An Investigation of *N*-person Prisoners' Dilemmas

Miklos N. Szilagyi
Department of Electrical and Computer Engineering,
University of Arizona, Tucson, Arizona 85721

This paper is an attempt to systematically present the problem of various *N*-person Prisoners' Dilemma games and some of their possible solutions. 13 characteristics of the game are discussed. The role of payoff curves, personalities, and neighborhood is investigated. Experiments are performed with a new simulation tool for multi-agent Prisoners' Dilemma games. Investigations of realistic situations in which agents have various personalities show interesting new results. For the case of pavlovian agents the game has nontrivial but remarkably regular solutions. Examples of nonuniform distributions and mixed personalities are also presented. All solutions strongly depend on the choice of parameter values.

1. Introduction

The Prisoners' Dilemma is usually defined between two players [1] and within Game Theory which assumes that the players act rationally. Realistic investigations of collective behavior, however, require a multi-person model of the game [2] that serves as a mathematical formulation of what is wrong with human society [3]. This topic has great practical importance because its study may lead to a better understanding of the factors stimulating or inhibiting cooperative behavior within social systems. It recapitulates characteristics fundamental to almost every social intercourse.

Various aspects of the multi-person Prisoners' Dilemma have been investigated in the literature but there is still no consensus about its real meaning [4–26].

Thousands of papers have been published about the two-agent iterated Prisoners' Dilemma game, a few of these are [27–30]. The interest in investigating various strategies for pair-wise interactions in multi-agent Prisoners' Dilemma computer tournaments is amazing because, as Rapoport rightly noted in [31], these "tournaments demonstrated neither evolution *nor* learning because nothing evolved and nothing was learned" in the succession of two-person games. Nevertheless, the obsession with these tournaments continues [32]. Even papers that claim the simulation of multi-agent games are usually based on dyadic interactions between the agents. A stochastic learning model was devel-

Complex Systems, **14** (2003) 155–174; © 2003 Complex Systems Publications, Inc.

https://doi.org/10.25088/ComplexSystems.14.2.155

oped by Macy in [33] to explain critical states where threshold effects may cause shifting the system of agents from a defective equilibrium to a cooperative one. Simple computer programs are written in [34, 35] that demonstrate the dynamics of deterministic social behavior based on pair-wise interactions between the participants. Akimov and Soutchansky [36] presented a multi-agent simulation (not a succession of two-person games) but their experiment was limited to six agents. Our own simulation tool [37] was designed to simulate social dilemmas with a wide range of user-defined parameters. It is suitable for an unlimited number of agents with various personalities. We were able to perform interesting nontrivial experiments with this tool [24, 26].

This paper is an attempt to systematically present the problem of the N-person Prisoners' Dilemma and some of its possible solutions.

2. N-person dilemmas

The N-person Prisoners' Dilemma considers a situation when each of N participants has a choice between two actions: cooperating with each other for the "common good" or defecting (following their selfish short-term interests). As a result of its choice, each participant receives a reward or punishment (payoff) that is dependent on its choice as well as everybody else's (Figure 1). The participants can be individuals, collectives of persons, organizations, or even computer programs. We will simply call them *agents*.

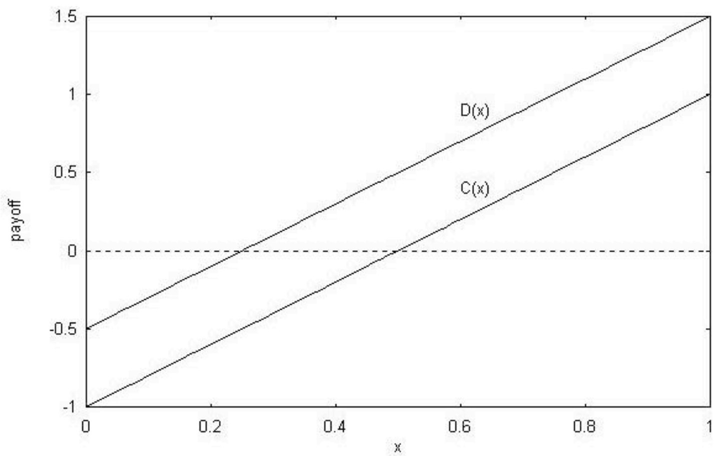

Figure 1. Reward/penalty functions for defectors (D) and cooperators (C). The horizontal axis (x) represents the ratio of the number of cooperators to the total number of neighbors; the vertical axis is the reward/penalty provided by the environment. In this figure, $D(x) = -0.5 + 2x$ and $C(x) = -1 + 2x$.

The dilemma can be formulated by the following two statements [11].

1. Regardless of what the other agents do, each agent receives a higher payoff for defecting behavior than for cooperating behavior.

2. All agents receive a lower payoff if all defect than if all cooperate.

If m of the N agents are cooperating and $C(m)$ and $D(m)$ are the payoffs to a cooperator and a defector, respectively, then the above conditions can be expressed as

$$D(m) > C(m+1) \tag{1}$$

and

$$C(N) > D(0). \tag{2}$$

$C(0)$ and $D(N)$ are undefined; therefore, the value of m is between 0 and $N-1$ in equation (1). The game has $N+1$ distinguishable outcomes: $0, 1, 2, \ldots, N-1, N$ participants may choose cooperation.

At a first glance, it looks like a well-defined problem. However, at least the following questions arise immediately.

1. Are the choices and actions of the agents simultaneous or distributed in time?

2. Can individual agents see and adapt to the actions of others?

3. Can they form coalitions?

4. What are the agents' goals in the game: to maximize their payoffs, to win a competition, to do better than their neighbors, to behave like the majority, or any other goal?

5. Is it a one-shot game or an iterated one? If it is an iterated game, how will the next action be determined?

6. What are the personalities of the agents? (Surely, different people react quite differently to the same conditions.) Can they change their personalities?

7. Can an agent refuse participation in the game?

8. Are the payoff curves the same for all agents?

9. What are the payoff curves?

10. How is the total payoff to all agents related to the number of cooperators?

11. How are the agents distributed in space? Can they move?

12. Do the agents interact with everyone else or just with their neighbors?

13. How is the neighborhood defined?

https://doi.org/10.25088/ComplexSystems.14.2.155

With so many open questions it is obviously quite difficult to create a general classification scheme for the N-person Prisoners' Dilemma and there is a great variety of possible games. It is, in fact, a whole family of quite different games. Even in the case of a *uniform* game the number of possible variations is infinitely large because of the infinite variety of the payoff curves. In a uniform game the payoff curves are the same for all agents, they are monotonically increasing functions of the number of cooperators, and there is some minimum number of cooperators that can gain by their cooperative choice [2]. It is, however, desirable to investigate at least the most characteristic cases because each possible variation may represent an important social situation. Let us first take a closer look at each of the 13 questions.

1. There is a huge difference between simultaneous actions and those distributed in time. In the first case all agents see the same environment at the moment of their simultaneous action. In most social settings, however, agents act at different and noncorrelated times. Therefore, each agent sees a slightly different world than another agent acting at a slightly different time [17]. Simulation of this case is a more sophisticated task than that of the previous case.

2. Even if the agents' actions are distributed in time, they may or may not have information about the actions of others. You may look out of the window and see how many cars are on the road before deciding if you are going to drive your car or take a bus but you do not know how many children will be born next year before deciding if you are going to have another child.

3. Obviously, if you do not know the other participating agents, you cannot form coalitions with them. Even if you know all of them, it is not certain that you can communicate with them, let alone form a coalition. However, coalitions may drastically change the outcome of the game.

4. The goals of the agents in the game is a critical issue. The game is totally different if the goals are different and in real-life situations different agents do have different goals. It is also possible that the agents simply react to their own and to their neighbors' payoffs without specific goals.

5. The one-shot game is less interesting than an iterated one where the agents act repeatedly on the basis of their personalities, their neighbors' situations, and the payoffs received for their previous actions. The next choices are determined by updating schemes that are different for different agents.

6. The personalities of the agents is one of the most important characteristics of the game. The psychological literature on the impact of personalities in social dilemmas is summarized in [19]. It is possible but not easy to quantify personality profiles in the traditional psychological sense. We will use the term "personality" in the sense of decision heuristics (repeated-game strategies) in this work, to represent the fact that different

agents react differently to the same stimulus from their environment. This is a rather primitive approach but it is still much better than the unjustified assumption of uniform response.

Personalities are usually neglected in the literature. In [24] N-person Prisoner's Dilemmas are considered with various personalities of the participating agents. Different agents may have quite different personalities in the same experiment. The agents' personalities may also change in time based on the influences of other agents.

Personalities of the agents may represent genetic as well as cultural differences between them. The simplest and probably most important personality profiles are the following.

- *Pavlovian*. The probability of taking a certain action p changes by an amount proportional to its reward or penalty from the environment. This personality is based on Pavlov's experiments and Thorndike's law [38]: If an action is followed by a satisfactory state of affairs, then the tendency of the agent to produce that particular action is reinforced.
- *Stochastically predictable*. p is a constant. Such an agent is not influenced by the environment at all. Here are the special cases.
  (a) *Negatively stubborn*. Never takes this action ($p = 0$).
  (b) *Positively stubborn*. Always takes this action ($p = 1$).
  (c) *Unpredictable*. Acts randomly ($p = 0.5$).
- *Accountant*. p depends on the average reward for previous actions.
- *Conformist*. Imitates the action of the majority of its neighbors.
- *Greedy*. Imitates the neighbor with the highest reward.

Other profiles may include properties like aggression, sensitivity, devotion, and so forth.

7. The iterated game may considerably change if an agent refuses participation in some iterations.

8. It is usually assumed that the game is uniform; therefore, the payoff curves are the same for all agents. This condition is, however, not always guaranteed.

9. When everything else is fixed, the payoff curves determine the game. There is an infinite variety of payoff curves. In addition, stochastic factors can be specified to represent stochastic responses from the environment. Zero stochastic factors mean a deterministic environment. Even in the almost trivial case when both payoff curves are straight lines and the stochastic factors are both zero, four parameters specify the environment. Attempts to describe it with a single variable are certainly too simplistic [34, 39]. As shown later, the relative position of the two payoff curves with respect to each other does not always determine the outcome of the game. Ordinal preference is not enough to represent the payoff functions:

https://doi.org/10.25088/ComplexSystems.14.2.155

the actual amounts of reward and punishment may be as important as the relative situation of the two curves.

The N-person game is a *compound* game (it can be reduced to a series of two-person games) if and only if both payoff functions are linear [8]. Therefore, a dyadic tournament where every agent plays two-person games against each of the $N-1$ other agents represents only a very limited subset of the N-person game.

10. The total payoff to all agents is related to the number of cooperators but the maximum collective payoff is usually not at maximum cooperation.

11. The agents may be distributed in space in many different ways. If there are fewer agents than locations in space or if more than one agent may occupy one location, then it is also possible that the agents move around in space and their neighborhood constantly changes.

12. The agents may interact with everyone else or just with their neighbors. In the latter case they behave like cellular automata (CA) [40].

13. The number of neighborhood layers around each agent and the agent's location determine the number of its neighbors. The depth of agent A's neighborhood is defined as the maximum distance, in three orthogonal directions, that agent B can be from agent A and still be in its neighborhood. An agent at the edge or in the corner of the available space has fewer neighbors than one in the middle. The neighborhood may extend to the entire array of agents.

To make our task manageable, in the following we assume that the game is uniform and iterated, the agents are distributed in and fully occupy a finite two-dimensional space, the updates are simultaneous, the agents have no goals, know nothing about each other, and they cannot refuse participation in any iteration. This restriction leaves the problem of payoff curves, personalities, and neighborhood open for investigation. These factors are mostly neglected in the literature. The role of the parameters fixed in this study will be addressed in a future work by the author.

We will use computer simulation to demonstrate the role of the above mentioned factors on the outcome of the game. If the parameters are selected appropriately, the simulation will exhibit behavior that is close enough to the behavior of real people when they are placed in a similar situation. It should be noted that even if only three factors are considered there is a huge number of different variations. Therefore, we can only show some characteristic examples in this paper.

3. The model

We have developed an agent-based model for the investigation of social dilemmas with a large number of decision-makers operating in a stochastic environment [37]. Our model has the following three distinctive new features.

1. It is a genuine multi-agent model, not a repeated two-person game.

2. It is a general framework for inquiry in which the properties of the environment, as well as those of the agents, are user-defined parameters. The number of interacting agents is theoretically unlimited.

3. Our agents have various distinct, user-defined personalities.

The participating agents are described as stochastic learning CA, that is, as combinations of CA [40, 41] and stochastic learning automata [42, 43]. The CA format describes the environment in which the agents interact. In our model, this environment is not limited to the agents' immediate neighbors: the agents may interact with all other agents simultaneously. Stochastic learning rules provide more powerful and realistic results than the deterministic rules usually used in CA. *Stochastic learning* means that behavior is not determined but only shaped by its consequences, that is, an action of the agent will be more probable but still not certain after a favorable response from the environment. The model is described in detail in [37] and only its most important features are briefly explained here.

A realistic simulation model of a multi-person game must include parameters such as the size and shape of the simulation environment, the payoff functions, updating schemes for subsequent actions, personalities of the agents, and the definition of the neighborhood.

Our simulation environment is a two-dimensional array of the participating agents. Its size is limited only by the computer's virtual memory. The behavior of a few million interacting agents can easily be observed on a computer screen.

The payoff (reward/penalty) functions are given as two curves: one (C) for a cooperator and another (D) for a defector. The payoff to each agent depends on its choice, on the distribution of other players among cooperators and defectors, and also on the properties of the environment. The payoff curves are functions of the ratio of cooperators to the total number of neighbors (Figure 1). The freedom of using arbitrary functions for the determination of the reward/penalty system makes it possible to simulate a wide range of dilemmas and other social situations, including those where the two curves intersect.

In an iterative game the aggregate cooperation proportion changes in time, that is, over subsequent iterations. The agents take actions according to probabilities updated on the basis of the reward/penalty received for their previous actions and of their personalities. The updating scheme may be different for different agents. This means that agents with completely different personalities can be allowed to interact with each other in the same experiment. Agents with various personalities

https://doi.org/10.25088/ComplexSystems.14.2.155

and various initial states and actions can be placed anywhere in the array. The response of the environment is influenced by the actions of all participating agents.

The updated probabilities lead to new decisions by the agents that are rewarded/penalized by the environment. With each iteration, the software tool draws the array of agents in a window on the computer screen, with each agent in the array colored according to its most recent action. The experimenter can view and record the evolution of the society of agents as it changes in time. The outcome of the game depends on the personalities of the agents. For example, agents with short-term rationality will always choose defection, benevolent agents will ignore their short-term interests and will all cooperate, and so forth.

4. Experiments

4.1 Pavlovian agents

It is realistic and interesting to consider pavlovian agents first. Their response is stochastic but their probability of cooperation p changes by an amount proportional to their reward/punishment from the environment (the coefficient of proportionality is called the *learning rate*). These agents are primitive enough not to know anything about their rational choices but they have enough "intelligence" to learn a behavior according to Thorndike's law. Such agents have been used in [29, 44, 45] and others for the investigation of iterated two-person games.

A linear updating scheme is used for these agents: the change in the probability of choosing the previously chosen action again is proportional to the reward/penalty received from the environment (payoff curves). Of course, the probabilities always remain in the interval between 0 and 1.

Let us assume that in a society of N pavlovian agents the ratio of cooperators is $x = m/N$ and the ratio of defectors is $(1 - x)$ at a certain time. We have shown [26] that when the cooperators receive the same total payoff as the defectors, that is,

$$x\, C(x) = (1 - x)D(x), \tag{3}$$

an equilibrium occurs. This may happen if $C(x)$ and $D(x)$ are either both negative or both positive. In the first case, a small number of cooperators are punished severely and a large number of defectors are punished little. This leads to a stable equilibrium. In the second case, a large number of cooperators are rewarded slightly and a small number of defectors are rewarded greatly. This point corresponds to an unstable equilibrium.

In the case of linear payoff functions the equilibrium equation is quadratic. For the payoff functions of Figure 1 the solutions are

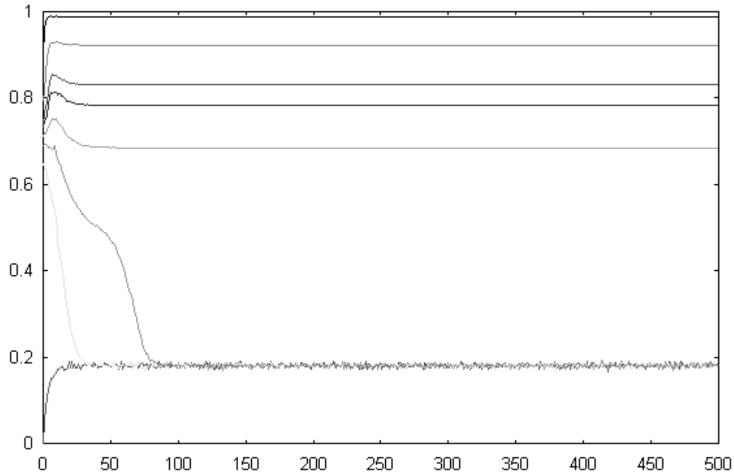

Figure 2. Evolution of the game for the case when all agents are pavlovian, the payoff curves are given by Figure 1, and the neighborhood is the entire collective of agents. The graphs show the proportion of cooperating agents as a function of the number of iterations. The initial cooperation ratios from top to bottom curves are 0.90, 0.80, 0.75, 0.73, 0.71, 0.69, 0.65, and 0.00.

$x_1 = 0.180$ (stable attractor) and $x_2 = 0.695$ (unstable repulsor). Simulation results are shown in Figures 2 and 3. The graphs represent the proportion of cooperating agents as a function of the number of iterations for different initial cooperation ratios.

Figure 2 refers to the case when the neighborhood is the entire collective of agents. When the initial cooperation ratio is below x_2, the solution of the game converges toward x_1 as an oscillation while it stabilizes exactly when the initial cooperation ratio is above x_2. The latter case does not result in the aggregate cooperation proportion converging to 1, as might be expected. This is because, for an individual agent that started off as a defector, there is always some likelihood that the agent will continue to defect. This probability is initially small but continues to increase because the agent is always rewarded for defecting. If the number of agents is sufficiently large, then there will be some agents that continue to defect until their cooperation probability reaches zero due to the successive rewards they have received, and these agents will defect forever.

The situation is different when the neighborhood is only one layer deep. In this case each agent has a maximum of eight neighbors whose behavior can influence its reward/penalty. Accordingly, the result is a more gradual dependence on the initial convergence ratio (Figure 3).

Naturally, the results are strongly dependent on the payoff functions. In the case of pavlovian agents the relative situation of the two payoff

https://doi.org/10.25088/ComplexSystems.14.2.155

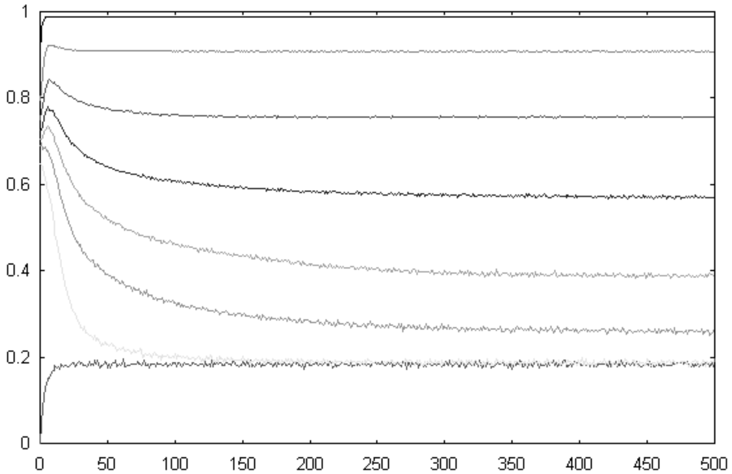

Figure 3. Evolution of the game for the case when all agents are pavlovian, the payoff curves are given by Figure 1, and the neighborhood is one layer deep. The graphs show the proportion of cooperating agents as a function of the number of iterations. The initial cooperation ratios from top to bottom curves are 0.90, 0.80, 0.75, 0.73, 0.71, 0.69, 0.65, and 0.00.

curves with respect to each other does not determine the outcome of the game. It is equally important to know the actual values of the payoff. For example, consider the simple payoff functions shown in Figure 1. If we shift the horizontal axis up and down, the following cases are possible.

(a) Both curves are positive for any value of x. In this case only the unstable equilibrium is possible and the solution of the game depends on the value of this equilibrium and on the initial ratio of cooperators. When the initial cooperation ratio is below x_2, the solution of the game stabilizes at a lower value between 0 and x_2. When the initial cooperation ratio is above x_2, the final stable ratio has a higher value between x_2 and 1.

(b) The $D(x)$ curve is entirely positive but $C(x)$ changes sign from negative to positive as the value of x grows. The situation is similar to case (a). The only difference is that in this case the region where both $C(x)$ and $D(x)$ are positive may be too narrow to produce a solution other than total defection.

(c) The most interesting case is when both $C(x)$ and $D(x)$ change sign. In this case both equilibria exist and we have the solutions discussed above (Figures 2 and 3).

(d) The $C(x)$ curve is entirely negative but $D(x)$ changes sign from negative to positive as the value of x grows. Only the stable equilibrium exists.

However, the region where both $C(x)$ and $D(x)$ are negative may be too narrow to produce a solution substantially different from total defection.

(e) Both $C(x)$ and $D(x)$ are negative for all values of x. In this case only the stable equilibrium exists and the solution always converges to x_1.

Our experiments totally confirm these findings.

Pavlovian solutions can be predicted for any situation. We have developed an algorithm that accurately predicts the final aggregate outcome for any combination of pavlovian agents and any payoff function [26]. The predictions are exact for an infinite number of agents but the experimental results of the simulation approximate the predictions very closely even for a few hundred agents.

A systematic formal analysis of equilibria of the pavlovian learning model in the N-person game shows that in the case of a linear payoff function and low initial cooperation rate the equilibrium cooperation rate cannot exceed 50 percent [25].

4.2 Conformist agents

The conformist agent imitates the action of the majority. If all agents are conformists and the neighborhood extends to the entire society of agents, then the outcome depends on the exact relationship between the initial number of cooperators and defectors: every agent will immediately imitate the majority and stay there. The behavior becomes quite interesting for the one-layer deep neighborhood. In this case, while the proportion of cooperators will not change substantially, their distribution will. Both cooperators (black spots) and defectors (white spots) will form mutually intertwined clusters (Figure 4).

4.3 Greedy agents

The greedy agent always imitates the behavior of the neighbor with the highest reward (this is the case investigated for dyadic interactions in [34]. If all agents are greedy and the neighborhood extends to the entire organization, they will all defect immediately at the first iteration because they will all imitate the defectors that received a higher reward for their initial action. The situation is not so hopeless for a one-layer deep neighborhood but the behavior will stabilize with a relatively small number of cooperators. Interesting oscillating patterns arise when the payoff functions are those shown in Figure 5 (see Figures 6 and 7).

4.4 Accountants

The accountant's payoff depends on the average reward for its previous actions. If initially the number of cooperators is approximately equal to the number of defectors for a one-layer deep neighborhood, the result

https://doi.org/10.25088/ComplexSystems.14.2.155

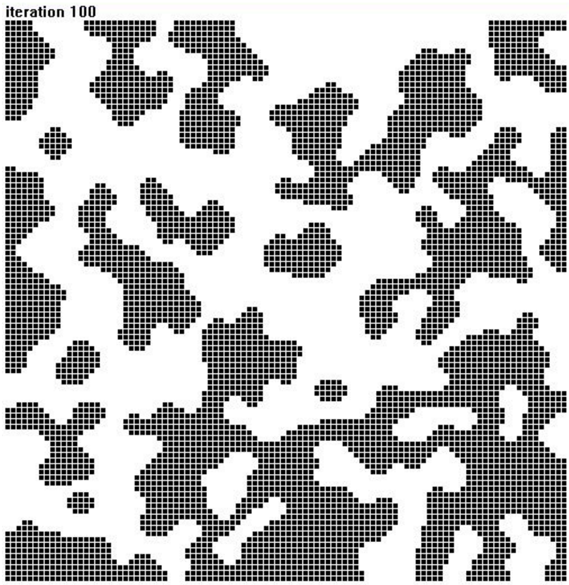

Figure 4. Graphics output of the 100th iteration for the case when all agents are conformists, the payoff curves are given by Figure 1, and the neighborhood is one layer deep. The black spots represent cooperators, the white spots are defectors. The initial ratio of cooperation is equal to 0.50 and the final ratio is 0.49.

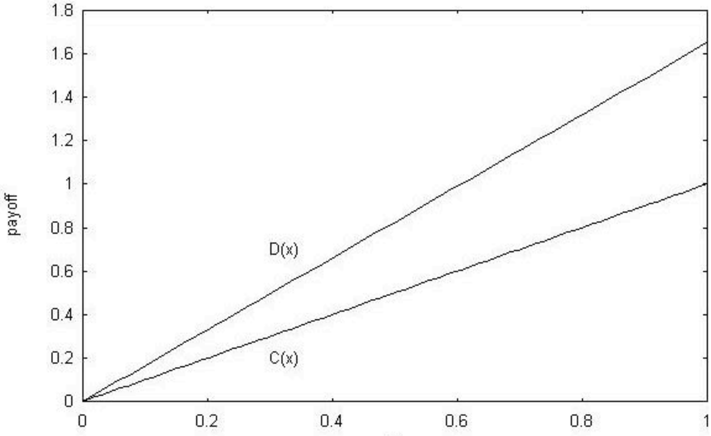

Figure 5. Reward/penalty functions for the case of $D(x) = 1.65x$ and $C(x) = x$.

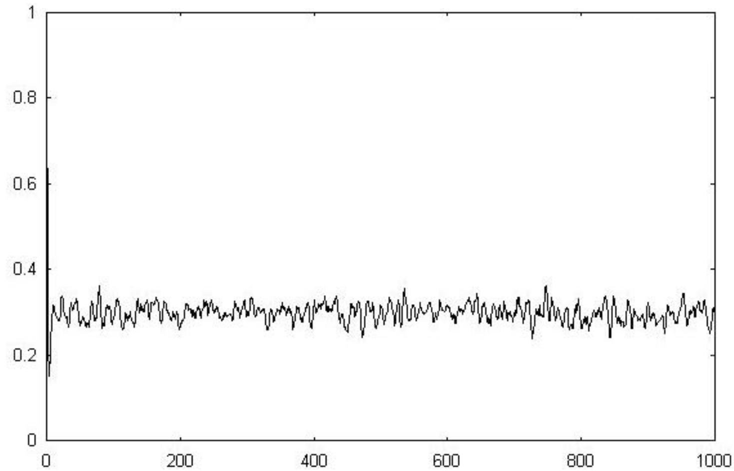

Figure 6. Evolution of the game for the case when all agents are greedy, the payoff curves are given by Figure 5, and the neighborhood is one layer deep. The graph shows the proportion of cooperating agents as a function of the number of iterations. The initial cooperation ratio is equal to 0.9.

Figure 7. A snapshot of the 1000th iteration for the case when all agents are greedy, the payoff curves are given by Figure 5, and the neighborhood is one layer deep. The black spots represent cooperators, the white spots are defectors. The initial ratio of cooperation is equal to 0.90 and the final ratio is 0.29.

https://doi.org/10.25088/ComplexSystems.14.2.155

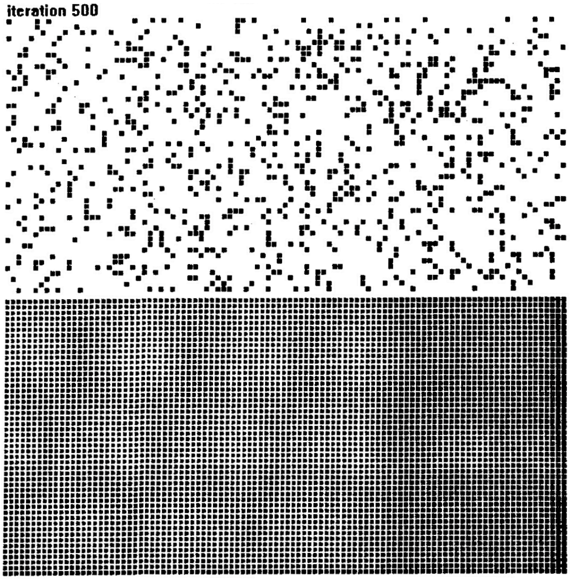

Figure 8. Graphics output of the 500th iteration for the case when the initial actions of all pavlovian agents are random but the array of agents is equally divided into two parts: agents in the upper half initially defect while those in the lower half initially cooperate. The neighborhood is one layer deep.

is universal defection because the defectors' payoff is always higher than that of the cooperators. If, however, the initial distribution is unequal, clusters will form. Agents situated at the borders of cooperative clusters will receive smaller and smaller payoffs. As a result, they will eventually defect, these clusters become smaller and smaller and after several thousand iterations universal defection takes over.

4.5 Nonuniform distributions and mixed personalities

Realistic simulations must take nonuniform distributions of different agents into account. Consider, for example, pavlovian agents with the payoff functions of Figure 1 for the case when the initial actions of all agents are random but the society is equally divided into two parts: agents in the upper half initially defect, those in the lower half initially cooperate. If the neighborhood is one layer deep, the upper half will be gradually infected with cooperators (Figure 8). On the other hand, if the neighborhood is the entire society of agents, change will start in the lower region (Figure 9) and gradually spread to the entire society.

Beautiful oscillating symmetric fractal patterns arise when a single defector initially occupies the middle spot in a sea of greedy cooperators with a one-layer deep neighborhood (Figure 10). It is instructional to

Figure 9. Graphics output of the 63rd iteration for the case of Figure 8 when the neighborhood is the entire array of agents.

Figure 10. Snapshot of the 1000th iteration for the case when a single defector initially sits in the middle of a sea of greedy cooperators. The payoff functions are given in Figure 5. The neighborhood is one layer deep.

investigate the emergence of these patterns. As the $D(x)$ curve is always above the $C(x)$ curve, a layer of defectors will surround the lonely defector after the first iteration. After the second iteration, however, further development depends on the actual shape of the payoff curves. Accordingly, the result may be universal defection, a small stable defection pattern around the center, oscillation in the same region, or the symmetric oscillating pattern of Figure 10. If we allow a small number of individual defectors randomly distributed among cooperators, these

https://doi.org/10.25088/ComplexSystems.14.2.155

patterns interact with each other and can produce other interesting patterns (Figure 11).

The number of variations is infinitely large. We can change all the parameters simultaneously and mix different personalities in arbitrary ways. Figure 12 shows situations similar to that of Figure 6 but with mixed personalities.

5. Conclusion

Experiments performed using a new simulation tool for multi-agent Prisoners' Dilemma games in realistic situations with agents that have various personalities show interesting new results. For the case of pavlovian agents the game has nontrivial but remarkably regular solutions. All solutions strongly depend on the choice of parameter values. Our results show that a viable model for the study of N-person Prisoners' Dilemmas must be based on a more careful selection of parameters than that offered in the literature.

This work is only the first step in analyzing a game with such a high-dimensional parameter space. A more thorough analysis will have to utilize new visualization methods such as the careful use of colors and three-dimensional animation techniques.

Many questions remain open. Future research will find answers to many of them. For example, we will learn the mechanism of cluster formation and the interactions of clusters with each other, the explanation of oscillatory behavior of greedy agents, the role of group size in the emergence of cooperation, and so forth. As a result, the study of N-person Prisoners' Dilemmas may lead us to a better understanding of some basic social dynamics, the emergence of social norms, and may even give some insight into changing human behavior.

Acknowledgment

This work was partially supported by the University of Arizona Foundation and the Office of the Vice President for Research of the University of Arizona.

References

[1] A. Rapoport and A. M. Chammah, *Prisoner's Dilemma* (University of Michigan Press, Ann Arbor, MI, 1965).

[2] T. C. Schelling, "Hockey Helmets, Concealed Weapons, and Daylight Saving," *Journal of Conflict Resolution*, **17**(3) (1973) 381–428.

[3] G. Hardin, "The Tragedy of the Commons," *Science*, **162** (1968) 1243–1248.

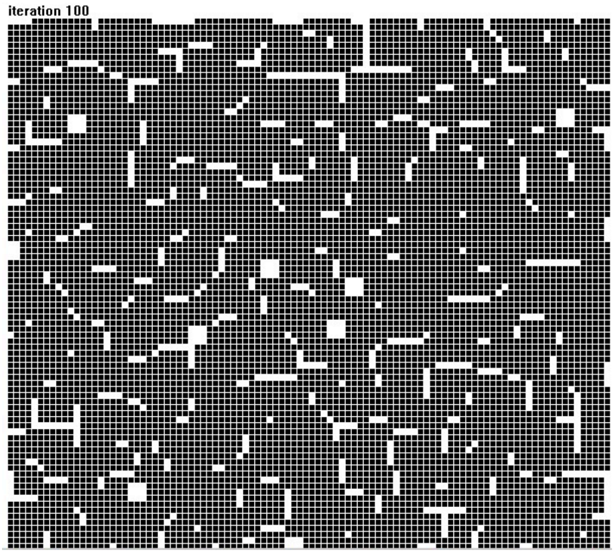

Figure 11. A snapshot of the 120th iteration for the case when all agents are greedy, the payoff curves are given by $C(x) = 5x - 1$ and $D(x) = 5x - 0.5$, and the neighborhood is one layer deep. The initial ratio of cooperation is equal to 0.90, the final ratio oscillates between 0.91 and 0.92.

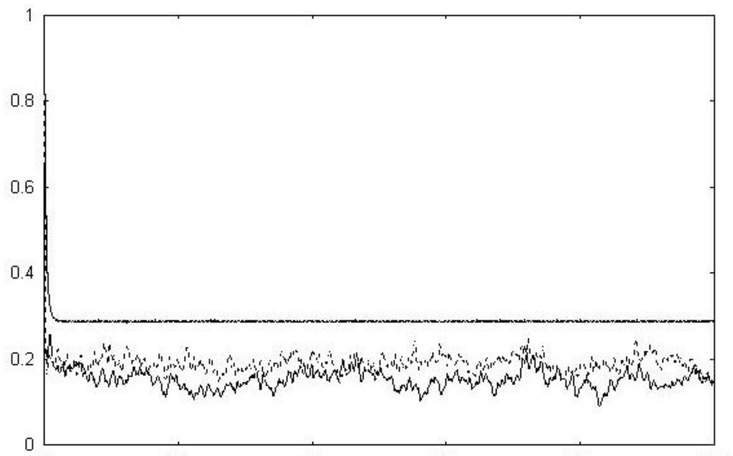

Figure 12. Evolution of the game for the case when the payoff curves are given by Figure 5 and the neighborhood is one layer deep. The graphs show the proportion of cooperating agents as a function of the number of iterations. The lower solid curve corresponds to the case when 97% of the agents are greedy, 3% are Pavlovian. For the middle dotted curve 97% of the agents are greedy, 3% are conformists. For the upper solid curve 45% of the agents are greedy, 45% of them are conformists, and 10% are pavlovian. The initial cooperation ratio is equal to 0.9.

[4] V. E. Bixenstine, C. A. Levitt, and K. V. Wilson, "Collaboration Among Six Persons in a Prisoner's Dilemma Game," *Journal of Conflict Resolution*, **10**(4) (1966) 488–496.

[5] R. L. Weil, "The *N*-person Prisoner's Dilemma: Some Theory and a Computer-oriented Approach," *Behavioral Science*, **11** (1966) 227–234.

[6] A. Rapoport, *N-Person Game Theory* (University of Michigan Press, Ann Arbor, MI, 1970).

[7] H. H. Kelley and J. Grzelak, "Conflict Between Individual and Common Interest in an *N*-person Relationship," *Journal of Personality and Social Psychology*, **21**(2) (1972) 190–197.

[8] H. Hamburger, "*N*-person Prisoner's Dilemma," *Journal of Mathematical Sociology*, **3** (1973) 27–48.

[9] J. M. Anderson, "A Model for 'The Tragedy of the Commons,'" *IEEE Transactions on Systems, Man, and Cybernetics*, (1974) 103–105.

[10] R. M. Dawes, "Formal Models of Dilemmas in Social Decision-making," in *Human Judgment and Decision Processes*, edited by M. F. Kaplan and S. Schwartz (Academic Press, New York, 1975).

[11] R. M. Dawes, "Social Dilemmas," *Ann. Rev. Psychol.*, **31** (1980) 169–193.

[12] P. Bonacich, G. H. Shure, J. P. Kahan, and R. J. Meeker, "Cooperation and Group Size in the *N*-person Prisoners' Dilemma," *Journal of Conflict Resolution*, **20**(4) (1976) 687–706.

[13] D. J. Goehring and J. P. Kahan, "The Uniform *N*-person Prisoner's Dilemma Game," *Journal of Conflict Resolution*, **20**(1) (1976) 111–128.

[14] J. Fox and M. Guyer, "Public Choice and Cooperation in *N*-person Prisoner's Dilemma," *Journal of Conflict Resolution*, **22**(3) (1978) 469–481.

[15] D. D. Heckathorn, "Collective Sanctions and the Creation of Prisoner's Dilemma Norms," *American Journal of Sociology*, **94**(3) (1988) 535–562.

[16] W. B. G. Liebrand, D. M. Messick, and H. A. M. Wilke, editors, *Social Dilemmas: Theoretical Issues and Research Findings* (Pergamon Press, Oxford, New York, 1992).

[17] B. A. Huberman and N. S. Glance, "Evolutionary Games and Computer Simulations," *Proceedings of the National Academy of Sciences USA*, **90** (1993) 7716–7718.

[18] A. Okada, "The Possibility of Cooperation in an *N*-person Prisoners' Dilemma with Institutional Arrangements," *Public Choice*, **77** (1993) 629–656.

[19] S. S. Komorita and C. D. Parks, *Social Dilemmas* (Brown and Benchmark, Madison, WI, 1994)

[20] U. Schulz, W. Albers, and U. Mueller, editors, *Social Dilemmas and Cooperation* (Springer-Verlag, Berlin, Heidelberg, and New York, 1994).

[21] D. A. Schroeder, editor, *Social Dilemmas: Perspectives on Individuals and Groups* (Praeger, Westport, CT, 1995).

[22] K. Nishihara, "A Resolution of *N*-person Prisoners' Dilemma," *Economic Theory*, **10** (1997) 531–540.

[23] R. Hegselmann, "Modeling Social Dynamics by Cellular Automata," in *Computer Modeling of Social Processes*, edited by W. B. G. Liebrand, M. A. Nowak, and R. Hagselmann (Sage, London, 1998).

[24] M. N. Szilagyi, "Solutions to Realistic Prisoners' Dilemma Games," *Proceedings of the 2001 IEEE Systems, Man, and Cybernetics Conference*, TA12/2 (2001) 841–846.

[25] F. Szidarovszky and M. N. Szilagyi, "An Analytical Study of the *N*-Person Prisoners' Dilemma," *Southwest Journal of Pure and Applied Mathematics*.

[26] M. N. Szilagyi and Z. C. Szilagyi, "Nontrivial Solutions to the *N*-Person Prisoners' Dilemma," *Systems Research and Behavioral Science*, **19**(3) (2002) 281–290.

[27] R. Axelrod, *The Evolution of Cooperation* (Basic Books, New York, 1984).

[28] L. Marinoff, "Maximizing Expected Utilities in the Prisoner's Dilemma," *Journal of Conflict Resolution*, **36**(1) (1995) 183–216.

[29] M. W. Macy, "Pavlov and the Evolution of Cooperation: An Experimental Test," *Social Psychology Quarterly*, **58**(2) (1995) 74–87.

[30] D. M. Messick and W. B. G. Liebrand, "Individual Heuristics and the Dynamics of Cooperation in Large Groups," *Psychological Review*, **102** (1995) 131–145.

[31] A. Rapoport, Editorial comments, *Journal of Conflict Resolution*, **38**(1) (1994) 149–151.

[32] R. Hoffmann, "Twenty Years On: The Evolution of Cooperation Revisited," *Journal of Artificial Societies and Social Simulation*, **3** (2000) 2.

[33] M. W. Macy, "Chains of Cooperation: Threshold Effects in Collective Action," *Am. Sociol. Rev.*, **56** (1991) 730–747.

[34] M. A. Nowak and R. M. May, "Evolutionary Games and Spatial Chaos," *Nature*, **359** (1992) 826–829.

[35] A. L. Lloyd, "Computing Bouts of the Prisoner's Dilemma," *Scientific American*, **272**(6) (1995) 110–115.

[36] V. Akimov and M. Soutchansky, "Automata Simulation of *N*-person Social Dilemma Games," *Journal of Conflict Resolution*, **38**(1) (1994) 138–148.

[37] M. N. Szilagyi and Z. C. Szilagyi, "A Tool for Simulated Social Experiments," *Simulation*, **74**(1) (200) 4–10.

[38] E. L. Thorndike, *Animal Intelligence* (Hafner, Darien, CT, 1911).

[39] S. S. Komorita, "A Model of the *N*-person Dilemma-type Game," *Journal of Experimental Social Psychology*, **12** (1976) 357–373.

[40] S. Wolfram, *Cellular Automata and Complexity* (Addison-Wesley, Reading, MA, 1994).

[41] R. Hegselmann and A. Flache, "Understanding Complex Social Dynamics: A Plea for Cellular Automata Based Modelling," *Journal of Artificial Societies and Social Simulation*, **1** (1998) 3.

[42] K. S. Narendra and M. A. L. Thathachar, *Learning Automata: An Introduction* (Prentice Hall, Englewood Cliffs, NJ, 1989).

[43] A. Flache and M. W. Macy, "Weakness of Strong Ties: Collective Action Failure in a Highly Cohesive Group," *Journal of Mathematical Sociology*, **21** (1996) 3–28.

[44] D. Kraines and V. Kraines, "Pavlov and the Prisoner's Dilemma," *Theory and Decision*, **26** (1989) 47–79.

[45] A. Flache and R. Hegselmann, "Rationality vs. Learning in the Evolution of Solidarity Networks: A Theoretical Comparison," *Computational and Mathematical Theory*, **5**(2) (1999) 97–127.