Chapter 17

The Evolution of Cooperative Behavior

Alejandro Lynch

Department of Zoology
University of Toronto
Toronto, Ontario M5S 1A1
lynch@zoo.utoronto.ca

Alejandro was born and raised in Buenos Aires, Argentina. He obtained his Bachelor's degree from the Universidad de La Plata, Argentina, and his Ph.D. from the University of Toronto, Canada. He has broad interests in evolutionary biology, in particular in the areas of cultural evolution and population genetics.

Reprinted from: Lynch, A. 1994. The evolution of cooperative behvaior. Pages 319–333, *in* Tested studies for laboratory teaching, Volume 15 (C. A. Goldman, Editor). Proceedings of the 15th Workshop/Conference of the Association for Biology Laboratory Education (ABLE), 390 pages.

Contents

Introduction

This laboratory exercise is a hands-on, interactive introduction to game theory and investigates the possibility of the origin of cooperation among selfish individuals. Students first learn about costs and benefits of behaviours, and then construct a payoff matrix for the Prisoner's Dilemma game. They then pair up with a partner and play single move, multiple move, and evolutionary versions of the Prisoner's Dilemma game.

The exercise is intended for an introductory biology course, although it can be easily adapted for use in higher-level courses. It runs best with a class size of 16 to 24. In larger classes, the data collection and computations required in each generation of the evolutionary game get unwieldy and the pace becomes too slow. The exercise can be completed in 2 hours, with 1 additional hour needed for discussion and viewing of one of the recommended videos (see Resource Material in the Notes for the Instructor). The only materials required are index cards (4 per student) and a calculator.

The success of the exercise depends on the skill of the instructor in being able to move the students through the three games at a pace which keeps them motivated. Since the instructor plays a very active role directing the exercise and encouraging student participation, she or he must be familiar with the flow of the exercise and its content (see suggested reading in the Resource Material section). I have developed a computer simulation program to give the students an opportunity to explore this experiment in more detail (see Computer Simulation in the Notes for the Instructor). The computer simulation game has essentially the same design as the evolutionary game the students play in the laboratory; however, there are 10 different strategies they can choose from, they can vary the initial conditions by specifying the number of players initially playing each strategy, and they can vary the number of generations.

Student Outline

Introduction

An important aspect of animal behavior involves the interactions between individuals, either of the same or of different species. Biologists have observed that in many instances during these interactions individuals help each other at their own expense. This cooperative behavior presents a difficulty for the theory of natural selection: if more successful organisms are the ones that transmit their characteristics to future generations, shouldn't organisms always be selfish and behave in ways that benefit only themselves and not others? One possible way to understand these situations is to realize that the results of the interactions depend on the presence or absence of common interests between the individuals. Two theories have been proposed to explain these observations: kin selection theory and reciprocity theory.

Kin selection theory argues that individuals help close relatives with which they share some of their genes, and therefore by helping them they are helping the survival of those shared genes. In this case the common interest of the interacting individuals is the transmission of their shared genes. However, in many instances organisms help unrelated individuals (sometimes individuals of different species!). **Reciprocity theory** argues that individuals help other non-related individuals in the hope that they will reciprocate. Individuals help each other not because they are being altruistic but because they can obtain benefits that are not obtainable by an individualistic behavior. In such cases, the benefit of helping is greater than the cost, and reciprocity is favoured. But with reciprocity comes a new problem. If individuals are not really interested in the welfare of others, is it possible for some to exploit others by receiving the benefits without reciprocating, thereby avoiding the cost involved in helping others?

In this lab we will be investigating the possibility of cooperative behavior to arise among unrelated individuals even in the presence of selfish individuals or individuals that do not care for the well-being of other organisms with which they interact. Is it possible for individuals to obtain the benefits of mutual cooperation but at the same time protect themselves from selfish individuals?

Game Theory

A useful way to study the interactions between individuals and their choices of behaviors is what is called “game theory,” so named because it originated from the study of the behavior of players in games of strategy, such as poker and bridge. In game theory the interactions between individuals are modelled as if they were a game where the players have a number of possible alternative behaviors, and the outcome of the interaction depends on the behavior chosen by both players.

There are two families of games. Games such as chess and other competitive games in which one player's loss is the other's gain are called **zero-sum games**. In these instances the interests of the individuals are in total conflict; possible biological examples are the interactions between a predator and its potential prey. In these competitive situations we would expect (and indeed we observe) totally selfish behavior.

The second family of games, and the type you will study in this exercise, are games in which the interests of the two individuals partly coincide and are partly opposed and are called **non zero-sum games**. Because there are at least some shared interests between the players, games like these can potentially lead to cooperation. Even such seemingly competitive interactions, such as between a parasite and its host, may have an aspect of cooperative behavior; it may benefit the parasite to maintain a healthy host and thus continue to benefit from it.

Prisoner's Dilemma Game

Many species of birds and mammals engage in mutual grooming, where individuals alternatively remove parasites. There are obvious benefits to be gained from getting someone else to pull off parasites from places that are not easily accessible to oneself, and therefore one can see the benefits of mutual grooming. However, one individual might do even better by not cooperating: if it let other individuals get rid of its own parasites without paying the costs of time and energy in reciprocating.

One game that seems to represent this type of situation in which individuals do not have strictly opposing interests is what is called Prisoner's Dilemma. This non-zero game allows the players to achieve mutual gains from cooperation, but it also allows for the possibility of one player exploiting the other, or the possibility that neither will cooperate.

The name “Prisoner's Dilemma” comes from the following imaginary situation: Two prisoners are in jail in different cells, with no possibility of communication. They have both been charged for a small crime for which they would have to spend 2 years in jail, but the authorities suspect (but have no evidence) that they are also guilty of a major crime for which they could spend 8 years in jail. To extract convictions for the major crime, the authorities offer a reduced sentence to each prisoner (charges will be dropped on the small crime) if they will betray the other prisoner and implicate them in the major crime (**defect**). If, on the other hand, the prisoners **cooperate** with each other and both remain silent, the authorities can only keep them in jail for 2 years.

Imagine you are Prisoner 1. You do not know what your partner in crime is going to do, so what should you do? The consequences or payoffs to you of your choice and your partner's choice can be summarised in a matrix as in Table 17.1.

Table 17.1. Payoff matrix for Prisoner 1.

		2	
		Cooperate	Defect
1	Cooperate	2 years	10 years
	Defect	Free (0 years)	8 years

The matrix in Table 17.1 is constructed from the point of view of *your* options. For instance, if you decide to remain silent (cooperating with your partner), but your partner defects, you have been suckered and you get jailed for 10 years while your partner goes free. If you both remain silent, both of you will only get the 2-year sentence. If you both tell on each other (defect), you both get jailed for 8 years (both collaborate with authorities and so do not get charged for the small crime). What should your choice be? If your partner cooperates by remaining silent, then your best move is to defect and go free. Alternatively, if your partner defects then your best move is to also defect, and you remain jailed for 8 years. And herein lies the dilemma: no matter what your partner chooses, you would do better by defecting, and the same is true for your partner. But if you both betray each other, you are both worse off than if you had cooperated and remained silent.

Costs and Benefits Payoff Matrix

In studies of animal behavior, rather than prison sentences, the payoff matrix is based on the costs of performing a behavior (c) and the benefits that behavior derives (b). The costs could be measured in energetic terms, or in terms of reduced reproduction or survivorship, etc. Similarly the benefits could be measured in terms of number of offspring, energy gain, increased probability of survival, economic gain, etc. For a payoff matrix for our Prisoner's Dilemma game, the exact numbers are not important, just so long as $b > c$. Thus, the payoff matrix for a Prisoner's Dilemma game can be given (as above, in terms of Player A) as in Table 17.2 If both players cooperate with each other, player A gets the benefits of cooperation but has also paid the cost of reciprocation ($b - c$); if A cooperates but B defects, player A paid the costs without deriving any benefits ($-c$); if A defects and B cooperates, player A gets the benefits without paying any costs (b); and if both defect, there are no costs of benefits (0).

Table 17.2. Payoff matrix for Player A.

		B	
		Cooperate	Defect
A	Cooperate	$b - c$	$- c$
	Defect	b	0

Laboratory Exercise

To familiarise yourselves with some game theory and its applications to the study of animal behavior, you will become the experimental animals! First, the payoff matrix for a Prisoner's Dilemma game must be constructed. In this exercise, the costs and benefits can be thought of as points which each player can amass.

The Payoff Matrix

1. Construct a **raw** payoff matrix (Table 17.3), where the benefits = 4 points and costs = 2 points.

Table 17.3. Raw payoff matrix.

		Partner's move	
		Cooperate	Defect
My move	Cooperate		
	Defect		

2. Now construct a **normalized** payoff matrix (Table 17.4) which makes it easy to tally points in the game. To do this, add *c* to each cell so that all payoffs are positive or zero.

Table 17.4. Normalized payoff matrix.

		Partner's move	
		Cooperate	Defect
My move	Cooperate		
	Defect		

Single-Move Game

1. Pair with one other student at random. Each of you will play with two cards, one marked C (for cooperate) and the other marked D (for defect).
2. Decide what card you will play, but do not give your partner any clue of your intended move.
3. When the TA says “ready” put your chosen card with the letter-side down on the desk.
4. When the TA says “go” turn your card over.
5. Record your move, your partner's move, and the points each of you made from the play in the table provided.
6. Class data will be tabulated on the board. Copy the results into Table 17.5. As a class, you will tally up the points and calculate the average points for cooperators, and the average for defectors.

   ☐ In a *single* interaction such as this, what is the best move in a Prisoner's Dilemma game?

Table 17.5. Summary of class data for single-move game.

Class data
CC =
CD =
DD =

	Number (%)	Average points
C		
D		

Multiple-Moves Game

Another version of this game is the Iterated Prisoner's Dilemma, where the probability of reencountering the same individual is high. You both know that it is to your advantage to collaborate, but there is always the danger of your partner “cheating” and defecting. Can you build up trust and elicit cooperation while preventing cheating? On the other hand, are there subtle ways of cheating while still managing to elicit cooperation?

You are going to play the same game, but this time an indeterminate number of times. Your objective is to obtain as many points as possible. You should keep in mind that what matters is not so much to make more points than the person you are playing with but to try to do as well as possible compared to the entire class. While you play try to think of possible game plans or strategies.

1. Pair randomly again. You will be playing the same game, this time between 20 to 30 times with the same person.
2. Follow the same procedure as before for each move: When the TA says “ready” put your chosen card on the table, and when the TA says “go” turn it over.
3. After each play, record your move and your partner's and the number of points you obtained in Table 17.6.
4. The TA will stop the game after somewhere between 20 and 30 times.

   ☐ Why might it be important that you do not know when the game is going to finish?

5. Tally up your points. The TA will tabulate each player's points on the board. Now you can compare how you did relative to others in the class.

   ☐ Did the players who scored best employ a game strategy?

   ☐ In retrospect, having played the game, can you come up with some game strategies which would increase your score?

Table 17.6. Data sheet for multiple-moves game.

	My move	Their move	My points
1			
2			
3			
4			
5			
6			
7			
8			
9			
10			
11			
12			
13			
14			
15			

	My move	Their move	My points
16			
17			
18			
19			
20			
21			
22			
23			
24			
25			
26			
27			
28			
29			
30			

Evolutionary Game

Now suppose that within any population, individuals adopt one of a number of possible game plans or strategies. Suppose that strategies are transmitted between individuals across generations, either by genetic transmission or cultural transmission (teaching and learning). The more successful strategies have higher probabilities of being transmitted. We are interested in seeing if there are some strategies that are more successful than others, and if this is so, what is it that makes them so. The rationale is that strategies that do well are more likely to persist in a given population and strategies that are not very successful will tend to disappear. This could be achieved in a number of different ways. For example, individuals copy the strategy of successful individuals, or successful individuals leave more offspring and transmit the strategy to them (either by genetic or cultural means).

1st Generation

1. You will be given a card with the name and a brief description of one of four possible strategies. All strategies will be represented by about the same number of players (for example, if there are 20 students in the class, each strategy will be played by 5 students).
2. Play the game 20 times with your partner using your assigned strategy. As before, record all moves as well as the number of points that you make at each move. Record your moves in Table 17.7.

☐ Why this time does it not matter that you know the number of times you will play?

3. When you finish, tally up your points and give them to the TA who will then tally up the total number of points for each strategy.

Calculate Strategies for the Next Generation

The number of strategies playing in the 2nd generation is based on the points each strategy made in the 1st generation. As an example, suppose that the total points amassed by the players using the "cut-and-run" strategy was 410 and together with the three other strategies there was a class total of 820 points. The relative success of the cut-and-run strategy is expressed as a frequency; that is, 410/820 = 0.50. So, one-half of the players in the 2nd round will be using the "cut-and-run" strategy.

2nd Generation and Onward

1. Strategies are now reassigned to players, based on their success in the previous generation, as described above.
2. Players pair up with a new partner.
3. Play 20 times with your new partner using your strategy (which may be a different one from the previous generation).
4. Tally up points as before, and the TA will calculate the success of each strategy after two generations.
5. This procedure will be repeated for a total of four generations and you should record the number of individuals playing each strategy in each generation in Table 17.8.

Table 17.8. Changes in frequency through time for the four strategies.

	Generation				
Strategy	0	1	2	3	4

Table 17.7. Data table for evolutionary game.

	1st generation			2nd generation			3rd generation			4th generation		
	My move	Their move	My pts	My move	Their move	My pts	My move	Their move	My pts	My move	Their move	My pts
1												
2												
3												
4												
5												
6												
7												
8												
9												
10												
11												
12												
13												
14												
15												
16												
17												
18												
19												
20												

Graph the Results

Plot the frequency of each strategy for each generation on a graph. Remember to include the initial frequencies as well as results for the four subsequent generations. This graph will track the evolution of each strategy in the population over four generations.

Review Questions

1. What is the difference between a zero-sum and a non-zero sum game?
2. Did any one of the four strategies employed come out a clear winner in the lab exercise?
3. What characteristics of the best strategies might enable them to persist in a population?
4. Can you identify some characteristics in common between the worst strategies?
5. Can you think of another strategy that could have done better?
6. Suggest some human behaviors which could be modelled by the Prisoner's Dilemma game.
7. Suggest some non-human behaviors which could be modelled by the Prisoner's Dilemma game.

Notes for the Instructor

Suggested Protocol

1. Give a short introduction as in the Student Outline. You might want to prompt your students with questions about kin selection and reciprocity theories. Fill up the payoff matrices in Tables 17.3 and 17.4 with the students' input.
2. Let the students pair at random. If there is an odd number of students the instructor will have to pair with one of them.
3. Collect results and complete Table 17.5.
4. Explain why in a single-move Prisoner's Dilemma the rational choice (i.e., safest, in the sense that maximizes potential payoffs) is to play D; no matter what the other person does, you can always get more points by playing D.
5. Introduce the iterated Prisoner's Dilemma.
6. Let the students pair at random again. Play the game somewhere between 20 and 30 times. The reason for not letting the students know how many times they are going to play is because if they knew they could play D at the end of the game.
7. Collect the results and write on the board the points obtained by each student. Discuss how did they decide their moves; did anyone have a well defined strategy? Were high-score students trying to get the other player to play C? What about low-scoring students?
8. Now you have to choose four strategies for the next part of the exercise. Get the students involved and try them to come up with well defined strategies. A list of the strategies that are available in the computer simulation exercise are given below. It is perhaps a good idea to try to get two “nice” strategies (strategies that never play D unless provoked by a previous D: tit for tat, tit for two tats, two tits for tat, always collaborate, grudger) and two not “nice” (random, sneaker, always defect, big memory, alternate). Explain the rationale for an evolutionary game: more successful strategies are likely to become more common as time goes by. Different mechanisms could be involved in this; it could be achieved by better survival and reproduction of the individuals playing more successful strategies, greater chance of being imitated, etc.

9. At the end of each generation collect the results and determine the frequency of players for each strategy for the next generation (see Appendix A for calculations). Add and delete strategy cards as needed and redistribute them among the students.

10. Plot the results for the four generations.

11. Discuss the results with the students. Have them answer the review questions as a class discussion.

In addition to this guide to the steps involved in conducting the lab, there is a flow chart in Figure 17.1 which shows the main steps in a condensed form. In the flow chart, the left column indicates Teaching Assistant input (by means of explanations, promoting discussion, etc.), the middle column shows the activities that the students perform, and the right column indicates data that the students will collect. The numbers given on this last column are real numbers obtained during a pilot test of the lab at the University of Toronto in April 1992.

There are several important points we hope to transmit to the students in this lab:

1. First of all, we hope to have shown a new way of looking at interactions between individuals, humans or other animals. Game theory can potentially teach us a lot about the ways that different behaviors can arise and exist in a population. The theory has wide applications not only in evolutionary biology, but also in psychology, economics, political, and social sciences.

2. Second, we show that it is sometimes possible to obtain cooperative behavior even in the presence of selfish individuals or individuals who do not care for the well being of others with whom they interact.

   Interactions where individuals can potentially benefit from mutual cooperation are widespread, both among members of the same species and also among individuals of different species. Relationships of mutual benefit between members of different species are called *symbiosis* or *mutualism.* These are of special interest to us because there is clearly no genetic relatedness involved. Some examples of these are birds that clean the mouths of crocodiles, or the close relationship between some fungi and algae that form lichens. It has been suggested that some of the organelles of multicellular organisms, such as the mitochondria, are the result of a symbiotic interaction between the two different types of bacteria.

3. The most successful strategies in evolutionary games such as the iterated Prisoner's Dilemma share the following properties:

   • Niceness: Here defined as not being the first to defect. It pays to cooperate as long as the other player is cooperating.

   • Provocability: Retaliate immediately after a defection by the other player.

   • Forgiveness: Don't hold a grudge once you've vented your anger. Show the other player that you will not be a sucker, but that you want to cooperate.

   • Not envious: In a non-zero sum game you don't have to beat the other player to do well yourself, especially if you are interacting with many different players. The success of the other player is almost a prerequisite for your doing well yourself. Note that TIT FOR TAT can never make more points than the other player; the best it can do is to make as well as the other, and often it loses, though not by much.

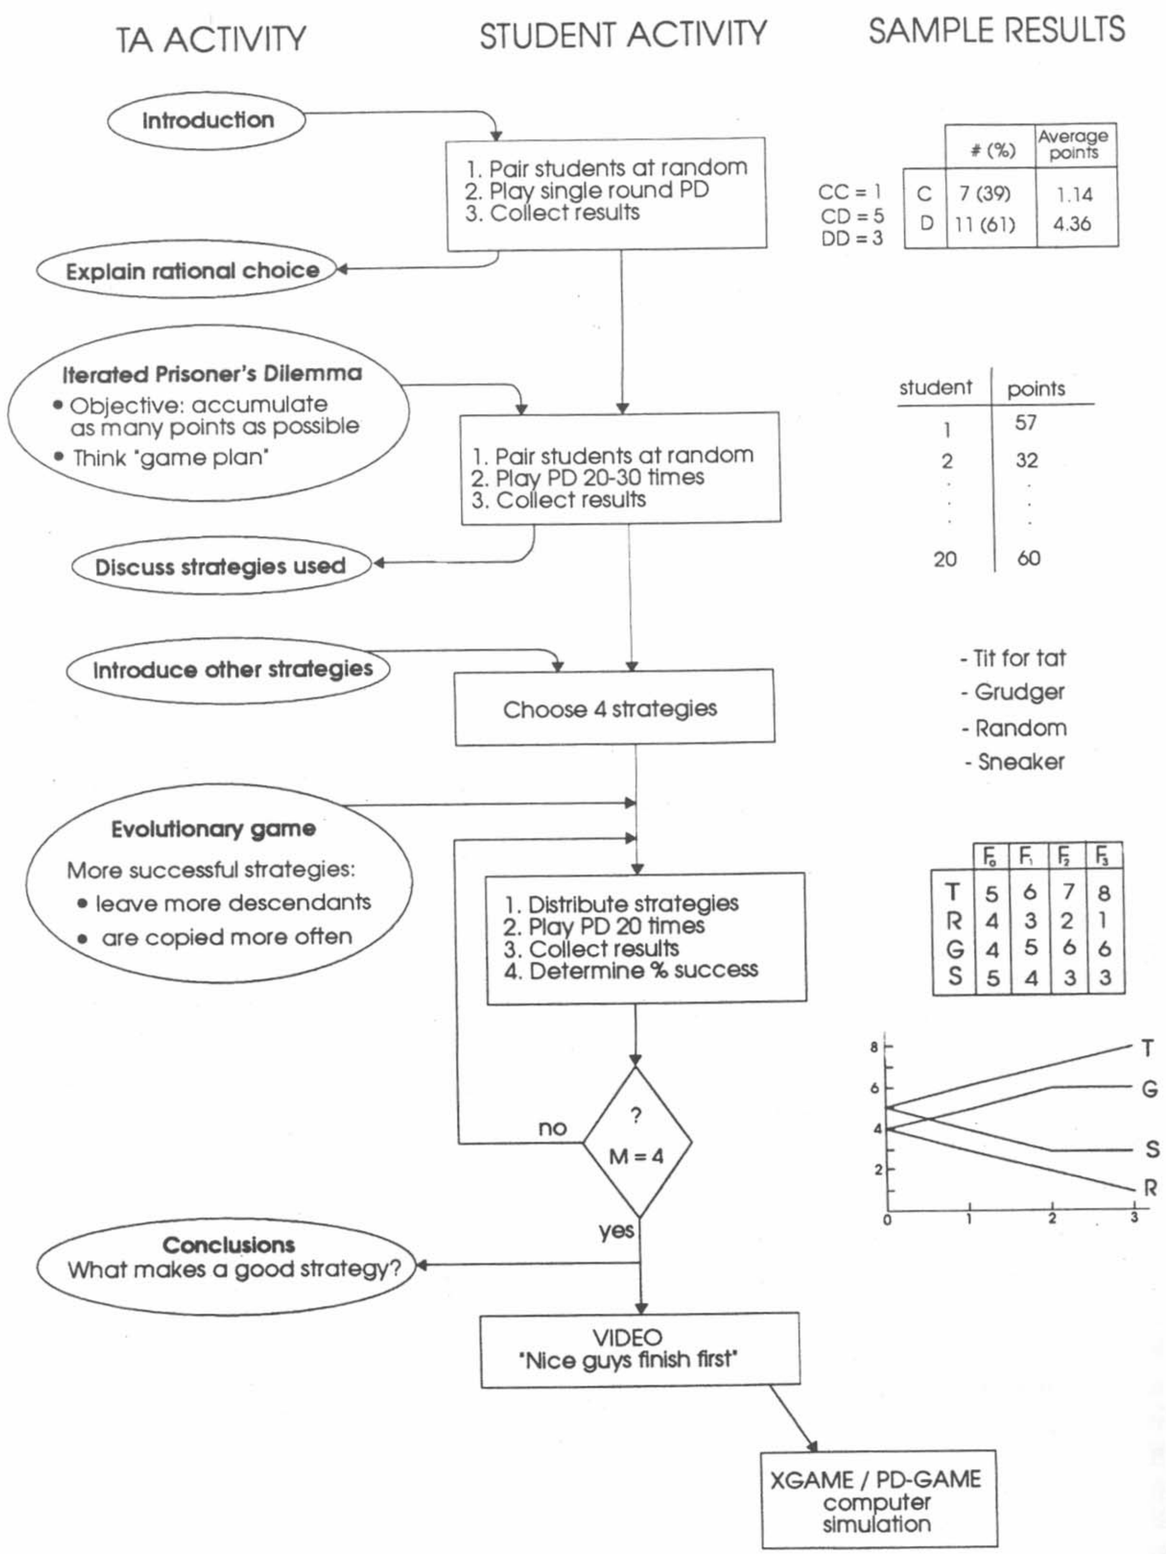

Figure 17.1. Flow chart for The Evolution of Cooperative Behavior laboratory exercise.

Resource Material

The following articles and videos contain additional background information that should be useful to instructors. One of the videos can profitably be incorporated into the lab and shown to the students at the end of the lab.

Axelrod, R. 1984. The evolution of cooperation. Basic Books, 241 pages. [ISBN 0-465-02121-2] (Describes the now classic computer tournaments that highlighted the efficacy and robustness of reciprocity in the evolution of cooperation. Chapter 5 is a slightly modified version of the award-winning paper by R. Axelrod and W. D. Hamilton (1981), *Science*, 211:1390–1396.)

Dawkins, Richard. 1989. Nice guys finish first. Chapter 12. Pages 202–233, *in* The selfish gene (New edition). Oxford University Press, 352 pages. (If you can put up with Dawkins' arrogance, this is a very good introduction to the topic from a biological point of view.)

Hofstadter, D. R. 1983. Computer tournaments of the Prisoner's Dilemma suggest how cooperation evolves. *Scientific American*, 248 (5):16–26. (An excellent introduction to the topic from a more human-oriented perspective. Reprinted as Chapter 29 in Hofstadter's book *Methamagical Themas*.)

Suzuki, David. 1990. *The Nature of Things: Little Wars.* An educational program on videocassette produced by the Canadian Broadcasting Corporation (Educational Sales, Box 500, Station A, Toronto, Ont. M5W 1E6, 416/975-6384). (Hosted by David Suzuki, with an interview of Anatole Rapoport, late professor at the University of Toronto, world expert of the Prisoner's Dilemma, and the person that submitted the winning Tit For Tat strategy to the two computer tournaments organized by Axelrod.)

Taylor, J. 1985. *Nice Guys Finish First*. A 50-minute educational program on videocassette produced by the British Broadcasting Corporation (Suite 111, 65 Heward Ave., Toronto, Ont. M4M 2T5, 416/469-1505). (Features Richard Dawkins presenting basically the same material as in his Chapter 12 of the New Edition of *The Selfish Gene*).

Computer Simulation

The class experiment is very simplified, and the instructor should stress that the results depend very much in the suite of strategies in the population and their respective frequencies. To give the students an opportunity to explore this experiment in more detail, I have developed a computer simulation for the IBM-compatible PC where students can expand the range of possibilities that we offered them during this lab. The computer simulation game has essentially the same design as the evolutionary game the students will play in the lab. However, there are 10 different strategies they can choose from (given in Table 17.8), they can vary the initial conditions by specifying the number of players initially playing each strategy (0–20; there has to be at least two players), and they can also vary the number of generations (1–100). Within a generation each player plays a variable number of times against two other randomly chosen individuals. The number of interactions between two players can also be varied (1–50). For more information on this program contact the author or the Intro Biology Programme, University of Toronto, 25 Harbord St. (Room 019), Toronto, Ont. M5S 1A1.

Table 7.8. Strategies available in computer simulation.

Alternate	The player alternates between C and D, starting with a C.
Always Cooperate	The player alternates between C and D, starting with a C.
Always Defect	Players using this strategy are always uncooperative, playing D in every move.
Big Memory	The player chooses C with probability *p*, where *p* is determined by the frequency of Cs in previous interactions with that particular player. The first move is either C or D with equal probability.
Grudger	Players using this strategy start playing C and continue to do so until the other player plays D. After that it plays D for the rest of the interactions with that particular player.
Random	The player chooses either C or D with equal probability.
Sneaker	The player starts with C and then plays whatever the other player did in the previous move, but tries to exploit the other player by playing D at random intervals with a probability of 0.3.
Tit for Tat	The player starts playing C and then plays whatever the other player did in the previous move.
Tit for Two Tats	The player plays C in the first and second moves. After that, if the other player played D in the two previous moves then plays D, otherwise he plays C.
Two Tits for Tat	The player starts with C, and then if the other player plays D, then plays D in the two following moves, otherwise plays C.

Acknowledgements

I thank Mike Dennison for his input and encouragement during all phases of the development of this lab exercise.

APPENDIX A

Calculating the Frequency of Players for Each Strategy

To calculate the number of players using a given strategy in a generation we use the relative success of the different strategies in the previous generation as follows:

1. Calculate totals for each strategy in generation i - 1. For example:

    TIT FOR TAT = 352 points
    GRUDGER = 260
    RANDOM = 212
    ALWAYS DEFECT = 184

2. Calculate grand total points = 1008 (352 + 260 + 212 + 184).

3. Calculate the relative success of each strategy at the end of generation i - 1 (total points/grand total).

    TIT FOR TAT = 352/1008 = 0.35
    GRUDGER = 260/1008 = 0.26
    RANDOM = 212/1008 = 0.21
    ALWAYS DEFECT = 184/1008 = 0.18

4. Players in generation i = relative success (in generation i) × total number of players.

    TIT FOR TAT = 0.35 × 18 = 6.3
    GRUDGER = 0.26 × 18 = 4.6
    RANDOM = 0.21 × 18 = 3.8
    ALL DEFECT = 0.18 × 18 = 3.24

5. Rounding up and down, gives us the number of players for generation i.

    TIT FOR TAT = 6 players
    GRUDGER = 5 players
    RANDOM = 4 players
    ALL DEFECT = 3 players