



[David Joyce, John Kennison, Owen Densmore, Steven Guerin, Shawn Barr, Eric Charles and Nicholas S. Thompson \(2006\)](#)

My Way or the Highway: a More Naturalistic Model of Altruism Tested in an Iterative Prisoners' Dilemma

Journal of Artificial Societies and Social Simulation vol. 9, no. 2
<<http://jasss.soc.surrey.ac.uk/9/2/4.html>>

For information about citing this article, click [here](#)

Received: 18-Aug-2005 Accepted: 08-Mar-2006 Published: 31-Mar-2006



Abstract

There are three prominent solutions to the Darwinian problem of altruism, kin selection, reciprocal altruism, and trait group selection. Only one, reciprocal altruism, most commonly implemented in game theory as a TIT FOR TAT strategy, is not based on the principle of conditional association. On the contrary, TIT FOR TAT implements conditional altruism in the context of unconditionally determined associates. Simulations based on Axelrod's famous tournament have led many to conclude that conditional altruism among unconditional partners lies at the core of much human and animal social behavior. But the results that have been used to support this conclusion are largely artifacts of the structure of the Axelrod tournament, which explicitly disallowed conditional association as a strategy. In this study, we modify the rules of the tournament to permit competition between conditional associates and conditional altruists. We provide evidence that when unconditional altruism is paired with conditional association, a strategy we called MOTH, it can outcompete TIT FOR TAT under a wide range of conditions.

Keywords:

Game Theory; Altruism; Prisoners' Dilemma; TIT FOR TAT; MOTH; Docking; Netlogo

The Problem of Altruism

1.1

In discussions of the biological evolution of social behavior, an organism is altruistic, by definition, when it enhances the relative reproduction of another organism at some cost to its own ([Sober and Wilson 1998](#)). The existence of altruistic behavior seems to present a paradox for evolutionary theorists because reproductive advantage is thought by many to be the principal engine of evolutionary change. The problem is, of course, that altruistic benefactors are at a reproductive disadvantage to their non-altruistic beneficiaries. In game theoretic discussion, this disadvantage is often embodied in a game in which 1) when altruists interact, they exchange the cost and benefit of altruism, 2) when altruists interact with selfish individuals, the altruist pays the cost and the selfish individual receives the benefit, 3) and when selfish individuals interact, they receive no benefit and expend no cost. The basic payoff structure of "Altruist Game" is shown in Table 1.

Standard Altruism Game

... .. from a meeting with...

		Altruist	Selfish
Pay off received by	Altruist	$b - c$	$-c$
	Selfish	b	0

1.2

Given that humans and other animals do in fact cooperate, the challenge to the evolutionary theorist has been to find some mechanism that will permit altruism to resist elimination or to eliminate more selfish alternatives. Two answers to this challenge, kin selection and trait-group selection, are conditional association strategies: They postulate a bias in the selection of social partners such that altruists meet altruists more often than by chance, and selfish individuals meet altruists less often than by chance. This is mathematically straightforward:

$$p = \text{the proportion of altruists in the population}$$

$$p_{A,A} = \text{the probability that an altruist encounters another altruist}$$

$$p_{S,A} = \text{the probability that a selfish individual encounters an altruist}$$

1.3

If selfish and altruistic individuals are equally represented in a population and encounters occur at random, then $p_{A,A} = p = p_{S,A}$. In these circumstances, selfish individuals will outperform altruists and eventually the entire population will become selfish. But suppose that altruists discriminate, with the result that $p_{A,A} > p_{S,A}$. Simple algebra shows that an unconditionally altruistic strategy will prevail over an unconditionally selfish strategy whenever the difference between $p_{A,A}$ and $p_{S,A}$ is greater than the ratio of costs (c) to benefits (b) (See Table 2). Kin selection theorists will recognize this principle as the basis of *Hamilton's inequality*, $r > c/b$, where r is defined as the degree of relatedness amongst associates. Group selection theorists will recognize it as the principle underlying trait group section ([Sober and Wilson 1998](#)) that partially segregates altruistic and selfish individuals into different ephemeral groups.

Payoff received by Individuals...		Against individuals playing		Total Payoff
		Discriminating Altruist with probability p	Selfish with probability $(1-p)$	
PLAYING	Discriminating Altruist	$(p+a)(b-c)$	$(1-p-a)(-c)$	$(p+a)(b-c)+(1-p-a)(-c)=pb-pc+ab-ac-c+pc+ac=pb+ab-c$
	Selfish	$(p-s)(b)$	$(1-p+s)(\theta)$	$(p)(b)-(s)(b)+(1-p)(\theta)=pb-sb$

Discriminating Altruists will come to characterize the population ONLY if the total payoff to them is greater than the total payoff to Selfish individuals: i.e., when $pb+ab-c > pb-sb$ or when the sum of the two discrimination coefficients exceeds the ratio of costs to benefits -- when $(a+s) = P_{AA} - P_{AS} > (c/b)$.

Tit for Tat Reciprocity

1.4

One remarkably influential solution to the altruism problem does not appear to be of that form. It has been based on a "TIT FOR TAT" strategy in an iterated prisoners' dilemma game. In prisoners' dilemmas, altruists are called Cooperators and selfish individuals, Defectors. A prisoners' dilemma game imagines a situation in which two arrested individuals are being coaxed to confess to doing a crime. In such a scenario, there is a reward for cooperation (R), a temptation to defect (T), a loss suffered when a loyal "sucker" is betrayed (S), and a punishment for non-cooperation (P). Any game can be categorized as a prisoners' dilemma so long as $T > R > P > S$ and $2R > (T + S)$ (See Table 3). It should be noted that all payoff patterns that are altruist games are also prisoners' dilemma games.

Payoff received by player 1, when		Against Player 2, playing	
		Cooperate	Defect (Cheat)
Player 1 Plays	Cooperate	<i>R</i> <i>Reward for Cooperation</i>	<i>S</i> <i>Sucker's Payoff</i>
	Defect (Cheat)	<i>T</i> <i>Temptation to Defect</i>	<i>P</i> <i>Punishment for failure to Cooperate</i>

A game is a PD game if (and only if) $T > R > P > S$ and $2R > (T + S)$. The values used in Axelrod's tournament were $R=3, S=0, T=5, P=1$.

1.5

TIT FOR TAT (hereafter, "TFT") was made famous by its repeated triumph against other strategies in the prisoners' dilemma tournaments summarized in Robert Axelrod's book, *The*

Evolution of Cooperation (1984). So influential was this work that it has served as the basis for hundreds of articles in which TFT reciprocity is used to explain various features of human and animal social behavior. However, this work did not consider in depth how the behavioral options of Axelrod's tournaments reflect the behavioral options of animals. We are specifically concerned that, in the Axelrod's tournaments, a player remained with its assigned partner until the number of moves stipulated for the game was completed. This constraint seems to violate an almost universal feature of animal psychology: *Animals tend to leave a situation that they find distasteful* (Thorndike 1911). Even theorists who favor use of the TFT strategy to explain animal behavior have implicitly conceded that this aspect of the tournament structure is quite artificial (Trivers 1971). When using TFT to explain natural social systems, theorists must speculate that partners recognize one another and remember what each did on previous occasions. That is, given that I am likely to move away from someone who acted against me, implementing a TFT strategy requires that I, at a minimum, recognize the individual when I meet them again and remember that they defected against me days, months, or even years ago. Maintaining this level of cognitive complexity is difficult, as neural tissue is the most energetically costly to maintain. But the TFT strategy in the simulated tournaments does not require complex cognition, only "knowledge" of what was done to it last, as it is bound to a single partner for prolonged periods. Thus, the strategy incurs none of the costs for maintaining cognitive complexity that would be suffered by an animal trying to implement a TFT strategy in nature.

Cooperation in "non-forced" games

1.6

Years ago researchers interested in economic decision-making became curious about how the option of leaving a market affected cooperation (Hirschman 1970). When implemented in iterative prisoners' dilemma games the "exit option" is meant to represent this freedom to walk away (Orbell and Dawes 1993). Yamagishi and Hayashi (1996) referred to such games as non-forced play paradigms and reviewed their history in depth. In addition to purely mathematical exploration, empirical studies of human behavior in prisoners'-dilemma type experiments have demonstrated that the presence of an exit option affects performance. For example, Boone and Macy (1999) showed that people do take advantage of the opportunity to leave partners, and cooperation does increase when participants in prisoners' dilemma tournaments are given the option of leaving.

1.7

However, these exit strategies were not well-suited to an evolutionary context and required extensive cognitive capacity on the part of agents. For example, Sherratt and Roberts' (1998) model required individuals to remember the outcome of every match they played, keep a running total of performance, and remember the partner for every previous match. Individuals then chose whether to sit-out a match if they encountered an undesirable player again. Using genetic algorithms, Sherratt and Roberts showed that such "choosiness" could evolve (assuming no cost of cognitive capacity). Similarly, Macy and Skvoretz (1998) implemented leaving strategies that required extensive cognitive ability. In their simulation, individuals displayed signs that could be attended to or ignored. If a strategy chose to "look at" its partner's signs, then the partner was accepted or rejected based on the average value of observed signs. If both partners found the other acceptable *a priori*, then the strategies played a single-game match in which they either cooperated or defected. Further, strategies could distinguish "neighbors" from "non-neighbors" and could respond to "greetings" by their partner. Macy and Skvoretz also used genetic algorithms so that any change in strategy between generations was the result of genetic "memory" for the value of signs, being a neighbor, and so on. The non-iterative scenario made it impossible to assess the success of their models against conditional prisoners' dilemma strategies (such as TFT). Similar levels of complication are found in numerous other simulations.

1.8

We seek to support the argument that cooperation can be explained by the most simple, evolutionarily feasible, cognitively undemanding mechanism of conditional association.

MOTH reciprocity

1.9

The most parsimonious conditional association model would simply have agents leave partners who defect against them but stay with partners who cooperate. That is, organisms must merely exhibit a tropism that leads them away from unpleasant stimuli. We called this strategy, My-way Or The Highway, or MOTH for short. Such a strategy requires far less "knowledge" than the previously considered leaving strategies. It requires the same amount of "knowledge" in simulated situations that TFT does, only needing to know what happened last. However, unlike TFT, MOTH requires no additional knowledge to be implemented in natural situations.

1.10

Far from requiring individual recognition and memory of long past events, implementing MOTH only requires that organisms respond to operant conditioning (i.e., not do things for which they were previously punished). Moreover, MOTH would be a powerful competitor because it would, in effect, sort through the population "looking for" other altruists. As a result, $p_{A,A} - p_{S,A}$ would increase steadily with successive iterations of the game. Note that MOTH switches partners following the same rule by which TFT switches roles. TFT is a conditional altruist that associates unconditionally; MOTH is an unconditional altruist that associates conditionally.

1.11

In order to determine the relative abilities of MOTH and TFT strategies, we performed a series of simulated tournaments.



Method

2.1

To explore these ideas, David Joyce prepared an evolutionary tournament in Java, in which players of eight different strategies played two-player matches consisting of a series of PD and/or altruist games. (See <http://aleph0.clarku.edu/~djoyce/Moth/battle/Battle.html>) At the end of each set of matches, the number of individuals playing each strategy was adjusted according to the results of the previous match. The eight strategies (See Table 4) varied in their ability to leave or stay with a partner and to cooperate or defect with the partner they have. Three of the eight are from Axelrod's tournament: ALL-C, ALL-D, and TFT. The five others were various different permutations of the two principles, leaving/staying, cooperating/defecting.

2.2

Because of the increasing use of NETLOGO (<http://ccl.northwestern.edu/netlogo/>) in social science modeling (Axelrod 1997; see also [Axtell, Axelrod, Epstein, and Cohen 1996](#)), Owen Densmore later "docked" our model in NETLOGO. Subsequent testing showed the two models to give identical results. Furthermore, preliminary explorations with the NETLOGO model showed that the differences between the outcomes of the 3, 0, 5, 1 prisoners' dilemma model and the altruist model are negligible. Thus, we report here only the outcome of experiments with the NETLOGO version of the model on standard 3, 0, 5, 1 prisoners' dilemma games. Readers may replicate the findings by running the model, which is available at <http://www.redfish.com/models/Moth.htm>.

The Tournament

2.3

The structure of the NETLOGO tournament is designed, insofar as possible, to mimic the structure Axelrod's evolutionary tournament. N players are assigned to S strategies at random and without replacement. A strategy is a set of rules that determines which of the options (cooperate, defect, or leave) a player will adopt under a specified set of circumstances. The tournament proceeds through a set number of matches. Each match begins with players being assigned partners at random, matches are made up of G iterations, or games, and the tournament is M matches long. Between matches, the N players are reassigned to the S strategies in proportion to their winnings in the previous match. The variables N, S, G, and M are all parameters that can be adjusted in our model. The payoffs of the game are also adjustable, although this report focuses, as we have said, on the 3,0,5,1 payoff matrix. Also adjustable are the initial numbers of the eight strategies, so that each strategy can begin with

anywhere from 0 to N players.

The Strategies

2.4

Eight strategies were available for comparison:

		Cooperate?	Leave?
All-C	Always cooperates; doesn't leave.	Yes	No
All-D	Always defects; doesn't leave.	No	No
Tit4Tat	Cooperates first; plays as partner did previous game; doesn't leave.	Conditional	No
Moth	Cooperates until partner defects, then leaves.	Yes	Conditional
Hit&Run	Defects, then leaves.	No	Yes
Santa	Cooperates, then leaves.	Yes	Yes
NasMoth	Always defects; leaves when partner defects.	No	Conditional
NNHRun	Defects twice then leaves; leaves immediately if partner defects on the first game.	No	Conditional

2.5

The variations between the eight strategies permitted us to explore the interactions between three of the Axelrod tournament strategies, ALL-C, ALL-D, and TFT, playing against strategies that could break partnerships any time during the match and resume playing with other partnerless players picked at random. The simplest version of this strategy was MOTH, described above, which is identical to ALL-C except that it left if its partner defected. Other variants were HIT AND RUN, which always defected and always left after the first round, SANTA, which always cooperated but always left after the first round, NASTY MOTH, which always defected but stayed around until its partner defected, and NASTY HIT AND RUN, which defected twice, but left immediately if its partner defected in the first game and then left unconditionally after the second game, no matter what its partner did.

2.6

The eight strategies can be thought of as varying simultaneously along three dimensions.

First Cooperate or First Defect Strategies: Because on the first game of a match, strategies have no information concerning the play of the other strategies, a first round decision must be made unconditionally whether to cooperate or defect.

Conditional or Unconditional Strategies: On the second and subsequent rounds, conditional strategies become possible, that is, strategies that were conditional on the other players' first round move. Some strategies make use of that information and others do not.

Leave-Stay or Cooperate-Defect Strategies: Conditional strategies can be further distinguished on the basis of what they did when their conditions were met.

Details of the Simulation

2.7

The players are paired off randomly. Simultaneously, each pair plays a single game. After that game, either player or both may decide to leave the partnership, in which case the partnership dissolves and both players are randomly re-paired with a partnerless player for the next game (they may be re-paired with the same partner); if neither wants to leave, the partnership

remains intact for the next game. At the end of a match, scores for each player are tallied, and the number of players of each strategy for the next match will be proportional to the total score for all the players of that strategy. In this way, the total number of players in the tournament is maintained at a constant. For instance, suppose that at the end of a match between 100 players, 45 ALL-C players and 55 ALL-D players. The 45 ALL-C players have a total score of 1000 while the 55 ALL-D players have a total score of 1500. In the next match there will be 40 ALL-C players and 60 ALL-D players, since $1000:1500 = 40:60$. If the proportion doesn't come out perfectly, the proportions will be rounded down and the remaining required players will be assigned to strategies with a probability determined by the same proportion.

2.8

Key variables in the simulations were numbers of games per match, the combinations of strategies in the simulation, and the starting frequencies of strategies in the simulation. As a general rule, 20 simulations were run to establish the value of parameter combination.

2.9

We conducted four types of tournaments: 1) MOTH vs. the 7 other strategies, 2) MOTH vs. the 3 non-leaving strategies, 3) MOTH and TFT separately vs. the unconditional strategies, and 4) MOTH vs. TFT invading the 6 other strategies. All tournaments began with equal numbers of all strategies, except type 4 tournaments, in which there were a smaller numbers of MOTH and TFT players initially.

 Results

Moth Versus the Other Seven Strategies

3.1

Simulations indicated that strategy success in these tournaments was most strongly affected by the number of games per match. A critical region was found between four and ten games per match. At shorter game lengths, first round defectors, such as HIT AND RUN, do well. At longer game lengths first round cooperators, including MOTH, do well.

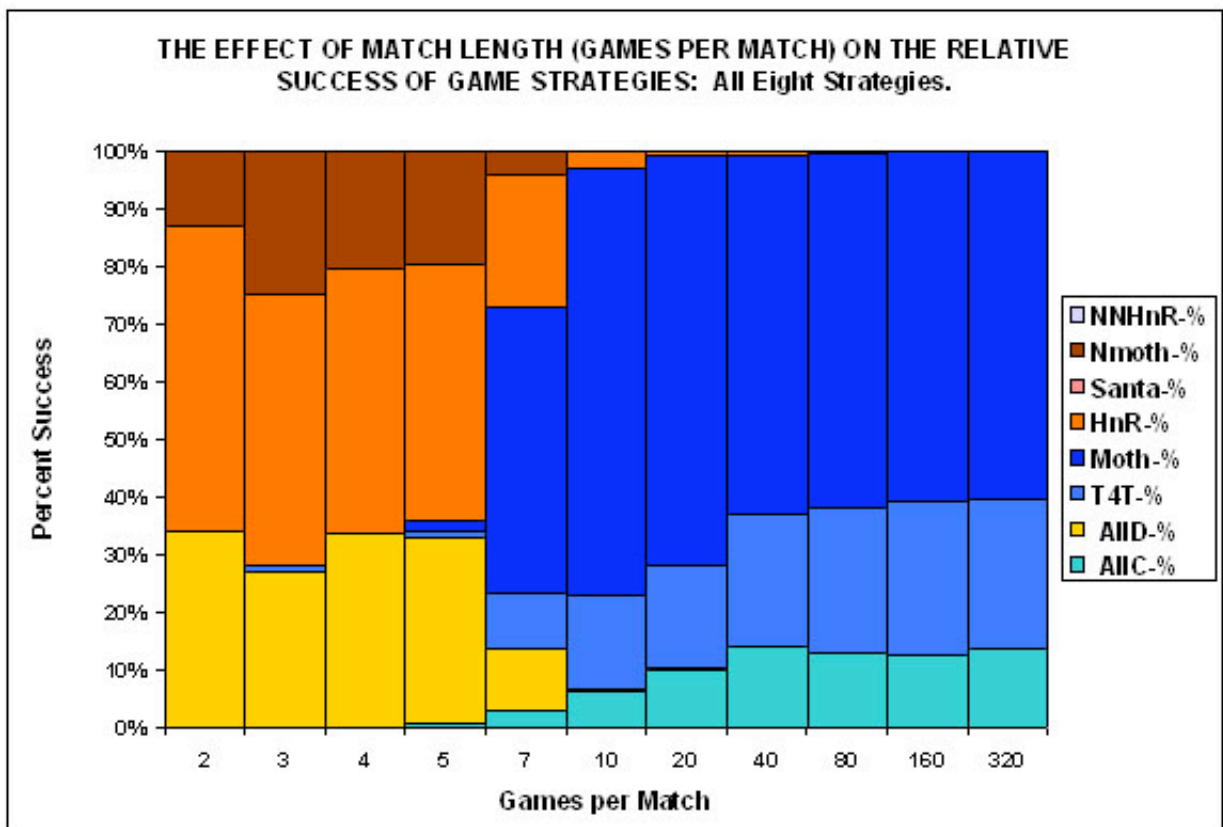


Figure 1.

3.2

These effects are illustrated in Figure 1, which shows the averages of 20 simulations. Each

simulation consisted of 40 matches, but with match lengths ranging from 2 to 320 games, varying on a roughly logarithmic scale. Each strategy started off with ten players. Each bar represents final percentage achieved by each strategy when matches consisted of the number of games indicated on the x-axis.

3.3

The graph dramatically shows the success of MOTH (and other first cooperator strategies) above a threshold of 10 games per match. It also shows the dependency of cooperative strategies, and MOTH in particular, on the game structure of the tournament. Interestingly, the graph also shows the power of leaving strategies, HnR and NasMOTH, against the pure defection strategy in matches with fewer games.

Moth Against Non-leaving Strategies

3.4

To test whether MOTH's success requires other conditional association strategies, we ran simulations pitting MOTH against the three non-leaving strategies from Axelrod's tournaments, TFT, ALL-D, and ALL-C. MOTH's performance was very similar to that found in the first simulation. Even in this limited field of competition, the outcome of the competition between first cooperator strategies and first defector strategies is highly dependent on match length. As match lengths increase, the first round cooperator strategies establish a strong collective majority over ALL-D at four games per match and nearly eliminate ALL-D in longer matches (See Figure 2). Amongst the two first cooperator conditional strategies, MOTH was consistently superior to TFT across parameters. It establishes a plurality at four games per match and a majority in longer matches. It should be noted, however, that MOTH does have a slightly weaker majority playing in this context than it does against the full field of eight strategies.

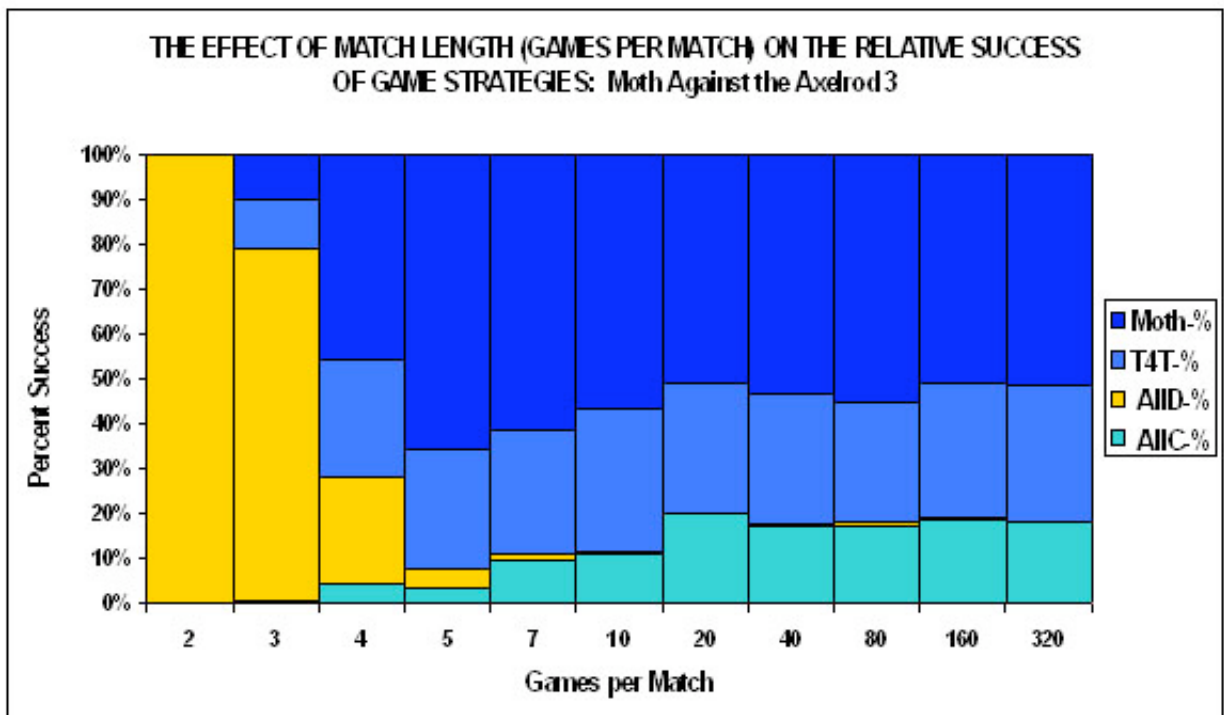


Figure 2.

MOTH and TFT vs. the Unconditional Strategies

3.5

Next, we ran TFT and MOTH separately against ALL-C and ALL-D to test whether the conditional leaving strategy fared better than the conditional defecting strategy in competition against the two unconditional strategies. MOTH had a slightly better record overall, achieving a majority at four games per match, whereas TFT's required at least five games per match. Further, MOTH averaged approximately 90% of the players in longer games, whereas TFT averaged 85%. (See Figures 3 and 4)

3.6

Interestingly, the fate of ALL-D in the two tournaments is slightly different. The combination of ALL-C and TFT eliminates ALL-D in any match of greater length than 20 games; however, the combination ALL-C and MOTH does NOT extinguish ALL-D players, even in the longest matches. If ability to drive competing strategies to extinction was considered desirable, this would be the only criterion we found by which TFT seem better than MOTH.

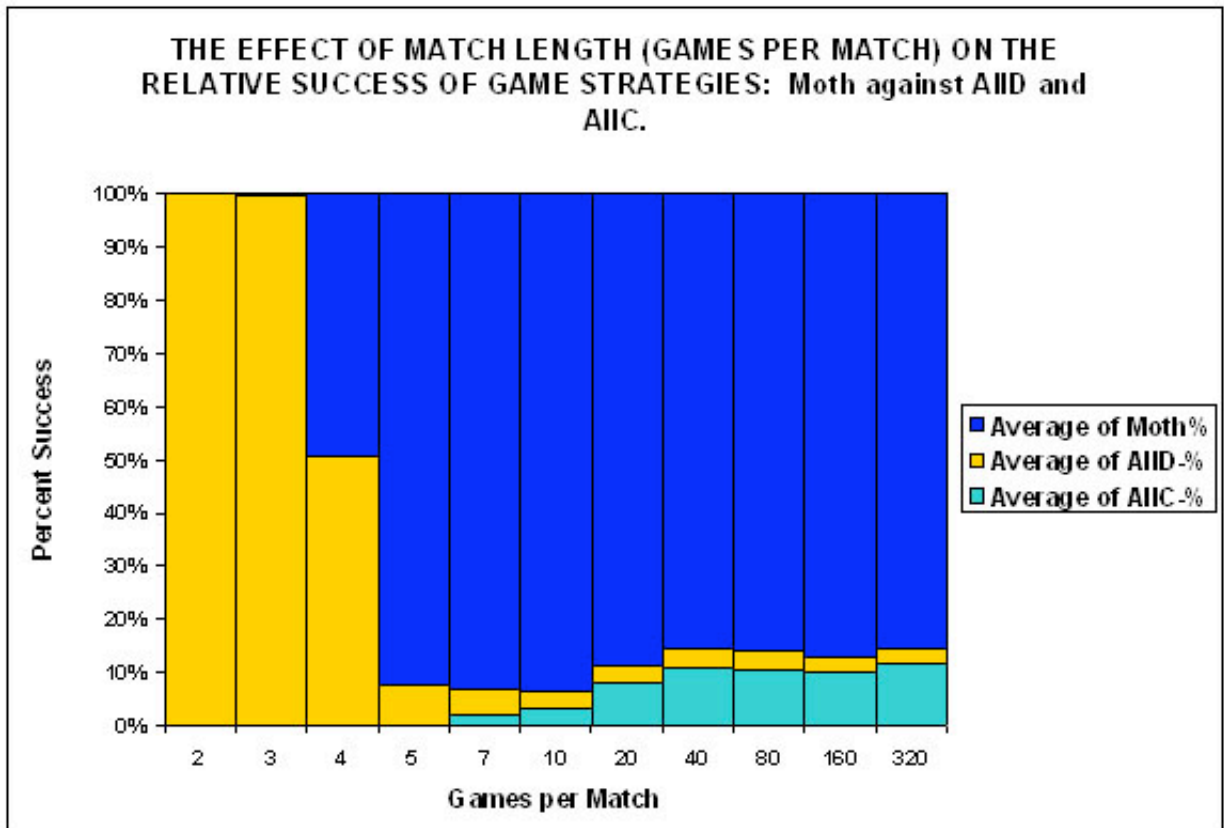


Figure 3.

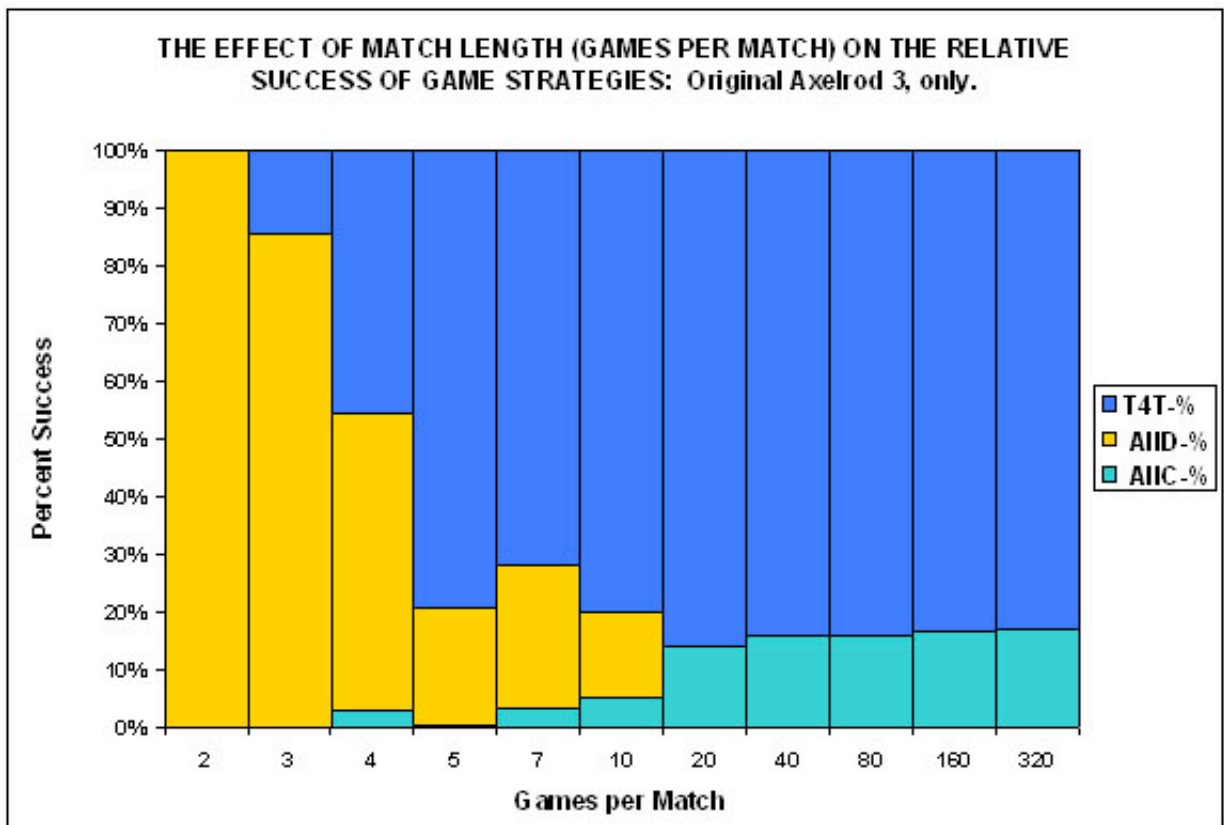



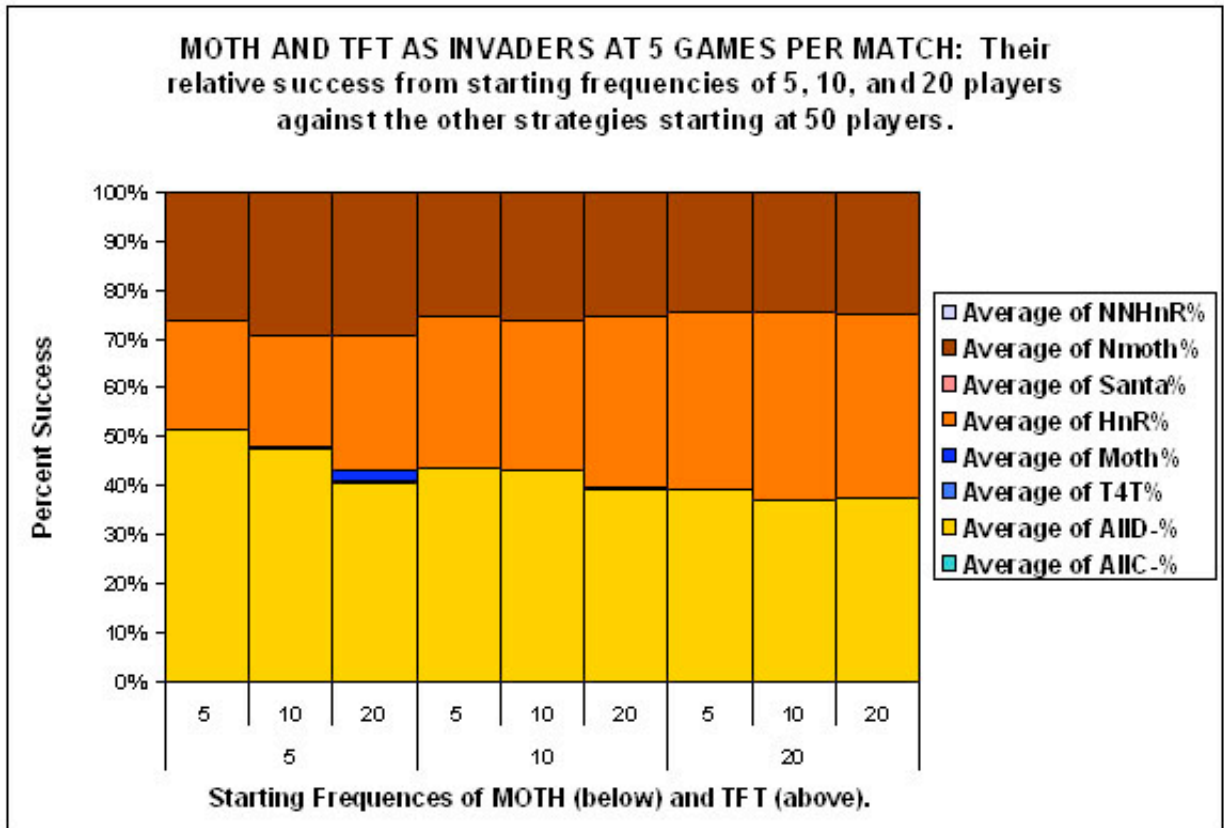
Figure 4.


3.7

One measure of the power of a strategy as an evolutionary competitor is its capacity to invade a population already well-occupied by players of other strategies. This capacity is significant because, as evolutionary change is usually understood, any novel type must begin its career at a great numerical disadvantage. To explore these issues, MOTh and TFT were run at starting frequencies of 5, 10, and 20 players against the six other strategies started at 50 players each. Because match length had proven to be such a potent variable above, these invasion simulations were run at 5, 10, and 20 games per match.

3.8

The results are presented in Figure 4, a-c, which affords three different sorts of comparisons.





Figures 4a, 4b and 4c.

3.9

Comparing the charts, 4a–c demonstrates, once again, the importance of match length as a variable. At the end of five–game matches, no first cooperator strategy is represented by more than one or two players. At the end of 10 or 20 game matches, first cooperator strategies have virtually eliminated all first defector strategies.

3.10

Each individual chart affords a comparison of the capacities of Moth and TFT to invade multi–strategic populations of 300 players. Note that (in all but one case) if TFT has surviving players at the end of a tournament, Moth has more. At 10 and 20 games per match, Moth outcompetes TFT at ratio of approximately ten to one. Further, in the face of competition from Moth, TFT is continually bested by ALL–C.

3.11

Comparing the heights of bars within each chart gives the effect of different starting frequencies on Moth and TFT success. The general outcome is that these effects are small. Both strategies do slightly better at higher starting frequencies, and both do slightly better when the starting frequencies of the other are lower. But these effects are modest compared to those described earlier.

Discussion

4.1

These results show that, when prisoners' dilemma games are modified so that both conditional association and conditional altruism are possible strategies, a conditional association outcompetes conditional altruism in a wide variety of conditions. To the extent that computer simulations can ever guide our understanding of human and animal sociality, this outcome implies that students of animal and human social behavior have been too quick to adopt conditional altruism as a principal solution to the problem of social evolution. The success of TFT in prisoners' dilemma tournaments has been largely an artifact of the structural limitations of those games. The results of Axelrod–type tournaments do not indicate the universality of conditional altruism as a principle of cooperation in situations where animals can leave their partners.

4.2

Beyond the robustness of conditional association in our computer tournaments, there are many reasons to prefer it as a foundation for animal and human societies. These include parsimony, generality, and heurism.

Parsimony

4.3

When attempting to explain animal sociality (rather than merely altruism), MOTH is more parsimonious than TFT. This is because TFT strategists must repeatedly come into contact with the same partners to gain a benefit from their strategy. Yet the logic of conditional altruism provides no mechanisms for such a social tropism. Any such mechanism would entail conditional altruists conditionally associating. That is, employing a strategy of conditional altruism *requires* partners to also conditionally associate. However, employing a strategy of conditional association does *not* require organisms to be conditional altruists. If one is an unconditional altruist and attaches oneself only to other altruists, no other principles are required to sustain sociality.

Generality

4.4

One does, of course, need some sort of principle to sustain conditional association. But that principle is readily available. Both classical and operant conditioning, widely dispersed in the animal kingdom, and both of which have their analogues in the plants (see [Barker 1993](#)), can be called upon to account for why creatures should stay where their surroundings please them. So, for instance, to use the example that Trivers ([1971](#)) relies on so heavily in his classical publication on reciprocal altruism, the cleaner fish and his clients need not recognize each other as individuals; to support the symbiosis, they need only to be territorial and to show up at the same place and the same time. So, in fact, what is usually taken as the classical example of conditional altruism is actually an excellent example of conditional association. This fact is clearly demonstrated by the success of the cleaner mimic, who does not need to mimic a specific individual. If he looks like ANY member of the cleaner species and shows up at the right time and place, he can feed off the flesh of the client species with impunity.

4.5

Individual recognition is a demanding threshold to meet in animal social interactions. As Konrad Lorenz's work has made abundantly clear ([Lorenz 1935](#)), the majority of creatures know each other only as disembodied sign stimuli, not as individuals. Individual recognition is not only a significant evolutionary accomplishment, but a late one. To build a theory of sociality on a principle that requires individual recognition is to limit the applicability of that theory to a relatively few organisms.

4.6

MOTH is a more general solution to the problem of reciprocation in another sense. MOTH, like kin selection and trait-group selection, is a conditional association strategy. It works by increasing the proportion of altruists that meet other altruists and decreasing the proportion of selfish individuals that meet altruists. Thus the success of MOTH holds out the possibility of unifying evolutionary theories of altruism into a single theory of conditional association being arrived at by three different mechanisms, rather than as three different kinds of "selection."

Final Words

4.7

The theory of conditional altruism has provoked much important research and some unexpected findings. Perhaps its most interesting success has been the manner in which Cosmides and Tooby ([1992](#)) and their colleagues have used it to rework our understanding of logical thinking and cognitive heuristics (c.f., [Buller 2005](#)). However, having conceded this benefit, we also must admit that the TFT principle may have been widely overextended. For instance, the most often cited example of TFT altruism, Wilkinson's vampire bats ([1984](#)), is arguably not an example of conditional altruism but rather an example of *unconditional* altruism

accompanied by conditional association.

4.8

MOTH has the value of raising research issues seldom considered by the reciprocal altruism literature. In natural coalitions, such as those observed in chimpanzees ([de Waal 1982](#)), there is a cost of breaking off a partnership. Partners may resist desertion and may inflict punishments on deserters. Furthermore, the possibility of leaving introduces a spatial dimension to social relations not reflected in Axelrod-type strategies. It is our intention, in future versions of this model, to introduce concepts of distance, interposition, and punishment for leaving.

4.9

Further work is surely needed to determine how MOTH fares against alternative strategies in ever more true-to-nature contests. Moreover, much work can be done to develop the mathematics of MOTH's success. Until that further research is done, the current research is sufficient to conclude that MOTH is capable of outcompeting TFT in many conditions. Given the seat of power granted to TFT by many evolutionary theorists, this should serve as a reminder that for the same reasons models are powerful aids to thought, they are also powerful limiters. Every one who uses a model should be deeply aware of the limitations of generality implied by its structure.



References

AXELROD R (1984) *The Evolution of Cooperation*. New York: Basic Books.

AXELROD R (1997) *The Complexity of Cooperation: Agent-Based Models of Competition and Collaboration*. Princeton, NJ: Princeton University Press.

AXTELL R, Axelrod R, Epstein J M, and Cohen M D (1996) Aligning simulation models: A case study and results. *Computation and Mathematical Organization Theory*, 1, 123–141.

BARKER G (1993), "Models of biological change: Implications of three studies of 'Lamarckian' change". In Bateson P P G et al. (Eds.), *Perspectives in Ethology, Volume 10: Behavior and Evolution*. New York: Plenum Press.

BOONE R T and Macy M W (1999) Unlocking the doors of prisoner's dilemma: Dependence, selectivity, and cooperation. *Social Psychology Quarterly*, 62, 32–52.

BULLER D J (2005) *Adapting Minds*. Cambridge, MA: MIT Press.

COSMIDES L and Tooby J (1992), "Cognitive adaptations for social exchange". In Barkow J H, Cosmides L and Tooby J. (Eds.), *The Adapted Mind: Evolutionary Psychology and the Generation of Culture*. New York: Oxford.

DE WAAL F (1982) *Chimpanzee Politics*. New York: Harper & Row, Publishers, Inc.

HIRSCHMAN A (1970) *Exit, Voice, and Loyalty*. Cambridge, MA: Harvard University Press.

LORENZ K (1935/1970), "Companions as factors in the bird's environment". In Lorenz K, (tr. Martin, R) *Studies in Animal and Human Behavior, Volume 1*. Cambridge, MA: Harvard University Press.

MACY M W and Skvoretz J (1998) The evolution of trust and cooperation between strangers: A computational model. *American Sociological Review*, 63, 638–660.

ORBELL J and Dawes R (1993) Social welfare, cooperator's advantage, and the option of not playing the game. *American Sociological Review*, 58, 787–800.

SHERRATT T N and Roberts G (1998). The evolution of generosity and choosiness in cooperative exchanges. *Journal of Theoretical Biology*, 193.

SOBER E and Wilson D S (1999) *Unto Others: The Evolution and Psychology of Unselfish Behavior*.

Cambridge, MA: Harvard University Press.

THORNDIKE, E L (1911) *Animal Intelligence, experimental studies*. New York: Macmillan.

TRIVERS, R (1971/1978) The Evolution of Reciprocal Altruism. *Quarterly Review of Biology*, 1971. As reprinted in Clutton-Brock T H and Harvey P H (Eds.), *Readings in Sociobiology*, pp 189–232. New York: WH Freeman.

WILKINSON G S (1984) Reciprocal Food Sharing in the Vampire Bat. *Nature*, Issue 308. pp. 181–184.

YAMAGISHI T and Hayashi N (1996) Selective play: Social embeddedness of social dilemmas. In Liebrand W and Messic D (Eds.) *Frontiers in Social Dilemmas Research*, Berlin: Springer.

[Return to Contents of this issue](#)

© [Copyright Journal of Artificial Societies and Social Simulation, \[2006\]](#)

