



Strategies for the Iterated Prisoner's Dilemma

Name	Abbreviation	Description
Unconditional Cooperator	Cu	Cooperates unconditionally.
Unconditional Defector	Du	Defects unconditionally.
Random	Random (=C.5 or R(.5,.5,.5) or S(.5,.5,.5,.5) below)	Defects unconditionally.
Probability p Cooperator	Cp for $0 \leq p \leq 1$	Cooperates with fixed probability p .
Tit for Tat	TFT (=R(1,1,0) or S(1,0,1,0) below)	Cooperates on the first round and imitates its opponent's previous move thereafter.
Suspicious Tit for Tat	STFT (=R(0,1,0) below)	Defects on the first round and imitates its opponent's previous move thereafter.
Generous Tit for Tat	GTFT (=R(1,1,g(R,P,T,S)) below)	Cooperates on the first round and after its opponent cooperates. Following a defection, it cooperates with probability $g(R, P, T, S) = \min\{1 - \frac{T-R}{R-S}, \frac{R-P}{T-P}\}$, where R, P, T and S are the reward, punishment, temptation and sucker payoffs.
Gradual Tit for Tat	GrdTFT	TFT with two differences: (1) it increases the string of punishing defection responses with each additional defection by its opponent (2) it apologizes for each string of defections by cooperating in the subsequent two rounds.
Imperfect TFT	ImpTFT	Imitates opponent's last move with high (but less than one) probability.
Tit for Two Tats	TFTT (or TF2T)	Cooperates unless defected against twice in a row.
Two Tits for Tat	TTFT (or 2TFT)	Defects twice after being defected against, otherwise cooperates.

Omega Tit for Tat	ΩTFT	Plays TFT unless measures of deadlock or randomness exceed specified thresholds. When deadlock threshold is exceeded it cooperates and resets the measure. When randomness threshold is exceeded, it switches to unconditional defection. For full specification see Slaney and Kienreich, p184. ΩTFT finished second in the 2005 reprise of the Axelrod IPD tournament.
GRIM (or TRIGGER)	GRIM (= S(1,0,0,0) below)	Cooperates until its opponent has defected once, and then defects for the rest of the game.
Discriminating Altruist	DA	In the Optional IPD, cooperates with any player that has never defected against it, and otherwise refuses to engage.
Pavlov (or Win-stay, Lose-shift)	WSLS (= P₁ below)	Cooperates if it and its opponent moved alike in previous move and defects if they moved differently.
n-Pavlov	P_n	Adjusts its probability of cooperation in units of $\frac{1}{n}$ according to its payoff on the previous round. More specifically it cooperates with probability $p_1 = 1$ on round 1 and probability p_{n+1} on round $n + 1$, where $p_{n+1} = p_n \left[\begin{array}{l} + \\ - \end{array} \right] \frac{1}{n} \text{ if payoff on last round was Reward (R)}$ $p_n \left[\begin{array}{l} + \\ - \end{array} \right] \frac{1}{n} \text{ if payoff on last round was Punishment (P)}$ $p_n \left[\begin{array}{l} + \\ - \end{array} \right] \frac{2}{n} \text{ if payoff on last round was Temptation (T)}$ $p_n \left[\begin{array}{l} + \\ - \end{array} \right] \frac{2}{n} \text{ if payoff on last round was Sucker (S),}$ <p>p_n is the probability of cooperation on round n, $x \left[\begin{array}{l} + \\ - \end{array} \right] y = \min(x + y, 1)$ and $x \left[\begin{array}{l} + \\ - \end{array} \right] y = \max(x - y, 0)$.</p>
Adaptive Pavlov	APavlov	Employs TFT for the first six rounds, places opponent into one of five categories according to its responses and plays an optimal strategy for each. Details described in Li pp 89-104. APavlov was the highest scoring strategy in the 2005 reprise of Axelrod's IPD tournament.
Reactive (with parameters y, p, q)	R(y,p,q)	Cooperates with probability y in first round and with probabilities p or q after opponent cooperates or defects
Memory-one (with parameters p, q, r, s)	S(p,q,r,s)	Cooperates with probabilities p, q, r or s after outcomes (C,C), (C,D), (D,C) or (D,D).
Zero Determinant	ZD	A class of memory-one strategies that guarantee that a player's long-term average payoff in the infinitely repeated, two-player prisoner's dilemma (2IPD) will be related to his opponent's according to a fixed linear equation.
Equalizer (or dictator)	SET-n (for $P \leq n \leq R$)	A ZD strategy that guarantees the opponent's long term average payoff is n . As it turns out, in a PD with payoffs 5,3,1 and 0, SET-2 = S(3/4,1/4,1/2,1/4) .
Extortionary	Extort-n	An extortionary strategy is a ZD strategy that guarantees that an opponent's average payoff can exceed the punishment payoff only if one's own long term average payoff is greater. Extort-n guarantees

that one's gain over punishment is n times one's opponent's. As it turns out, for a PD with the payoffs above, **EXTORT-2**= $S(\frac{7}{8}, \frac{7}{16}, \frac{3}{8}, 0)$.

Generous **Gen-n**

A generous strategy is a **ZD** strategy that guarantees that an opponent's average payoff can be lower than the reward payoff only if one's own long term average payoff is even lower. **GEN-n** guarantees that one's loss relative to the reward is n times one's opponent's. As it turns out, for a PD with the payoffs above, **GEN-2**= $S(1, \frac{9}{16}, \frac{1}{2}, \frac{1}{8})$.

Good **GOOD**

A good strategy for the infinitely-repeated, two-player PD is a strategy with the following properties: (1)its use by both players ensures that each gets reward as long-term average payoff, (2)it is a nash-equilibrium with itself, and (3)if it is employed by both, any deviation by one that reduces the average payoff of the other will also reduce its own average payoff. Aikin, 2013 provides a simple characterization of the memory-one strategies that are good.

[Return to Prisoner's Dilemma Entry](#)

[Copyright © 2019](#) by

[Steven Kuhn <kuhns@georgetown.edu>](mailto:kuhns@georgetown.edu)

Open access to the Encyclopedia has been made possible, in part, with a financial contribution from the Portland State University Library. We gratefully acknowledge this support.

The Stanford Encyclopedia of Philosophy is [copyright © 2023](#) by [The Metaphysics Research Lab](#), Department of Philosophy, Stanford University

Library of Congress Catalog Data: ISSN 1095-5054