Available online at www.sciencedirect.com
SCIENCE DIRECT®

Journal of Theoretical Biology 231 (2004) 249–260

Journal of Theoretical Biology

www.elsevier.com/locate/yjtbi

Know when to walk away: contingent movement and the evolution of cooperation

C. Athena Aktipis*

Department of Psychology, University of Pennsylvania, 3720 Walnut St., Philadelphia, PA 19104, USA

Received 21 April 2004; received in revised form 22 June 2004; accepted 23 June 2004
Available online 3 August 2004

Abstract

Models of the evolution of cooperation suggest that an important characteristic of successful strategies is the ability to respond contingently to the social environment. A number of mechanisms by which this can be accomplished have been suggested, some of which require relatively complex information processing systems. This research explores relaxing the requirements on information processing while preserving the evolvability of a cooperative strategy. The agent-based computer simulations reported here show that 'Walk Away,' a behavioral rule of extremely limited complexity (move after partner defects), can outperform more complex strategies under a number of conditions. Previous simulations of exit strategies have not examined the effect of implicit and explicit movement costs, different error rates, or the simultaneous presence of TFT and PAVLOV. The simulations reported here establish that the Walk Away strategy resists invasion and can invade a population of defectors at a lower initial frequency than any other strategy. The Walk Away strategy was successful, despite its simplicity, because it exploited aspects of the physical and social environment.

Keywords: Evolution; Cooperation; Walk away; Tit-for-tat; Agent-based; Movement; Exit

1. Introduction

For decades, evolutionary biologists and game theorists have pondered the origins of cooperation, aided by the powerful tools afforded by the game theoretical approach (Axelrod, 1984, 1997; Trivers, 1971). Such work has illustrated that simple strategies such as tit-for-tat (TFT) can be successful. Of course, the success of a given strategy depends on the frequency of other strategies in the social environment, an idea that is borne out in computer simulations of the evolution of cooperation (Axelrod, 1984, 1997).

In most past work, strategies interact with one another according to a stochastic matching process, simulating random encounters with others in the physical world. However, for many species, encounters with others will not only be non-random, but depend systematically on the way that organisms move in space. Because assortment can promote the evolution of cooperation (Wilson and Dugatkin, 1997), spatial models have much to offer. Recently there has been increased interest in spatial models among a number of researchers (Brauchli et al., 1999; Ferriere and Michod, 1995, 1996; Killingback and Doebeli, 1996; Nowak and May, 1992) including evolutionary psychologists (Kenrick et al., 2002, 2003). In the current model, space is simulated as a lattice of occupiable patches, populated with multiple mobile agents. The simulations reported here describe the performance of a cooperative strategy that uses contingent movement to avoid repeated interaction with defectors. The success of this strategy is explored under a variety of parameter values and with a number of competing strategies such as TFT and PAVLOV.

*Corresponding author. Tel.: +1-267-972-5500; fax: +1-215-898-7301.
E-mail address: aktipis@psych.upenn.edu (C.A. Aktipis).

0022-5193/$ - see front matter © 2004 Elsevier Ltd. All rights reserved.
doi:10.1016/j.jtbi.2004.06.020

1.1. Simple strategies

Despite the existence of more complex processes in humans (such as memory for repeated interactions, punishment of defectors and reputation), cooperation might have originally evolved through much simpler mechanisms. Similarly, cooperation in other animals might be realized through relatively simple decision rules. As Anatol Rappoport's TFT strategy (Axelrod, 1984) and Nowak and Sigmund's PAVLOV strategy (Nowak and Sigmund, 1993) illustrate, a simple strategy can outperform more complex strategies. TFTs simplicity lies in the fact that it simply copies the last behavior of its partner. PAVLOV uses a 'win stay, lose shift' strategy, meaning that it switches its behavior (from C to D or from D to C) whenever its partner defects. In the simulations reported here, the viability of another simple strategy, 'Walk Away,' is examined. This strategy could be described as 'win stay, lose move,' and it is different from TFT and PAVLOV in that it always cooperates and it moves away from a defecting partner instead of changing its behavior.

The simulations reported here and elsewhere use the prisoner's dilemma paradigm (Fig. 1). Both individuals are better off if they both cooperate than if they both defect, but each individual would be better off defecting (because defection has a higher payoff no matter what the other player does). In Axelrod's (1984) round-robin simulation he found that the cooperative strategy submitted by Anatol Rappoport, TFT, outperformed all other strategies.

TFT was well suited for the environment in which it competed, but the environment of Axelrod's simulations abstracted from the real world in several important regards. Most relevant to the work described here, the environment in these simulations was not spatial (as the real biological world is) and the interactions took place in a somewhat unrealistic round-robin fashion (which is difficult to interpret spatially). The work reported here aims to create a more realistic environment in which agents can interact.

	Cooperate	Defect
Cooperate	**3**,3	**-1**,5
Defect	**5**,-1	**0**,0

Fig. 1. Prisoner's dilemma payoff matrix. Row player is in bold.

1.2. Game-theoretic models

Axelrod's round-robin tournament is only one of several classes of simulations. Some simulations do not represent space at all (Axelrod, 1984, 1997; Nowak and Sigmund, 1993), and others represent space implicitly by including search cost (Dugatkin, 1992; Dugatkin and Wilson, 1991; Enquist and Leimar, 1993; Eshel and Cavalli-Sforza, 1982; Peck and Feldman, 1986). Spatial game theoretic models represent space explicitly, but do not necessarily involve *movement* in space (Brauchli et al., 1999; Killingback and Doebeli, 1996; Nowak and May, 1992). For these models, spatiality is only relevant in that successful strategies can expand into areas previously occupied by other strategies; the individuals themselves never move.

Another class of simulations, agent-based simulations, are explicitly spatial and can involve movement in space (Ferriere and Michod, 1995, 1996; Pepper and Smuts, 1999). One of the advantages of using spatial agent-based models is that they can approximate many of the features of the biological and social world in which we (and other animals) evolved, while providing the opportunity to carefully control and monitor the changes in the system over time.

In most models, whether they are spatial or not, all individuals have the same number of total interactions. In the present model, this is not the case; agents move in the spatial world until they encounter another agent. After interacting once, they either continue moving or stay on the patch. If both individuals stay, they will continue to interact. If one or both individuals leave, no further interaction takes place between those two individuals.

1.3. Spatiality and movement

Previous research has shown that spatiality can favor cooperation because it generates assortment, resulting in a "conspiracy of cooperators" (Brauchli et al., 1999; Ferriere and Michod, 1995, 1996; Killingback and Doebeli, 1996; Nowak and May, 1992)- although this depends on the payoff structure (Hauert and Doebeli, 2004), the initial parameters and the level of stochasticity (Hauert, 2002). In that work, unlike the present research, individuals do not move in space. Instead, all patches are occupied by individuals who interact only with neighbors. Changes in the population structure are determined by replacing each individual with the most successful strategy among the nearby cells (Nowak and May, 1992), or by similar rules that involve imitating neighbors in proportion to their relative success (Hauert, 2002). These models suggest that spatiality favors the evolution of cooperation, or at least limits the ability of defectors to take over a population.

In contrast, for models with mobile individuals, the conclusion has generally been that mobility constrains the evolution of cooperation because defectors can find cooperators to exploit (Dugatkin, 1992; Dugatkin and Wilson, 1991; Enquist and Leimar, 1993). These models focused on how defectors can use movement for their own gain, but not how cooperators might use movement for their own gain. In the present study, cooperative agents used a contingent movement strategy to avoid repeated interactions with defectors.

1.4. Exit and contingent cooperation

Exit strategies, such as 'out-for-tat' (Yamagishi et al., 1994) are similar to the Walk Away strategy in that they use information about a partner's previous behavior to determine whether or not to interact with that partner. However, the exit strategies explored in other simulations often assumed a fairly complex agent (Vanberg and Congleton, 1992; Yamagishi et al., 1994), or imposed an unrealistic scheme of reassigning agents to another partner immediately after exiting an interaction (Schuessler, 1989). In the present study, it is shown that physical movement away from defectors can substitute for behavioral complexity.

2. Strategy description

By exploiting space, the Walk Away strategy obviates memory and recognition. The Walk Away cooperator moves through space and interacts when it encounters another individual. If its partner cooperates, Walk Away stays on that same patch, but if its partner defects, Walk Away moves. This enables Walk Away to engage in repeated interactions with cooperators and avoid repeated interactions with defectors without using memory. This strategy requires only two states and four transition rules (Fig. 2).

In the first simulation, this strategy is pitted against 3 other simple strategies of equivalent complexity. One is a defecting Walk Away strategy, which employs the same rules but always defects instead of always cooperating (essentially preying on cooperators). Also included were naïve cooperators and defectors, which move when they are without a partner and stay when they find a partner regardless of that partner's behavior (Fig. 2).

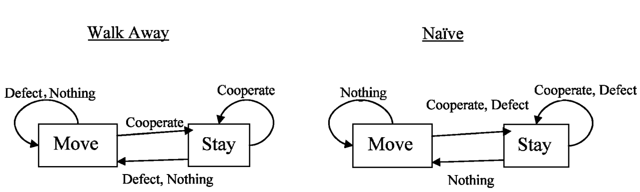

Fig. 2. State transition figures show the simplicity of the Walk Away strategy. Boxes indicate possible states the agent can occupy and arrows show possible transitions between states. Words next to arrows indicate the partner behavior associated with each transition. Walk Away and Naïve are pure strategies with all C and all D versions. When partner behavior on arrow is 'nothing,' this indicates the state change rule for when there is no partner.

After comparing the performance of Walk Away to other moving strategies of minimal complexity, it is then compared to the well-known TFT strategy (Axelrod, 1984). TFT follows the simple rule of copying the last behavior of its partner. Two versions of TFT were implemented in these simulations, one was a mobile TFT and the other was a stationary TFT. Mobile TFT moves in the environment when it is without a partner and stays when it finds a partner, employing the 'tit-for-tat' rule. This strategy is more computationally complex than 'Walk Away,' requiring three states and nine transition rules (Fig. 3). Stationary TFT was included largely because is equivalent to Walk Away in its computational complexity.

Walk Away is also compared to the PAVLOV strategy, another simple strategy that is successful in certain environments (Nowak and Sigmund, 1993). PAVLOV uses a 'win-stay, lose-shift' strategy, which means it switches its strategy whenever its partner defects. This enables it to exploit cooperators but leaves it fairly vulnerable to defectors (because it continually switches back and forth from D to C when it is with a defector). PAVLOV also has the capacity to 'correct' mistakes, avoiding cycles of retaliation that occur with stochastic TFT players, making it more resilient to noise than strategies like TFT (Nowak and Sigmund, 1993). PAVLOV was also implemented as both a mobile strategy (requiring three states and nine transition rules) and stationary strategy (requiring two states and four transition rules; Fig. 3).

3. Methods (simulation description)

In this simulation, energy was the currency that determined reproduction and death. There was no other source of 'energy' besides the interactions with other agents and the initial energy of each agent. At the beginning of the simulation, agents started with an energy level chosen from a uniform distribution between 0 and 49. When the energy of an agent reached 0, the agent died. If the energy of the agent reached 100, that agent reproduced without mutation, creating a copy of itself and splitting its energy with its offspring. Offspring were placed on a random patch.

At the beginning of each simulation, 25 agents of each type were introduced (unless otherwise specified). A carrying capacity equivalent to the starting number of agents for each simulation was implemented to ensure that the simulations would run in a reasonable amount of time. In order to keep the number of initial agents of

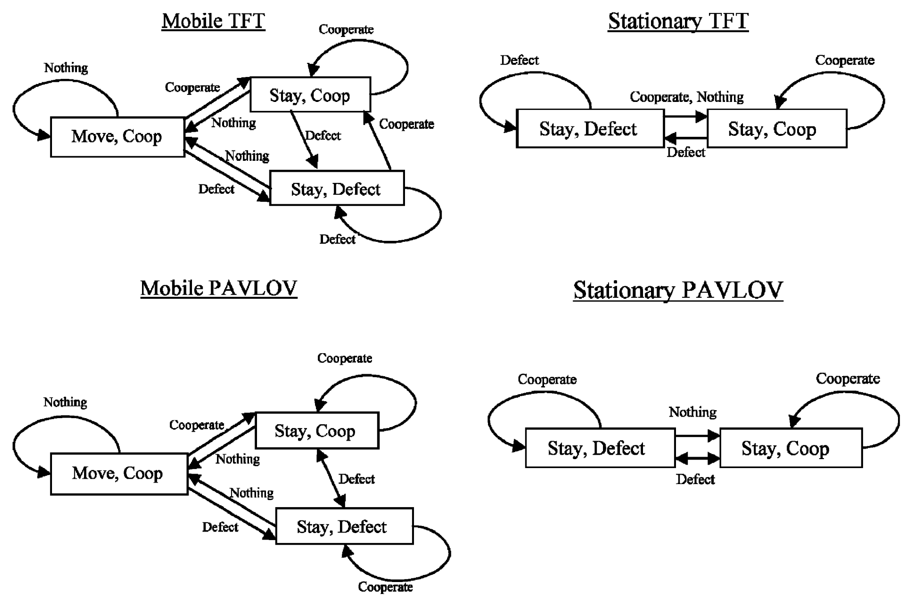

Fig. 3. State transition figures show both versions of TFT and PAVLOV used in these simulations. Boxes indicate possible states the agent can occupy and arrows show possible transitions between states. The 'testing' behavior of PAVLOV is not embodied in this figure (see Sections 2 and 4.1.3). When partner behavior on arrow is 'nothing,' this indicates the state change rule for when there is no partner.

each type the same across simulations, the initial number of total agents (and therefore the carrying capacity) differed across some of the simulations. When the carrying capacity of the environment was reached, a random agent's energy was decreased by 10. These random energy decrements continued until enough agents died and the population was again within the carrying capacity. Occasionally, the total number of agents was lower than the carrying capacity (as can be seen in Figs. 4 and 5) because agents died both when their energy was decremented by the carrying capacity subroutine and when their energy reached 0 due to their interactions with partners. The number of agents of each strategy usually achieved stability within 500 time steps, and the length of all runs was 1000 time steps.

Agents inhabited a spatial world with 25×25 unique patches that could be inhabited by one or more agents. The patches wrapped from bottom to top and left to right, forming a torroidal lattice. All simulations were implemented in Starlogo 1.2.2 (MIT Media Laboratory, 2001), an agent-based programming environment.

In each simulation, agents moved randomly, moving forward and changing their heading randomly to the left or right during each time step. When two agents encountered each other, an interaction took place and payoffs were assigned according to the prisoner's dilemma payoff matrix (Fig. 1). Agents were updated in virtual parallel; during each time step each agent had one interaction with another agent. After this interaction, each agent could choose to stay on the current patch or move to a neighboring patch. Agents either stayed in the same patch unconditionally, or stayed only if their partner cooperated. When both partners stayed on the patch, they interacted again in the next round. If another agent landed on a patch with two interacting agents already on it, two of the three agents on that patch would be randomly paired. This means that an agent could potentially split up an interacting dyad by interacting with one of the agents in the existing pair.

In order to examine how robust a strategy is to 'mistakes,' there was a .001 probability that an agent would make a 'mistake' by cooperating instead of defecting, or vice versa. This error could potentially break up a cooperative pair if the partner was a contingently moving cooperator. In simulation 7, several different error rates were studied in order to investigate the viability of Walk Away under greater noise.

4. Results

4.1. Basic model (simulation 1)

The first simulation used 4 types of agents: 'naïve' cooperators that cooperate and stay even if their partner just defected, Walk Away cooperators that cooperate and then stay only if their partner just cooperated, 'naïve' defectors that defect and stay regardless of their partner's behavior, and Walk Away defectors that defect and only stay with a cooperator. The simulation began with 100 agents, 25 of each type.

During a typical run of this simulation, the proportion of defecting agents (Walk Away defectors and

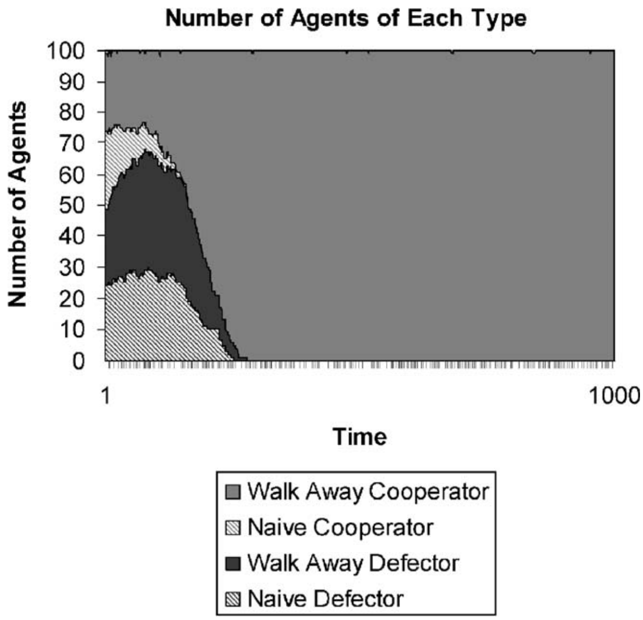

Fig. 4. Area plot of the number of agents of each strategy over time in a typical run of Simulation 1.

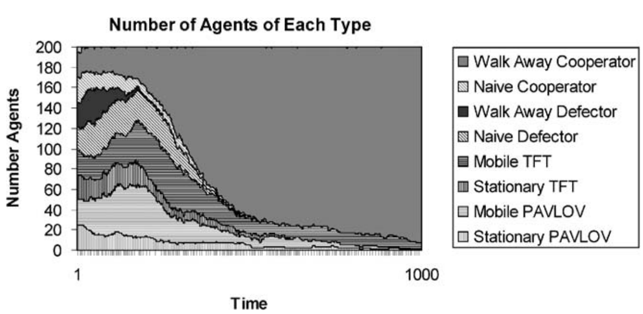

Fig. 5. Area plot of the number of agents of each strategy over time in a typical run of Simulation 4.

'naïve' defectors) initially increased as they exploited 'naïve' cooperators. This caused 'naïve' cooperators to decrease in number and eventually die out. With no 'naïve' cooperators left to exploit (and an inability to exploit Walk Away cooperators for more than one time period), defecting agents decreased in frequency. Walk Away cooperators were then left as the only agents capable of maintaining mutually beneficial interactions and they completely overtook the defecting strategies in each of the 10 runs (see Table 1). Fig. 4 shows the change in number of agents of each strategy over time in a typical run.

4.2. *Comparative performance of TFT and PAVLOV (simulations 2–4)*

The second, third and fourth simulations compared the performance of the four strategies in the basic model with two TFT strategies and two PAVLOV strategies. All features of these simulations were identical to simulation 1 except for the strategies included and the total number of agents; the initial number of agents and the carrying capacity were increased to 150 (simulations 2 and 3) or 200 (simulation 4). Although changing the total number of agents changes the density which has an effect on the outcomes of the eight strategy model, the density increases in these simulations are within a range that has little qualitative effect on the general outcomes (see Section 4.6)[1].

4.2.1. *Simulation 2: Tit-for-Tat*

The second simulation included two types of agents that used the TFT strategy: one that moved when it was without a partner and one that never moved (see Fig. 3). At the end of the simulations, most of agents were Walk Away cooperators ($M = 58.7\%$, $SD = 17.1\%$). However, the mobile TFT strategy always maintained some percentage of the population ($M = 24.4\%$, $SD = 10.6\%$). 'Naïve' cooperators also achieved limited success ($M = 14.4\%$, $SD = 14\%$) (see Table 1.)

4.2.2. *Simulation 3: PAVLOV*

The third simulation was identical to the first simulation, except that it included two PAVLOV strategies. On occasion, PAVLOV attempts defection, which allows it to exploit certain cooperators. In this simulation, the likelihood of PAVLOV 'testing' its partner was .01, which replaced the lower error rate of .001 used by the other strategies. As with TFT, there were two PAVLOV strategies, one that moved unless it had a partner, and another that never moved (see Fig. 3).

Again, the Walk Away cooperator strategy attained the highest frequency in each run ($M = 92.4\%$, $SD = 11.7\%$) (see Table 1). Mobile PAVLOV also maintained a positive frequency in some runs ($M = 6.0\%$, $SD = 10.6\%$).

4.2.3. *Simulation 4: Tit-for-Tat and PAVLOV*

Simulation 4 included all four strategies in simulation 1, as well as the two TFT strategies from simulation 2 and the two PAVLOV strategies from simulation 3, for a total of eight strategies and a carrying capacity of 200 agents. At the end of these runs, Walk Away cooperators were the most successful ($M = 78.1\%$, $SD = 13.7\%$), with mobile TFT second ($M = 10.9\%$, $SD = 10.0\%$) and mobile PAVLOV next ($M = 6.8\%$, $SD = 4.1\%$). 'Naïve' cooperators also had limited success ($M = 3.6\%$, $SD = 5.4\%$). The results of simulation 4 are summarized

[1]In truth, it is not possible to fully determine this based on the simulations run in Section 4.6. In order to do this, simulations with varying density levels would have to be conducted with the basic model as well as the models that include TFT or PAVLOV.

Table 1
Agents of each type in simulations 1–5 M% (SD in %)

	1	2	3	4	5
Naïve coop	0.0 (0.0)	14.4 (14)	0.0 (0.0)	3.6 (5.4)	4.7 (3.9)
Walk away coop	100 (0.0)	58.7 (17.1)	92.4 (11.7)	78.1 (13.7)	76.9 (9.3)
Naïve defect	0.0 (0.0)	0.0 (0.0)	0.0 (0.0)	0.0 (0.0)	0.1 (0.2)
Walk away defect	0.0 (0.0)	0.0 (0.0)	0.0 (0.0)	0.0 (0.0)	0.0 (0.0)
Mobile TFT	N/A	24.4 (10.6)	N/A	10.9 (10.0)	10.7 (8.0)
Stationary TFT	N/A	1.2 (1.5)	N/A	0.5 (0.7)	1.8 (1.3)
Mobile PAVLOV	N/A	N/A	6.0 (10.6)	6.8 (4.1)	5.2 (3.2)
Stationary PAVLOV	N/A	N/A	7.3 (0.2)	0.2 (0.5)	0.7 (0.8)

Results are averaged from 10 runs. Simulations 1–4 started with 25 agents of each included type and changes in the population occurred only through reproduction and death. In Simulation 5, five agents of each type were introduced into the population every 100 time periods.

in Table 1 and Fig. 5 shows the change in frequency of each strategy over 1000 time periods in a typical run.

4.3. Invasion (simulations 5–6)

Simulations 5 and 6 examine Walk Away's resistance to invasion by other strategies and the comparative ability of Walk Away to invade a population of defectors.

4.3.1. Simulation 5: resistance to invasion

Resistance to invasion is a necessary property of an Evolutionarily Stable Strategy (ESS) (Axelrod, 1984). Although the term ESS, as originally formulated, does not strictly apply to this model because this model is made up of a finite number of agents (Nowak et al., 2004), the resistance of Walk Away to invasion by other strategies is nonetheless investigated here. This simulation included all eight strategies used in simulation 4 and was identical in all respects except that every 100-time periods, five agents of each strategy were introduced (with starting energy of 50). This provided an opportunity for strategies that had died out to invade the population. However, this does not constitute a comprehensive test of invadability; such a test would need to include separate simulations to examine whether any strategies can invade Walk Away without any other strategies present.

Results of simulation 5 are very similar to those of simulation 4. Walk Away cooperators were the clear winner ($M = 76.9\%$, $SD = 9.3\%$), with Mobile TFT next ($M = 10.7\%$, $SD = 8.0\%$) and 'naïve' cooperators ($M = 4.7\%$, $SD = 3.9\%$) and Mobile PAVLOV ($M = 5.2\%$, $SD = 3.2\%$) also attaining a limited frequency in the population (see Table 1).

4.3.2. Simulation 6: invasion of cooperation

Previous work has shown that TFT can invade a population of defectors if there are a sufficiently high number of TFT pairings (Ferriere and Michod, 1996).

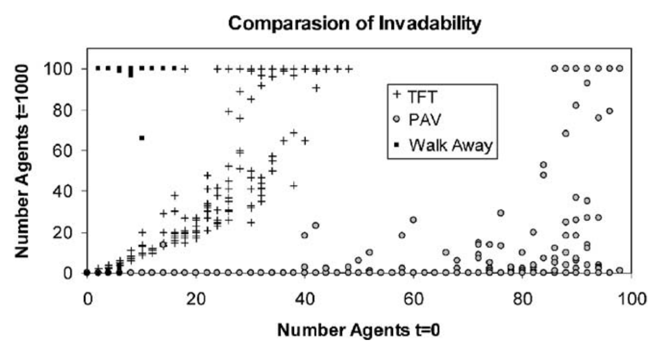

Fig. 6. Minimum initial frequency necessary for agents to invade a population of defectors. Initial number of agents of a particular type (out of 100 total) is plotted on the x-axis and the number of agents after 1000 time periods is plotted on the y-axis. Ten runs were performed for each initial frequency, but simulations were not run for initial numbers of Walk Away over 16 or TFT agents over 48 because it was clear that the strategies successfully invaded at these frequencies.

The simulations in this section compare the abilities of TFT, PAVLOV and Walk Away to invade a population of defectors[2]. The population in all of these simulations was made up 100 agents, including some initial number of one of the following strategies: Walk Away Cooperate, Mobile TFT, or Mobile PAVLOV (indicated on the x-axis of Fig. 6). The remaining agents consisted of half 'naïve' Defectors and half Walk Away Defectors. Ten runs were carried out for each even number of initial frequencies over a range of initial frequencies (i.e., 2, 4, 6,…, 98 initial Walk Away, TFT or PAVLOV agents). As can be seen from Fig. 6, Walk Away cooperators invaded a population of defecting agents at a much lower initial frequency than either TFT or PAVLOV. Walk Away invaded the population of defectors at initial frequencies as low as 2 (out of 100),

[2]The ability of Walk Away to invade populations of cooperative strategies (TFT and PAVLOV) was not examined because Walk Away is functionally identical to these other strategies in the absence of defectors (in that they will always cooperate, except for PAVLOV's occasional 'testing'). This means that there would be no directional selection in such a population of mixed cooperative strategies.

and, at initial frequencies of 8 or more, Walk Away invaded the population in all runs. In contrast, the lowest frequency at which TFT invaded the population was 18, and only at initial frequencies of 40 or more did TFT agents always invade the population of defectors. PAVLOV performed worse than both Walk Away and TFT; the lowest frequency at which PAVLOV invaded was 86, and even when the initial number of PAVLOV agents was 98, they did not always invade.

4.4. Error (simulation 7)

Recall that there was a .001 probability that agents would make a 'mistake' defecting instead of cooperating, or vice versa. Investigations of the role of error indicated that the Walk Away strategy is sensitive to error rates, doing extremely well at low levels of error and doing poorly at high error rates. When the error rate is increased, mobile TFT agents are most successful over an interval of error rates. At yet higher error rates, the 'naïve' defector strategy is most successful. Table 2 summarizes these results.

4.5. Movement cost (simulation 8)

In the real world, there is a cost associated with moving. The viability of Walk Away with various movement costs was examined here. Movement cost was subtracted from the energy of the agent every time that agent moved from one patch to another.

Walk Away was successful at relatively low levels of movement cost (less than 3). However, it is important to take into account the ratio of movement cost to the payoffs associated with partner interactions. The highest payoff possible from an interaction is 5 (when an agent defects with a cooperator) and the mutually cooperative payoff is only 3. It appears that the Walk Away strategy continued to be successful in the face of movement costs (at this particular inverse density level of 3.13), as long as the movement costs were less than the benefit from one round of cooperation.

When movement cost was between 3 and 9, Mobile TFT was the most successful strategy. Presumably, Mobile TFT outperforms Walk Away at intermediate levels of movement cost because it stays with a partner after a defecting 'mistake,' while Walk Away moves after defection even if that defection is the result of error. This means that Walk Away will incur movement costs much more often than Mobile TFT.

When the movement cost was between 9 and 45 (movement cost of greater than 45 were not explored in these simulations), Stationary TFT and Stationary PAVLOV dominated the population in approximately equal numbers (although mobile strategies continued to make up approximately 25% of the population). Because the stationary strategies did not move at all, their success in the face of high movement costs is not surprising. However, the continued existence of mobile strategies is more surprising and the reasons for it more complex. Because agents are placed on random patches when they 'hatch,' stationary strategies only get partners if other agents find them. Even when the cost of moving is very high, it can still be worth moving if the likelihood of obtaining a partner as a stationary individual is sufficiently low. This results in frequency dependent selection on mobile strategies, which accounts for their continued existence in the face of high movement costs.

4.6. Density (simulation 9)

Several different densities were investigated in order to determine what density conditions cause Walk Away to lose its advantage. Ten simulations were run for 1000 time steps at several density levels. Number of patches per agent (the inverse of density) was varied, but because of software constraints and the need to keep the grid square, these values are not whole numbers. In the simulations reported in earlier sections of this paper, the number of patches per agent ranged from 6.25 (simulations with 100 agents) to 3.13 (simulations with 200 agents). However, in the present section, a much larger range of inverse density values was applied (.85–41.41).

Table 2
Agents of each type in simulations 1–5 M% (SD%)

	.001[a]	.002	.005	.01	.02	.1
Naïve coop	2.2 (2.9)	2.1 (3.5)	1.5 (2.5)	3.0 (5.0)	0.7 (1.7)	0.0 (0.0)
Walk away coop	**79.2 (10.6)**	**75.2** (15.7)	**62.9** (20.3)	39.4 (24.8)	20.7 (33.6)	0.0 (0.0)
Naïve defect	0.0 (0.0)	0.0 (0.0)	0.0 (0.0)	0.0 (0.0)	1.1 (3.5)	**93.9 (10.9)**
Walk away defect	0.0 (0.0)	0.0 (0.0)	0.0 (0.0)	0.0 (0.0)	0.0 (0.0)	0.0 (0.0)
Mobile TFT	13.6 (9.6)	16.9 (15.0)	26.2 (21.1)	**44.5** (27.3)	**68.1** (35.0)	5.6 (10.5)
Stationary TFT	0.4 (.9)	0.7 (1.1)	0.9 (1.2)	1.0 (1.9)	2.3 (3.4)	0.2 (0.7)
Mobile PAVLOV	4.6 (4.7)	5.1 (6.0)	8.5 (7.5)	11.8 (9.8)	7.0 (7.1)	0.3 (0.8)
Stationary pavlov	0.1 (0.2)	0.0 (.1)	0.1 (.3)	0.4 (0.9)	0.2 (0.6)	0.0 (0.0)

The most successful strategy at each error rate is in bold.
[a]All other simulations reported in this article use this error rate.

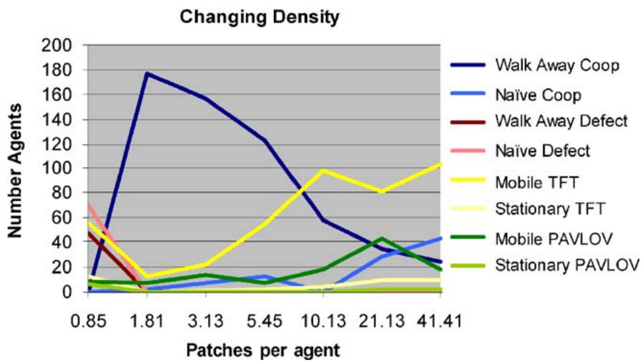

Fig. 7. Number of agents at various densities.

At very high densities (when there was less than one patch per agent), Walk Away did extremely poorly and the defecting types were most successful. Mobile TFT also did relatively well.

Although Walk Away does poorly at extremely high densities (i.e., low inverse densities) because dyads are continuously broken up by less cooperative agents, Walk Away does very well over a fairly large range of densities (see Fig. 7). Walk Away cooperators achieved an average 88.6% of agents (SD = 10.9%) when the number of patches per agent was as low as 1.81. Walk Away cooperators also did well when the number of patches per agent was much higher; at 5.45 patches per agent, Walk Away cooperators still achieved an average of 61.5% (SD = 15.1%).

Only when the number of patches per agent was increased to 10.13 did mobile TFT (M = 49.1%, SD = 15.3%) agents achieve a greater frequency than Walk Away cooperators (M = 29.1%, SD = 8.9%). 'Naïve' cooperators also faired better when the number of patches per agent was high, although this was not apparent until the number of patches per agent increased to 21.13. Walk Away's strategy involved leaving a partner after one defection, even if that defection was the result of a mistake (which occurred with probability .001). The most likely reason that the strategy did not do as well at extremely low densities is that it left these potentially cooperative and very hard to find partnerships after just one defection. In these circumstances TFT had an advantage, since it did not leave after its partner defected.

4.7. Increased mutual defection payoff (simulation 10)

In simulation 10, the payoff for mutual defection was increased from 0 to 1. Relative success of each strategy in both the four-strategy model and the more complex eight-strategy model were investigated. Results are reported for 100 runs with a length of 1000 time steps each.

In the basic four-strategy model with Walk Away cooperators, naïve cooperators, Walk Away defectors and naïve defectors, the results over 100 runs were slightly different. In 91 runs, only Walk Away cooperators were left after 1000 time periods. Of the remaining nine runs, six ended with only naïve defectors, two ended with mostly Walk Away cooperators and a small number of naïve cooperators, and one ended with equal numbers of Walk Away cooperators and naïve defectors. In six of the runs, naïve defectors were able to interact with each other and out-compete Walk Away cooperators because of the small but positive mutual defection payoff. In the runs that ended with both Walk Away cooperators and naïve defectors, all dyads (lasting longer than one time period) were made up of homogenous individuals (i.e., either both Walk Away cooperators or both naïve defectors).

In contrast, results of the eight-strategy model (with TFT and PAVLOV) showed that Walk Away cooperators dominated the population in every run, averaging 75.6% (SD = 17.6%) at the end of each run. Mobile TFT agents were present in the population at the end of most runs (M = 19.6%, SD = 16.2%), whereas defecting agents were absent at the end of every run.

5. Discussion

Walk Away agents had no memory of the action of other agents; they only responded to their partner's most recent behavior. Despite these limited capabilities, Walk Away was successful because agents employing this strategy were able to achieve assortative interactions. Walk Away was able to reap the benefits of repeated interactions with cooperators because other cooperators stayed on the same patch and this strategy was able to avoid continued interactions with defectors by leaving the patch.

5.1. Contingent movement and assortment

The movement rules employed by Walk Away resulted in behavioral assortment. When Walk Away interacted with a defector, it moved, but when it interacted with another cooperator, it stayed. The cooperative dyads that resulted from this assortment were far more stable and fecund than dyads made up of agents that did not cooperate in every round. In fact, when agents use movement rules to avoid defectors, dyads of defectors or mixed dyads will be much less stable (often existing for only one round), than groups of cooperators. In essence, the Walk Away cooperators outperformed other strategies due to the effects of between-group selection. As Wilson (Wilson, 1983; Wilson and Dugatkin, 1997) notes, selection can act at the group level when there is behavioral assortment because this leads to greater variation between groups. Because of the movement rule employed by Walk Away, behavioral assortment led to greater fitness of the cooperative dyads (compared to non-cooperative ones). This led to greater evolutionary success for agents employing the Walk Away strategy[3].

Like TFT, strategies based on altruistic punishment (Fehr and Gachter, 2002; Price et al., 2002), indirect reciprocity (Nowak and Sigmund, 1998a, b; Panchanathan and Boyd, 2003), and other types of assortment (Eshel and Cavalli-Sforza, 1982; Wilson and Dugatkin, 1997) enable the formation of a "conspiracy of cooperators," limiting potential benefits to defectors. The PAVLOV strategy accomplishes this to a lesser extent, as defectors can exploit PAVLOV's switching behavior.

5.2. Comparing walk away to TFT and PAVLOV

In several ways, the Walk Away strategy is similar to TFT. After interacting with an agent who defects, both Walk Away and TFT stop cooperating with that individual: TFT punishes the defector by defecting in turn; Walk Away 'punishes' the defector by not interacting in the next time period. Because certain conceptions of the TFT strategy equate defection and non-action, often implicitly through having mutual defection payoffs of 0, 0 (e.g., Nowak and May, 1992), TFT and Walk Away can be construed as fundamentally the same strategy. However, the difference between Walk Away and TFT becomes clear by virtue of their unique instantiations in a spatial, multi-agent world. While TFT stays with a defecting partner, Walk Away leaves a partner who defects (and seeks out a new partner), no matter how many times a partner has cooperated before. After an interaction with a defector, Walk Away pursues a strategy of contingent movement (switching from stay to move) while TFT pursues a strategy of contingent behavior (switching from cooperate to defect). While Walk Away's strategy enables it to find a new partner, TFT is essentially stuck in a partnership with whichever agent happens to be on the same patch.

An additional difference between TFT and Walk Away is that Walk Away needs no memory, while TFT needs memory for at least one round of play. When Walk Away interacts with a defector, it simply responds to that defection by moving. TFT, on the other hand, must remember that interaction so it can respond appropriately in the next round (unless TFT is somehow responding to the behavior of its partner without actually remembering it). Although this point might seem trivial, the alternative conception of contingent cooperation provided by the Walk Away strategy might provide deeper insight into the evolution of cooperation in non-human animals (see Section 5.4).

Like PAVLOV, the Walk Away cooperator and Walk Away defector strategies use contingent cooperation, but instead of employing a 'win-stay, lose-shift,' strategy, Walk Away uses a 'win-stay, lose-move' strategy, never switching to defection. PAVLOV responds to the behavior of its partner by continuing to do the same thing (either cooperating or defecting) if its partner cooperates and changing state (to cooperation or defection) if its partner defects (Axelrod, 1997; Nowak and Sigmund, 1993). Both PAVLOV and TFT differ from Walk Away in that they defect under certain conditions, while Walk Away is a purely cooperative strategy.

5.3. Similar models of the evolution of cooperation

Interestingly, the results of the current simulations largely contradict the conclusions of a similar study in which Enquist and Leimar (1993) investigated the evolution of cooperation in mobile organisms using a mathematical model. They concluded that mobility constrained the evolution of cooperation (in their simulation the cooperative strategy played TFT). However, their model did suggest that long search times (for partners) and long coalition times (i.e., the time during which a particular dyad stays together) would favor the evolution of cooperation. This suggests that the different conclusions of the simulations presented in this paper and Enquist and Leimar's model turn on differences in search times and the length of coalitions, both being longer in the simulation reported here. Indeed, the results reported in Section 4.6 show that under very high densities (where search time is short and coalition time is short), defecting agents are most successful.

Several models of exit strategies bear resemblance to the current study, but have fundamentally different assumptions. Yamagishi et al. (1994) and Vanberg and Congleton (1992) showed that a cooperative strategy that refrains from interacting with defectors does relatively well. However, they assumed that these strategies had extensive memories for past interactions, while the present model assumes no such memory.

Schuessler's (1989) CONCO strategy was similar to Walk Away except that it immediately entered a new dyad after leaving a partner. Walk Away, on the other hand, must search for a new partner after leaving one. In the real world, there is usually some cost associated with leaving an interaction partner, even if it is just the cost of not interacting for several rounds (until one can find a new partner). When there is no cost associated with finding a new partner (as is the case in Schuessler's

[3] It is important to note that Walk Away does not need to form clusters of dyads in order to succeed. Dyads of Walk Away agents were instead relatively evenly distributed over the spatial grid. The only assortment taking place in this simulation was within dyads (i.e., Walk Away cooperators formed dyads with each other).

simulations, but not in the simulations reported in this paper), the conditions are more favorable for a deserting strategy, so its success should be less surprising.

5.4. Simple strategies revisited

The work reported in this paper shows that not only can contingent movement promote the evolution of cooperation, but that it can do so with minimal cognitive complexity. Some recent work has focused on complex and memory-intensive strategies in the evolution of cooperation, rather than exploring other simple strategies that might promote the evolution of cooperation. Researchers have studied the role of reputation, finding that indirect reciprocity (Nowak and Sigmund, 1998a, b; Panchanathan and Boyd, 2003) and the ability to gossip (Nakamaru and Kawata, 2002) increase the likelihood of the evolution of cooperation. Nonetheless, these strategies require both large memory capacities and communication abilities.

Other work suggests that costly punishment of cheaters (by first and third parties) can evolve, increasing the viability of the cooperative strategy (Axelrod, 1997; Boyd and Richerson, 1992). Although punishment seems to be important in human interactions (Fehr and Gachter, 2002; Price et al., 2002), it requires fairly complex behavioral rules so it is not as likely to have played a role in the first stages of the evolution of cooperation in humans and other species. Subjective commitment has also been explored as a possible means to the evolution of cooperation (Frank, 1988; Nesse, 2001), but this again requires complex cognitive capacities.

As is apparent from the state transition figures (Figs. 2 and 3), Walk Away requires only two states, moving and staying, while TFT and PAVLOV require three (the two-state stationary versions of these fared poorly). In this respect, Walk Away is even simpler than TFT and PAVLOV. Recent research suggests that a number of 'simple heuristics' can work better than complex, information intensive strategies (Gigerenzer et al., 1999). These strategies often succeed because they exploit the information structure of the environment directly rather than storing all the relevant information in order to perform computations on it. The Walk Away strategy outperforms more complex strategies because it does just that—exploiting the structure of the environment by staying on patches with cooperators and moving away from defectors.

Given that a strategy based on movement rules does not require any memory, it can also be used by the simplest of organisms, including those without brains. It could, for instance, be implemented by responses to chemical gradients that are products of different types of interactions. If we assume that organisms originally used simple decision rules and small (or non-existent) memory capacities when making decisions about cooperation, the Walk Away strategy might be a likely candidate because it is both simple and successful under a variety of parameter values. Indeed, Walk Away requires neither the notion of intentionality nor the ability to represent conspecifics. An organism using the Walk Away strategy can respond to others just like it would respond to the inanimate environment: by moving away if it is incurring a cost.

Humans might also use something akin to the Walk Away strategy in their social lives. Instead of persisting in a mutually defecting relationship, individuals often seek out new social partners. Previous research has suggested that the TFT and PAVLOV strategies might approximate human behavior relatively well (Milinski and Wedekind, 1998), that people tend to make use of exit strategies when given the option to desert a partner (Yamagishi and Hayashi, 1996), and, intriguingly, that the option to exit increases cooperation in individuals who tend to be cooperative in the first place (Boone and Macy, 1999; Hauk, 1999; Orbell and Dawes, 1993). If humans do indeed 'walk away' from uncooperative partners—and this increases cooperation among individuals with cooperative tendencies—then this could have important implications for the role of contingent movement in the evolution of cooperation in humans. Although bringing these two findings together in this way is speculative, the combined effect suggests the following: behavioral assortment resulting from contingent movement might increase between group variation in fitness (because those with cooperative tendencies are more cooperative when they have the option to exit), resulting in stronger selection pressures for both cooperation and contingent movement.

6. Conclusion

In sum, the results of the simulations reported here show that cooperative agents with simple contingent movement rules can outperform TFT and PAVLOV under a wide array of parameter values. The Walk Away strategy avoids repeated interactions with defectors and reaps the benefits of interacting with cooperators without compromising its simplicity. These simulations demonstrate that Walk Away is successful when movement cost, error rates and number of patches per agent (inverse density) are low, although extremely low inverse densities favor defectors. Furthermore, the present study shows that Walk Away is resistant to invasion by many other strategies, and that it can invade populations of defectors at lower initial frequencies than either TFT or PAVLOV.

Both the adaptationist approach (Williams, 1966) and the 'simple heuristics' approach (Gigerenzer et al., 1999) emphasize the importance of an organism's environment in understanding its cognitive and behavioral adaptations. The simulation results reported in this paper illustrate the point that it is often important to model aspects of the social environment, such as spatiality and mobility. The Walk Away strategy would have been impossible to implement in an environment that lacked these features. Indeed, it is these very features—spatiality and mobility—which allowed Walk Away to avoid repeated interactions with defectors and maintain interactions with other cooperators without employing complex strategies.

Acknowledgements

I would like to thank Robert Kurzban, David Sloan Wilson, Karthik Panchanathan, John Orbell and two anonymous reviewers for their valuable comments on earlier drafts of this paper. This material is based upon work supported under a National Science Foundation Graduate Research Fellowship.

References

Axelrod, R., 1984. The Evolution of Cooperation. Basic Books, New York.

Axelrod, R., 1997. The Complexity of Cooperation: Agent-Based Models of Competition and Collaboration. Princeton University Press, Princeton.

Boone, R.T., Macy, M.W., 1999. Unlocking the doors to prisoner's dilemma: dependence, selectivity, and cooperation. Social Psychology Quarterly 62, 32–52.

Boyd, R., Richerson, P.J., 1992. Punishment allows the evolution of cooperation (or anything else) in sizable groups. Ethol. Sociobiol. 13 (3), 171–195.

Brauchli, K., Killingback, T., Doebeli, M., 1999. Evolution of cooperation in spatially structured populations. J. Theor. Biol. 200, 405–417.

Dugatkin, L.A., 1992. The evolution of the "con artist". Ethol. Sociobiol. 13, 3–18.

Dugatkin, L.A., Wilson, D.S., 1991. ROVER: a strategy for exploiting cooperators in a patchy environment. Am. Naturalist 138 (3), 687–701.

Enquist, M., Leimar, O., 1993. The evolution of cooperation in mobile organisms. Animal Behav. 45, 747–757.

Eshel, I., Cavalli-Sforza, L.L., 1982. Assortment of encounters and evolution of cooperativeness. Proc. Natl Acad. Sci. 79 (4), 1331–1335.

Fehr, E., Gachter, S., 2002. Altruistic punishment in humans. Nature 415, 137–140.

Ferriere, R., Michod, R.E., 1995. Invading wave of cooperation in a spatial iterated prisoner's dilemma. Proc.: Biol. Sci. 259 (1354), 77–83.

Ferriere, R., Michod, R.E., 1996. The evolution of cooperation in spatially heterogeneous populations. Am. Naturalist 147 (5), 692–717.

Frank, R.H., 1988. Passions Within Reason: The Strategic Role of the Emotions. W.W. Norton & Company, Inc, New York.

Gigerenzer, G., Todd, P.M., The ABC Research Group, 1999. Simple heuristics that make us smart. Oxford University Press, New York.

Hauert, C., 2002. Effects of space in 2×2 games. Int. J. Bifurcation Chaos 12 (7), 1531–1548.

Hauert, C., Doebeli, M., 2004. Spatial structure often inhibits the evolution of cooperation in the snowdrift game. Nature 428, 643–646.

Hauk, E., 1999. Multiple prisoner's dilemma games with(out) an outside option: an experimental study, Universitat Pompeu Fabra.

Kenrick, D.T., Maner, J.K., Butner, J., Li, N.P., Becker, D.V., Schaller, M., 2002. Dynamical evolutionary psychology: mapping the domains of the new interactionist paradigm. Personality Social Psychol. Rev. 6 (4), 247–356.

Kenrick, D.T., Li, N.P., Butner, J., 2003. Dynamical evolutionary psychology: individual decision rules and emergent social norms. Psychol. Rev. 110 (1), 3–28.

Killingback, T., Doebeli, M., 1996. Spatial evolutionary game theory: hawks and doves revisited. Proc.: Biol. Sci. 263 (1374), 1135–1144.

Milinski, M., Wedekind, C., 1998. Working memory constrains human cooperation in the Prisoner's Dilemma. Proc. Natl Acad. Sci. 95, 13755–13758.

MIT Media Laboratory, 2001. StarLogo, Cambridge, MA.

Nakamaru, M., Kawata, M., 2002. Evolution of rumors that discriminate lying defectors. Evol. Ecol. Res. 6, 261–283.

Nesse, R.M., 2001. Natural Selection and the Capacity for Subjective Commitment. In: Nesse, R.M. (Ed.), Evolution and the Capacity for Commitment Russel Sage Foundation. New York, pp. 1–44.

Nowak, M.A., May, R.M., 1992. Evolutionary games and spatial chaos. Nature 359, 826–829.

Nowak, M., Sigmund, K., 1993. A strategy of win-stay, lose-shift that outperforms tit-for-tat in the Prisoner's Dilemma game. Nature 364, 56–58.

Nowak, M.A., Sigmund, K., 1998a. The dynamics of indirect reciprocity. J. Theor. Biol. 194, 561–574.

Nowak, M.A., Sigmund, K., 1998b. Evolution of indirect reciprocity by image scoring. Nature 393, 573–577.

Nowak, M.A., Sasaki, A., Taylor, C., Fudenberg, D., 2004. Emergence of cooperation and evolutionary stability in finite populations. Nature 428 (8), 646–650.

Orbell, J.M., Dawes, R.M., 1993. Social welfare, cooperators' advantage, and the option of not playing the game. Am. Socio. Rev. 58, 787–800.

Panchanathan, K., Boyd, R., 2003. A tale of two defectors: the importance of standing for evolution of indirect reciprocity. J. Theor. Biol. (224), pp. 115–126. doi.10.1016/S0022-5193(03)00154-1.

Peck, J.R., Feldman, M.W., 1986. The evolution of helping behavior in large, randomly mixed populations. Am. Naturalist 127 (2), 209–221.

Pepper, J.W., Smuts, B., 1999. The evolution of cooperation in an ecological context: an agent-based model. In: Kohler, T., Gumerman, G. (Eds.), Dynamics in Human and Primate Societies. Oxford University Press, Oxford, pp. 45–76.

Price, M.E., Cosmides, L., Tooby, J., 2002. Punitive sentiment as an anti-free rider psychological device. Evol. Human Behav. 23, 203–231.

Schuessler, R., 1989. Exit threats and cooperation under anonymity. J. Conflict Resolution 33 (4), 728–749.

Trivers, R.L., 1971. The evolution of reciprocal altruism. Quart. Rev. Biol. 46 (March), 35–57.

Vanberg, V.J., Congleton, R.D., 1992. Rationality, morality and exit. Am. Political Sci. Rev. 86 (2), 418–431.

Williams, G.C., 1966. Adaptation and Natural Selection. Princeton University Press, Princeton, NJ.

Wilson, D.S., 1983. The group selection controversy: history and current status. Ann. Rev. Ecol. Syst. 14, 159–187.

Wilson, D.S., Dugatkin, L.A., 1997. Group selection and assortative interactions. Am. Naturalist 149 (2), 336–351.
Yamagishi, T., Hayashi, N., 1996. Selective play: social embeddedness of social dilemmas. In: Liebrand, W., Messick, D. (Eds.), Frontiers in Social Dilemma Research. Springer, Berlin, pp. 363–384.
Yamagishi, T., Hayashi, N., Jin, N., 1994. Prisoner's dilemma networks: selection strategy versus action strategy. In: Schulz, U., Albers, W., Mueller, U. (Eds.), Social Dilemmas and Cooperation. Springer, Berlin, pp. 233–250.