

COOPERATION GAME

The Signal Loom · Field Guide

Sahilwho / Sahilverse · A print companion for the Iterated Cooperation Arena

An illustrative teaching model, not a test of human character or a scientifically validated score.

01 / The dilemma is a structure, not a personality test

Two people can each have a reason to make a choice that leaves both worse off. That is the central tension of the Prisoner's Dilemma. In its familiar form, each player independently chooses to cooperate or defect. Mutual cooperation rewards both. A unilateral defector does better than the cooperator, but mutual defection pays less than mutual cooperation. The dilemma is not proof that humans are selfish. It describes incentives under particular assumptions: fixed payoffs, limited communication, and choices made without enforceable agreements. Real relationships are more complex. Cooperation may be rewarded by institutions, identities, social norms, or future opportunities. A model is useful when its boundaries remain visible. The arena below lets you inspect one particular repeated-game model rather than diagnosing anyone's character. The score is illustrative game currency, not a measure of trust, morality, or intelligence. A cooperative action may be wise in one setting and costly in another. A defection may reflect fear, error, incentives, or something else entirely. The game does not tell us which.

02 / Read the payoff architecture

The arena uses four outcomes: mutual cooperation gives each player three points; mutual defection gives each one point; a defector facing a cooperator gets five points while the cooperator gets zero. Game theorists often write these as $R=3$, $P=1$, $T=5$, and $S=0$. The inequalities $T > R > P > S$ create the dilemma. Here twice the reward for cooperation also exceeds the combined temptation and sucker payoffs. These numbers are conventional teaching values, not measured human rewards. They let us hold incentives constant while changing strategy and accidental mistakes. Each round, the two strategies independently choose an intended action from the available history. A possible execution error then flips an action. Both executed actions, not intentions, determine the score. Watch for the difference: an apparent betrayal might be noise rather than a strategic choice. Repeated games invite a second question: will the next encounter change what makes sense now?

03 / Six strategies, six different commitments

Always Cooperate offers cooperation on every turn, even after being exploited. Always Defect chooses defection every turn and never tests reconciliation. Tit-for-Tat starts cooperatively and then copies the opponent's previous executed action. Generous Tit-for-Tat sometimes forgives a previous defection; the probability is fixed at one in four for this exhibit, driven by a reproducible seed. Win-Stay Lose-Shift starts cooperatively and repeats its previous executed move after receiving either the reward for mutual cooperation or the temptation payoff; after either other payoff, it switches. Grim Trigger cooperates until the opponent has defected once, then defects forever. These are formal decision rules, not personality categories. Their strengths depend on the opponent, error rate, horizon, and payoff structure. Tit-for-Tat can be interpretable and responsive but can also lock into retaliation after a mistake. Grim Trigger can discourage exploitation but may make a single error permanent. Generosity can repair a relationship but can also be exploited. There is no universal champion.

04 / Time changes the question

A single round invites immediate temptation: if the other cooperates, defection pays five instead of three; if the other defects, defection pays one instead of zero. Repetition changes the setting because today's move can influence tomorrow's response. If players expect many future encounters, the value of maintaining cooperation may outweigh a one-time gain. Yet the conclusion depends on what players know about the future. A known last round can create incentives to defect and unravel cooperation in finite games under narrow rationality assumptions. In practice, people may have incomplete information, social commitments, or preferences not represented by a simple payoff table. In the exhibit you choose a fixed number of rounds. Strategies here do not reason backward from the ending; that is a deliberate simplification. Increasing rounds therefore shows how these particular rules interact over time, not what perfectly rational agents must do.

05 / Noise is not the same as malice

Imagine a message is lost, a button is misread, or an agreement is implemented incorrectly. In a repeated game, even a small execution error can change the trajectory. Tit-for-Tat against itself can alternate retaliation after a single accidental defection. More forgiving strategies may recover, though forgiveness can have costs against persistent exploitation. Our noise slider controls the independent probability that each player's intended action flips before the round is scored. Zero percent means intentions and actions match. Twenty percent means each action has a one-in-five chance of flipping. A seeded pseudo-random generator makes runs reproducible; it does not make the model predictive. If you repeat a run with identical settings and seed, you should get identical results. Try the same pairing with noise at zero and then ten percent. Ask not only who scored more, but what happened to mutual cooperation and how many turns were spent in a retaliatory cycle.

06 / Case file: Axelrod's computer tournaments

Political scientist Robert Axelrod organized computer tournaments in which submitted strategies played repeated Prisoner's Dilemma games against one another. His 1984 book *The Evolution of Cooperation* popularized findings about how simple contingent cooperation can perform well in a particular tournament ecology. Tit-for-Tat, associated with Anatol Rapoport, became a prominent example. The result is often flattened into "be nice and copy the other person," but that misses the role of tournament rules, opponents, scoring, and the absence or presence of error. A strategy that thrives in one field of competitors can fare differently in another. The tournaments were computational demonstrations, not randomized tests of universal human behavior. Their enduring value is methodological: specify rules, expose assumptions, compare outcomes, and test how conclusions change when the environment changes. This arena follows that spirit on a smaller scale, without claiming to reproduce Axelrod's original tournament design or results.

07 / Case file: governing shared resources

Political scientist Elinor Ostrom studied how communities govern shared resources, including irrigation systems, forests, and fisheries. Her work challenged the claim that shared resources inevitably require either privatization or centralized control to avoid overuse. She documented institutional arrangements involving locally adapted rules, monitoring, graduated sanctions, and ways to resolve conflicts. These settings are not two-player Prisoner's Dilemmas. They involve multiple actors, changing ecological conditions, unequal power, history, and rule-making itself. That difference is precisely why the comparison is useful. The arena holds the rules fixed; communities can redesign rules and institutions. A cooperative outcome may depend less on finding the perfect individual strategy than on creating trustworthy information, credible commitments, and fair ways to repair breaches. The case invites you to ask what is missing from a tiny game when it is used as a metaphor for real cooperation.

08 / Competing lenses and contested interpretations

One perspective emphasizes strategic reciprocity: repeated encounters, reputations, and credible responses can make cooperation stable without assuming saintly motives. Another emphasizes institutions, social preferences, communication, and power: cooperation is not reducible to a sequence of anonymous choices. Experiments in public-goods games have found that punishment can sustain contributions under some conditions, but punishment can also be costly, unequal, or misused. The importance of a finding depends on the population, laboratory rules, cultural context, and the institutions surrounding a choice. Evolutionary models offer yet another lens, asking under what conditions cooperative traits can spread. These explanations are not interchangeable. A computer strategy winning points does not establish an ethical rule for friendship, work, or politics. In particular, retaliation is not automatically justice, and forgiveness is not always safe. Interpret model output as a demonstration of conditional mechanisms, not a prescription.

09 / What the arena cannot prove

This simulation cannot establish whether a person is trustworthy, what motivates an individual action, whether cooperation is morally required, or which institutional design is best. Its players have no emotions, negotiation, identity, unequal resources, learning beyond their fixed rule, or outside options. The payoffs are made-up teaching numbers. The randomness is algorithmic and seeded, not sampled from the world. Even the strategy names hide variations across academic implementations. “Generous” here means one specific forgiveness probability; other authors may implement generosity differently. “Win-Stay Lose-Shift” here follows a defined threshold rule. These details are visible so you can challenge them. The useful question is not “which strategy wins life?” but “what assumptions make this result happen, and what would change it?” That habit of inspecting assumptions is the real learning objective.

10 / Guided experiment: the repair test

Start with Tit-for-Tat against Tit-for-Tat, twenty rounds, zero noise. Step through five turns and observe how mutual cooperation persists. Reset and raise noise to ten percent. Keep the same seed and run all twenty rounds. Note when intended and executed actions diverge. Undo a turn and replay it to see that the result is reproducible. Then replace the second strategy with Generous Tit-for-Tat and compare the number of cooperative rounds, the two scores, and the visible sequence of responses. Finally choose Always Defect as one player. Does forgiveness still help the same way? Write down a claim, the exact settings supporting it, and one reason the claim might fail in a different setting. You can copy the transcript for your own notes; the site does not upload it. If JavaScript is unavailable, use the printed payoff matrix and the walkthrough in the Field Guide to conduct the first five rounds manually.

11 / Beyond winning: a better set of questions

The highest individual score is not always the same as the highest combined score. Two players can cooperate to produce six points together, while unilateral defection produces five in total. That distinction matters because competition over relative advantage can destroy shared value. Likewise, a rule that resists exploitation may sacrifice opportunities to repair an honest mistake. Look at both players’ scores, their combined score, and the cooperation rate before declaring an outcome successful. Then imagine a third player, a mediator, a public record, or a way to renegotiate payoffs. Each addition changes the game itself. This is how simple models can support sophisticated thinking: they force us to state which features we have included and which we have left out. The goal is not to romanticize cooperation. It is to recognize that cooperation can be strategically fragile, institutionally supported, and ethically complicated at the same time.

12 / Design a different game, not just a different move

Suppose players can communicate before choosing. Suppose a referee can verify agreements. Suppose a third party can compensate someone who cooperates after a loss. Suppose the temptation payoff is smaller, or

repeated cooperation produces a growing shared resource rather than three fixed points. Each change alters the problem, and some changes remove the original Prisoner's Dilemma entirely. That is not cheating: designing better institutions often means changing incentives rather than hoping people will become more cooperative by willpower alone. In a workplace, a transparent procedure for correcting errors can be more effective than expecting everyone to perfectly infer intentions. In a community, monitoring can support trust when rules are legitimate and accountability is reciprocal, but surveillance can also undermine it. In a family, the ability to discuss expectations may matter more than any abstract retaliation rule. These analogies are prompts for analysis, not direct predictions from a two-player model. As you experiment, keep three notebooks in mind: the action that happened, the rule that generated it, and the context that made the rule sensible or foolish. When those notebooks are confused, simple models become slogans. When they are separated, a game can become a precise language for asking better questions. The most productive outcome may be to discover that a dilemma was framed too narrowly. If the only available buttons are cooperate and defect, the model cannot represent a request for clarification, a pause, a negotiated repair, or a decision to leave the game. Real-world agency often begins with inventing those missing buttons.

Play without a screen

Choose strategies and play one round at a time. A strategy selects C or D based on the preceding round. Use the matrix below to add points. Tit-for-Tat starts C and thereafter copies the other player's previous actual action. Always Defect chooses D every time. Against each other, round one is C/D and scores 0/5; subsequent rounds are D/D and score 1/1 each. Compare this with Tit-for-Tat against itself, which maintains C/C in the absence of error.

A \ B	Cooperate	Defect
Cooperate	3, 3	0, 5
Defect	5, 0	1, 1

Reflection prompts

- What changes when the same two strategies meet repeatedly?
- Which strategy appears forgiving, and what is its cost?
- What does an accidental defection do to subsequent moves?
- Is maximizing individual points the same as maximizing combined points?
- Which features of real-world cooperation are absent from this model?

Further reading · Bibliographic leads (URL verification pending)

Axelrod, R. (1984). *The Evolution of Cooperation*. Basic Books. — <https://books.google.com/books?id=NZx9AAAAMAAJ>

Axelrod, R., & Hamilton, W. D. (1981). The evolution of cooperation. *Science*, 211, 1390–1396. — <https://doi.org/10.1126/science.7466396>

Rapoport, A., & Chammah, A. M. (1965). *Prisoner's Dilemma*. University of Michigan Press. — <https://press.umich.edu/Books/P/Prisoner-s-Dilemma>

Ostrom, E. (1990). *Governing the Commons*. Cambridge University Press. — <https://doi.org/10.1017/CBO9780511807763>

Nowak, M. A. (2006). Five rules for the evolution of cooperation. *Science*, 314, 1560–1563. — <https://doi.org/10.1126/science.1133755>

Fehr, E., & Gächter, S. (2002). Altruistic punishment in humans. *Nature*, 415, 137–140. — <https://doi.org/10.1038/415137a>

Trivers, R. L. (1971). The evolution of reciprocal altruism. *Quarterly Review of Biology*, 46, 35–57. — <https://doi.org/10.1086/406755>

Dal Bó, P. (2005). Cooperation under the shadow of the future: Experimental evidence from infinitely repeated games. *American Economic Review*, 95, 1591–1604. — <https://doi.org/10.1257/000282805775014434>

Milinski, M., Semmann, D., & Krambeck, H.-J. (2002). Reputation helps solve the tragedy of the commons. *Nature*, 415, 424–426. — <https://doi.org/10.1038/415424a>

Return to the Sahilverse: <https://sahiltayal.com> · Write to Sahil: sahil@sahiltayal.com