Biology 4415/5415
Evolution

LABORATORY EXERCISE 6:

THE PRISONER'S DILEMMA

The **Prisoner's Dilemma** originally was a classic problem in the mathematical field of game theory. In the original version of the Prisoner's Dilemma problem, there are two criminals who are in prison for a relatively minor crime, serving, say, a three-year sentence. The prosecutor thinks that both are guilty of a much more serious crime—perhaps a murder—but can't prove it with the evidence he has. So he meets with each prisoner individually, without the other one knowing of the meeting, and he makes the prisoners a deal: If one of them confesses to the serious crime, and rats on the other one, the one who confesses will go free and the one who doesn't will get twenty-five years in prison. However, if they both confess, both will get ten-year sentences. If neither one does, of course, the prosecutor won't be able to make a case, and both will end up serving three years for their more minor offense.

We can draw up a matrix of payoffs for the four possible outcomes. Each prisoner has to choose whether to **cooperate** (that is, cooperate with his fellow prisoner, and refuse to snitch) or **defect** (tattle to the authorities). If we call the prisoners A and B, then the outcomes are:

	A Cooperates	**A Defects**
B Cooperates	**A: -3**	**A: 0**
	B: -3	**B: -25**
B Defects	**A: -25**	**A: -10**
	B: 0	**B: -10**

(The outcomes are written as negative numbers because they're penalties. If they were bonuses—say, if a defector got not only freedom but a $25,000 reward—then we'd write positive numbers.)

As the rules stand, the best thing to do is to defect—*if* the Dilemma is presented only once. Why? Briefly: the worst that can happen to a defector (ten years in prison) is better than the worst that can happen to a cooperator (25 years in prison), and the best that can happen to a defector (freedom) is better than the best that could happen to a cooperator (three years in prison). This applies whether the payoffs are negative, as in the example above using prison time, or positive, as it would be if winning money were the aim of the game. And if the game is played only once, neither A nor B can retaliate if he thinks he's been suckered.

Another way to look at it is this: Suppose that A knew that B would decide what to do by flipping a coin, then if A defected, he'd have a 50% chance of going free, a 50% chance of serving ten years—for an *expectation* of (0.5)(0) + (0.5)(-10) = -5. If A cooperated, his expectation would be (0.5)(-3) + (0.5)(-25) = -14. A's best move is to defect. For the same reason, so is B's.

A single Prisoner's Dilemma isn't very interesting, and isn't all that biologically realistic. However. . . what would happen if, after both were out of prison, A and B were caught again and put in the same situation? What should a prisoner choose if he knows that his buddy has cooperated with him before? Or ratted on him before? What if the Dilemma is presented to both criminals, each with different partners? What if the goal isn't to minimize a single penalty (here, a term in jail) but to minimize the total penalty over a long lifetime?

If both players in a Prisoner's Dilemma know that the game will be repeated a finite number of times, and know how many times the game will be repeated, defection is still the best option. But tn a game like this that's repeated indefinitely, in which neither player knows how many times the game will be repeated, the "winning strategy" is not so simple. In fact, it turns out that cooperation works better than defection over the long term. This indefinite game is called an **Iterated Prisoner's Dilemma**, and it has applications in everything from economics and political science to evolutionary biology.

In evolution, the Iterated Prisoner's Dilemma can be used to look at the question: How can cooperation evolve? Does natural selection shape organisms in ways that benefit them at the expense of others? The answer, surprisingly, turns out to be "no". In computer simulations in which "virtual organisms" interact, either in random pairings or in "round-robin tournaments", limited cooperation is the most successful strategy over the long term. Many organisms interact in ways that can be modeled as Iterative Prisoner's Dilemmas.

MATERIALS:

strategy cards (three types, labeled AC, AD, TFT)
encounter cards (two types, labeled Cooperate and Defect)
"life energy" tokens (beans, beads, marbles, pennies, etc.)

1. We start the game with twelve players. Each player starts with twelve tokens, representing some sort of life resource (energy, food, etc.) Each player also starts with two cards labeled "Cooperate" and "Defect".

2. Each player randomly draws a card that gives him or her a strategy to play. **DO NOT SHOW YOUR CARD TO ANYONE, OR TELL ANYONE WHAT YOUR STRATEGY IS!** There are three options:

- Always Cooperate (AC). A player playing the AC strategy always cooperates, no matter what has gone before.

- Always Defect (AD). A player playing the AD strategy always defects. Again, this is completely regardless of anything that has happened before.
- Tit for Tat (TFT). A Tit for Tat player must cooperate on the first move, and afterwards must copy whatever was done to him on the previous move.

Notice that the TFT strategy involves "remembering" past experiences, at least for one round—and is capable of both "aggression" (defecting on others) and "forgiveness" (cooperating even after having been defected on).

3. Players wander aimlessly and randomly until the referee shouts "Pair up!", when each player pairs with the closest other player. As soon as all players have paired up, on the referee's signal each player simultaneously shows the other either the "Cooperate" or "Defect" card, depending on which strategy is being played.

4. As soon as the strategy cards are played, the referee will award tokens according to the following rules:

- If both players cooperate, each player gets two tokens.
- If both players defect, each player loses two tokens.
- If one player cooperates and the other defects, the defector gets four tokens and the cooperator loses four.

We can represent the payoffs in a matrix like this:

	A Cooperates	**A Defects**
B Cooperates	**A: +2** **B: +2**	**A: +4** **B: -4**
B Defects	**A: -4** **B: +4**	**A: -2** **B: -2**

5. Any player who gets twenty-four tokens or more automatically "reproduces". In other words, he or she calls in a new player from the sidelines, who enters the game and must play the "parent's" strategy. ("Parents" should make sure that no one else but their "offspring" find out what their strategy is.) The "parent" gives twelve of his or her tokens to the "offspring" and retains the rest. As various players leave and enter the game, make certain that both outer and inner circles have the same number of players

6. A player who loses all tokens is out of the game.

7. The referee will keep track of how many players are in the game, and which strategies they are playing. The game will continue until only one strategy is left, or until the referee decides to quit, whichever comes first. It is likely that we will have to play for several dozen rounds in order to see any pattern, and we may have to repeat the entire game several times over. If the population becomes larger than the number of available students, we may have to start over with a slight change of the rules—such as increasing the threshold for "reproducing", or increasing the penalty for being defected on.

8. Other strategies are possible (although perhaps not all are biologically realistic). Some that have been tried in simulations—that we might try, if there's time—include:

- Mistrustful—like Tit for Tat, but its first move is always defecting, not cooperating.
- Remorseful—like Tit for Tat, but if it defects and its opponent cooperates, it will automatically cooperate on its next two moves
- Suspicious—like Tit for Tat, but if it cooperates and its opponent defects, it will automatically defect on its next two moves
- Pavlov—if you won tokens on the last move, play the same move as you did before; if you didn't win any tokens, switch to the other option.
- Periodic CD—alternates cooperating and defecting, regardless of what the other players do. [Further variations include Periodic DC, Periodic CCD, Periodic DDC, etc.—these are self-explanatory.]
- Random—determines move by a coin flip or other randomization method.

More complex strategies exist that involve knowing your opponent's past behavior. Examples include:

- If your opponent has Defected before at least twice, Defect; otherwise Cooperate.
- If your opponent has Defected on >50% of all his/her previous encounters, Defect; otherwise, Cooperate.
- If the last round had more players cooperating than defecting, Defect; otherwise Cooperate.

These more complex strategies are probably more realistic for modeling the behavior of more intelligent animals. Humans, for instance, tend to keep a mental score of who in their circle of acquaintances is trustworthy and who is not; so do chimps and some others. These strategies have been studied in computer simulation. However, they are somewhat cumbersome to simulate in this sort of experiment.

10. At the end of the day, the referee will collate all data, and provide you with copies at the next class meeting.

11. Turn in a full lab writeup that includes:

- a brief introduction that sets the context for what we did
- a description of our materials and methods, including anything we did that was different from this handout
- the class data
- a general discussion of the results, and how this simulation would apply to biological evolution

Embarrassing moments at gene parties