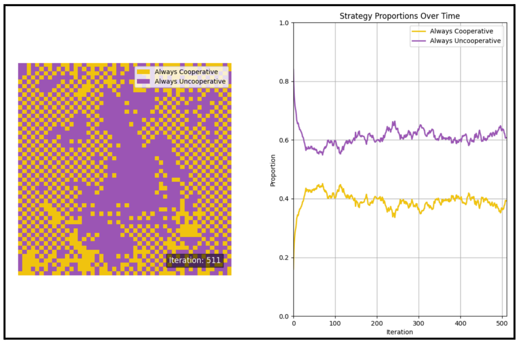

Figure 5: Results of running BoS with the moran strategy while heavily favoring cooperation as an initial condition.

The results from Rock-Paper-Scissors were more aligned with what we expected than BoS, demonstrating that our simulation works correctly with more than two options. When only the three strategies of Rock, Paper, and Scissors were utilized, the outcome was symmetric no matter what learning dynamic was selected; an equilibrium state emerges where any strategy that pulls ahead due to random chance is quickly pulled back down by the strategy that counters it; the system as a whole is *stable*. This is in line with both expected results from past simulations and real life results from nature, and can be seen in figure 6a. For exploration purposes, we introduced a fourth strategy that randomly chooses one of the three options equally. This led to some interesting results; with the replicator and the moran learning dynamics, the random strategy quickly exceeded the other strategies and became dominant, as seen in figure 6b. This could make sense since all strategies have an amortized ⅓ chance of being the winning strategy, but the spatial environment complicates this finding. The random strategy has that chance against every opponent while the other strategies only have that chance in the amortized state, and thus the random strategy is less susceptible to being invaded and becomes more dominant in the learning dynamics that lead to easier changes in strategy. This doesn't hold true for the Fermi and aspiration strategies - in fact, with Fermi it quickly becomes a worse strategy as can be seen in 6c. This is probably because the random strategy takes longer to grow its overall score and the Fermi randomness allows the other agents to build a buffer against the advantages of the random strategy making them more resilient against it in this case.

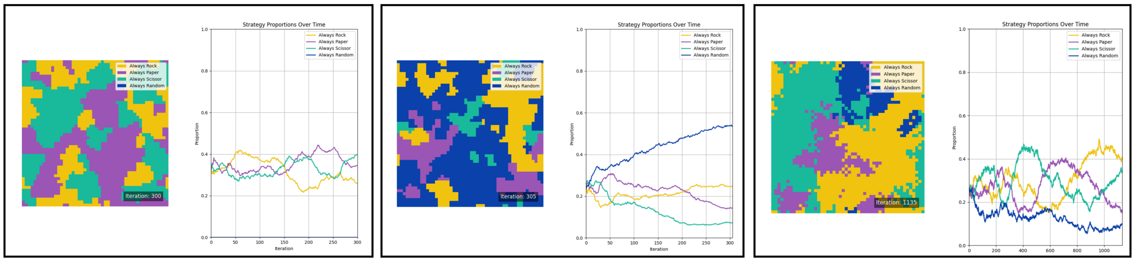

Figure 6: Results of running RPS. 6a is the initial result w/o the random strategy, 6b is with the random strategy and the replicator learning dynamic, and 6c is w/ the random strategy and Fermi learning dynamic..

Future Work

Future work will focus on increasing the complexity, accuracy, and interpretability of our simulation. A primary next step is to conduct systematic parameter sensitivity analyses to understand how small changes in payoff structures, mutation rates, interaction rules, and other initial factors can affect the stability of cooperation. This will help identify threshold effects and define the boundaries within which cooperative behavior is evolutionarily viable.

Expanding beyond uniform spatial grids, future simulations could explore more realistic network structures which better approximate the uneven and clustered nature of biological and social interactions. Incorporating these features could reveal how network architecture affects the diffusion and resilience of cooperation.

Environmental variability is another valuable addition. Introducing resource fluctuations, seasonal cycles, or random shocks could test whether cooperation remains stable in unpredictable settings. Incorporating mechanisms like kin selection or group selection would add biological depth, allowing exploration of how genetic relatedness or group-level fitness influences cooperative dynamics. This aligns with empirical observations where cooperation often evolves in groups with shared lineage or common interests. Together, these extensions aim to make the simulation not only more realistic but also more broadly applicable to the study of cooperation in biological, ecological, and social systems.

Conclusion

We have endeavored to study the emergence of cooperation in nature and create a simulation to model this emergence. Through existing literature research we identified various games that can express cooperation as a strategy. We have also found real world instances of those games to get a baseline of how nature has evolved to form cooperation. Utilizing those baselines and a simulation we created, we have experimented to try and find starting conditions that cause the games to arrive at the endpoint seen in nature. Our simulation is modular and flexible enough to model multiple different games with various starting conditions. Our data driven findings should help bridge the gap between theoretical behaviour and real-world behaviours seen in nature.

References

Alex McAvoy, Julian Kates-Harbeck, Krishnendu Chatterjee, Christian Hilbe, Evolutionary instability of selfish learning in repeated games, PNAS Nexus, Volume 1, Issue 4, September 2022, pgac141, https://doi.org/10.1093/pnasnexus/pgac141

Arnold, C. (2020). Biodiversity May Thrive through Games of Rock-Paper-Scissors. *Quanta Magazine*. www.quantamagazine.org/biodiversity-may-thrive-through-games-of-rock-paper-scissors-20200305/

Axelrod, R., & Hamilton, W. D. (1981). The evolution of cooperation. *Science.* https://websites.umich.edu/~axe/research/Axelrod%20and%20Hamilton%20EC%201981.pdf

Chen, et. al. (2017). Evolutionary dynamics of N-person Hawk-Dove games. *Scientific Reports*. https://doi.org/10.1038/s41598-017-04284-66

Cowden, C. C. (2012) Game Theory, Evolutionary Stable Strategies and the Evolution of Biological Interactions. Nature Education Knowledge 3(10):6

Hoffman, M., Suetens, S., Gneezy, U. et al. (2015). An experimental investigation of evolutionary dynamics in the Rock-Paper-Scissors game. *Sci Rep 5, 8817*. https://doi.org/10.1038/srep08817

Kahneman, D., & Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica, 47*(2), 263–291. https://doi.org/10.2307/1914185

Kerr, B., Riley, M., Feldman, M. *et al.* Local dispersal promotes biodiversity in a real-life game of rock–paper–scissors. *Nature* 418, 171–174 (2002). https://doi.org/10.1038/nature00823

Killingback, T., & Doebeli, M. (1996). Spatial evolutionary game theory: Hawks and Doves revisited. *Proceedings: Biological Sciences*, *263*(1374), 1135–1144. https://www.jstor.org/stable/50511

Leimar Olof and McNamara John M. 2023 Game theory in biology: 50 years and onwardsPhil. Trans. R. Soc. B37820210509 http://doi.org/10.1098/rstb.2021.0509

Littman, M. L. (1994). Markov games as a framework for multi-agent reinforcement learning. *Proceedings of the Eleventh International Conference on Machine Learning*. https://courses.cs.duke.edu/spring07/cps296.3/littman94markov.pdf.

Lu, J. (2023). *Computational validation of a mathematical model of stable multi-species communities in a Hawk Dove game*. Preprint.

https://www.researchgate.net/publication/369759382_Computational_Validation_of_a_Mathematical_Model_of_Stable_Multi-Species_Communities_in_a_Hawk_Dove_Game

Maynard Smith J, Price GR. 1973 The logic of animal conflict. Nature 246, 15-18. (doi:10.1038/246015a0)

Nowak, M. A., & Sigmund, K. (1993). A strategy of win-stay, lose-shift that outperforms tit-for-tat in the Prisoner's Dilemma game. *Nature*, 364, 56–58. https://doi.org/10.1038/364056a0

Nowak, M. A. (2012). Evolving cooperation. *Journal of Theoretical Biology.* https://projects.iq.harvard.edu/files/ped/files/jtb12e_0.pdf

Nowak MA, Sigmund K. Evolutionary dynamics of biological games. Science. 2004 Feb 6;303(5659):793-9. doi: 10.1126/science.1093411. PMID: 14764867.

Postma, Ivar. (2012). A Spatial Simulation of the Battle of the Sexes, University of Groningen. https://fse.studenttheses.ub.rug.nl/10537/1/Spatial_Simulation_of_the_BotS_1.pdf.

Sadekar, Onkar & Civilini, Andrea & Gómez-Gardeñes, Jesus & Latora, Vito & Battiston, Federico. (2023). Evolutionary game selection creates cooperative environments. 10.48550/arXiv.2311.11128.

Trivers RL. 1971 The evolution of reciprocal altruism. Q. Rev. Biol. 46, 35-57. (doi:10.1086/406755)

Wei, H., Zhang, J., & Zhang, C. (2024). Evolutionary dynamics of direct and indirect reciprocity on networked populations. *Swarm and Evolutionary Computation, 88*, 101611. https://doi.org/10.1016/j.swevo.2024.101611

Wilson, D. S., & Wilson, E. O. (2007). Rethinking the theoretical foundation of sociobiology. *The Quarterly Review of Biology, 82*(4), 327–348. https://doi.org/10.1086/522809