Lab 7: The Iterated Prisoner's Dilemma

Two Players Model

Start the NetLogo model "PD Two Person Iterated", which is located in the Models Library under Social Science → (unverified) → Prisoner's Dilemma. The following strategies are available:

- *act-randomly* – Randomly cooperate or defect.
- *cooperate* – Cooperate always.
- *defect* – Defect always.
- *tit-for-tat* – If the opponent cooperates this round, cooperate next round. If the opponent defects this round, defect next round. Initially cooperate.
- *tit-for-two-tats* – If the opponent cooperates this round cooperate next round. If the opponent defects two rounds in a row, defect the next round. Initially cooperate.
- *unforgiving* – Cooperate always unless the opponent defects once. Upon opponent defection retaliate by defecting always.
- *custom-strategy* – Not used.

1. Set the SELECT-COMPUTER-STRATEGY? switch to ON, which will allow you to specify the strategy used by the computer (when the switch is OFF, the computer's strategy is chosen at random). Manually select your own strategy from the *human-strategy* menu, and the computer's strategy from the *computer-strategy* menu. The PLAY ONCE button will play one round of the game, following the selected strategies. Clicking PLAY REPEATEDLY will automate the game for as long as the button is pressed (click it again to stop). The payoff matrix is shown below:

	Computer cooperates	**Computer defects**
You cooperate	You both get 3 points	You get 0, Computer gets 5
You defect	You get 5, Computer gets 0	You both get 1 point

2. Experiment with playing different strategies against one another. A plot of the average score over time is shown for both you and the computer, with the maximum possible average score being 5 points per interaction. Which strategies do the best? Which do the worst?

3. If the computer always plays strategy *cooperate*, then which strategy for the user results in the highest average long-term score? Why?

4. If the computer always plays strategy *defect*, what does the average score plot look like over time when the user plays strategy *tit-for-tat*, *tit-for-two-tats*, or *unforgiving*, respectively? Why?

5. What does the average score plot look like when the computer plays strategy *tit-for-tat* and the user plays strategy *cooperate*, *tit-for-tat*, *tit-for-two-tats*, and *unforgiving*, respectively? Why?

N Players Model

6. Next, open the model "PD N-Person Iterated". The previous Two-Person model demonstrated an interesting concept: When interacting with someone over time in a prisoner's dilemma scenario, it is possible to tune your strategy to do well with theirs. Each possible strategy has unique strengths and weaknesses that appear through the course of the game. For instance, the *defect* strategy does the best of any against *act-randomly*, but poorly against itself. *Tit-for-tat* does poorly with the *act-randomly* strategy, but well with itself.

   This makes it difficult to determine a single "best" strategy. One such approach to doing this is to create a world with multiple agents playing a variety of strategies in repeated prisoner's dilemma situations. The N-Person model does just that. The turtles with different strategies wander around randomly until they find another turtle to play with. Note that each turtle remembers their last interaction with each other turtle. While some strategies don't make use of this information, other strategies do. When two turtles interact, they display their respective payoffs as labels.

   For these experiments, you should set *n-unknown* to 0, since we won't be using the *unknown* strategy.

7. Observe the results of running the model with a variety of populations and population sizes. For example, can you get *cooperate*'s average payoff to be higher than *defect*'s? Can you get *tit-for-tat*'s average payoff higher than *cooperate*'s? What do these experiments suggest about an optimal strategy?

8. Set all the numbers of players for each strategy to be equal in distribution. For which strategy does the average payoff seem to be highest? Do you think this strategy is always the best to use, or will there be situations where other strategies will yield a higher average payoff?

9. Set the number *n-cooperate* to be high, *n-defect* to be equivalent to that of *n-cooperate*, and all other players to be 0. Which strategy will yield the higher average payoff?

10. Set the number *n-tit-for-tat* to be high, *n-defect* to be equivalent to that of *n-tit-for-tat*, and all other players to be 0. Which strategy will yield the higher average payoff? What do you notice about the average payoff for *tit-for-tat* players and *defect* players as the iterations increase? Why do you suppose this change occurs?

PD Basic Evolutionary Model

11. The "PD Basic Evolutionary" model explores how cooperators can spread through a population dominated by defectors, or vice versa, using a simple spatial configuration where each player interacts with its eight surrounding neighbors in a 2-dimensional grid world. Each player is color-coded as follows: RED players defected on the last 2 time steps; BLUE players cooperated on the last 2 time steps; GREEN players cooperated but then defected; YELLOW players defected but then cooperated. The *initial-cooperation* parameter controls the percentage of pure cooperators (BLUE) versus pure defectors (RED) in the initial population. The *defection-award* parameter controls the size of the "Temptation payoff", that is, the amount of reward received by a player if it defects when its opponent cooperates. Read through the description under the *Info* tab for further information about the model.

12. Experiment with this model by observing the effects of different initial settings. How does varying the "Temptation payoff" influence the outcome? Under what conditions does the population evolve to a stable uniform state consisting of all players of the same type? Can stable "mixed" states arise?