

Cross-level Interactions Between Conflict Resolution and Survival Games

Matthias Scheutz

Artificial Intelligence and Robotics Laboratory
Department of Computer Science and Engineering
University of Notre Dame
Notre Dame, IN 46556
mscheutz@nd.edu

Abstract

We investigate the interactions between conflict resolution and survival in multi-agent environments, where agents compete for resources. We define strategies for the “conflict game” and the “survival game,” analyze their properties, and compare their performance in agent-based simulations. The results demonstrate that inferior strategies in conflict games can be turned into superior strategies in the survival game if combined with the right parameters for foraging.

Introduction

Conflicts arise when two or more agents require a non-sharable resource at the same time. In a biological context, such conflicts (e.g., for food) occur as part of a larger competition for survival and procreation, where agents need to live long enough to procreate and pass on their genes.

In this paper, we investigate the interactions of strategies to resolve “conflict games” in the context of larger “survival games” in which conflicts occur. We start with a description of the conflict game and discuss different strategies for playing it. Then we discuss the setup of the larger survival game, define an agent model for it, and perform extensive simulation studies to examine the relationship between strategies for the two games. We show how the performance of poor strategies in the embedded conflict game can be completely reversed by small parameter changes in strategies for the larger game.

Conflict and Survival

Conflicts among agents can be viewed as games, where the participating agents can perform two actions: they can “continue” the game (i.e., fight for the resource) or they can “leave” the game (e.g., give up the resource and leave the scene). If two or more agents decide to continue, the game itself continues. Various one-round games like the *Prisoner’s Dilemma* (Axelrod and Hamilton, 1981) or the *Dove-Hawk Game* (Maynard-Smith and Price, 1973), and iterated games like the *Iterated Prisoner’s Dilemma* (Axelrod, 1984), or the *Sequential Assessment Game* (Enquist and Leimar, 1983) have been proposed to study conflicts as

they arise in nature. While game-theoretic models of conflicts and conflict resolution strategies are particularly studied in economics (e.g., (Hirshleifer, 1988; Hirshleifer, 2001; Garfinkel and Skaperdas, 2000; Reuveny and Maxwell, 2001)), there is also a vested interest in conflict resolution for multi-agent systems in distributed AI (e.g., (Klein, 1991; Sycara, 1991; Adler et al., 1990; Lander and Lesser, 1990)), where agents compete for shared resources.

We now describe the particular conflict and survival games employed and investigated in this paper, which differ from the above models in several important ways.

Conflict Games and Conflict Resolution Strategies

Many real life incarnations of the above games occur as special iterated games of what is called *Dynkin games* or *stopping games* (Shmaya et al., 2003; Touzi and Vieille, 2002), where players with finite budgets have to pay a cost C_P (“the cost for playing”) ¹ in every round; they also have to pay a cost C_L (“the cost for leaving”) for quitting. The payoff is only determined when at most one player is still in the game, at which point the resource’s value (B_P for “benefit of playing”) is added to that player’s budget. Losers might still get other resources (e.g., ones that are not contested) later with a certain probability—we call this the “benefit of leaving” B_L and assume that the utility of winning outweighs the utility of losing in a one-round game (i.e., $B_P + C_P > B_L + C_L$). We will also take another important aspect of many real-world conflicts into account (see below), that of signaling—truthfully or not—(aspects of) behavioral propensities, which we construe as the probabilities with which agents continue a game.²

A *strategy for the conflict game* is defined as the probability P_A that agent A will continue a game. Each conflict can then be described by a sequence of interactions (or “rounds”) which continue as long as at least two players continue the

¹This cost corresponds to the basic bet in poker games or the energy spent for fighting for a period of time.

²Many animals that are about to fight for resources such as food, mates, or territory typically show displays of aggression, e.g., see (Adamo and Hanlon, 1996; Hofmann and Schildberger, 2001).

game.³ Intuitively, the expected utility of a conflict game should be independent of its duration. To see this, first define an expected utility $U_{S,O}(n)$ for player S in a two-player game with players S and O after n rounds:

$$\begin{aligned} U_{S,O}(n) = & P_S \cdot P_O \cdot (C_P + U_{S,O}(n+1)) + \\ & P_S \cdot (1 - P_O) \cdot (B_P + C_P) + \\ & (1 - P_S) \cdot P_O \cdot (B_L + C_L) + \\ & (1 - P_S) \cdot (1 - P_O) \cdot (B_L + C_L) \end{aligned}$$

Note that in case both players continue in round n , the utility of both continuing for S is given in terms of the utility of the $n+1$ round game, otherwise the respective utilities can be directly computed based on player's strategies and the benefits/costs alone.

Let $Rest = U_{S,O}(n) - P_S \cdot P_O \cdot U_{S,O}(n+1)$, which is a number independent of $U_{S,O}(n)$ for any n . Then combining $U_{S,O}(n) = Rest + P_S \cdot P_O \cdot U_{S,O}(n+1)$ and $U_{S,O}(n+1) = Rest + P_S \cdot P_O \cdot U_{S,O}(n+2)$, we get $U_{S,O}(n) = Rest + Rest \cdot P_S \cdot P_O + (P_S \cdot P_O)^2 \cdot U_{S,O}(n+2)$. In general, we get $U_{S,O}(n) = Rest \cdot \sum_{k=0}^m (P_S \cdot P_O)^k + (P_S \cdot P_O)^{m+1} \cdot U_{S,O}(n+m+1)$, hence the utility of the overall game $U_{S,O} = \lim_{m \rightarrow \infty} U_{S,O}(n)$ (for any n). Since $\sum_{k=0}^n (P_S \cdot P_O)^k$ can be written as $\frac{(P_S \cdot P_O)^{n+1} - 1}{P_S \cdot P_O - 1}$ and $\lim_{m \rightarrow \infty} (P_S \cdot P_O)^{m+1} = 0$, and assuming that $P_S \neq 1$ or $P_O \neq 1$ (otherwise the game has no value), we get $U_{S,O} = \frac{1}{1 - P_S \cdot P_O} \cdot Rest$. Note that $U_{S,O}$ has two maxima at $P_S = 0, P_O = 1$ and $P_S = 1, P_O = 0$, and that because long games only incur costs (since for each round C_P has to be paid) and have no added benefits, the best game is a one-round game.

So what is a good strategy for the conflict game? Player S can reason based on player O 's choice in the following way: suppose P_S is fixed, then player O should play $P_O = 1$ if $(1 - P_S) \cdot B_P + C_P > B_L + C_L$.⁴ Conversely, if $(1 - P_S) \cdot B_P + C_P < B_L + C_L$, then player O should play $P_O = 0$. For $(1 - P_S) \cdot B_P + C_P = B_L + C_L$ player O will be indifferent between $P_O = 0$ and $P_O = 1$ (any value in $[0, 1]$ will pay $B_L + C_L$). Similarly, player S will be indifferent if player O plays P_O such that $(1 - P_O) \cdot B_P + C_P = B_L + C_L$. Therefore, the pair (P_O, P_S) such that $(1 - P_S) \cdot B_P + C_P = (1 - P_O) \cdot B_P + C_P = B_L + C_L$ is a Nash equilibrium.

It may be tempting for player S to play $P_S = 1$ as this promises the highest payoff. However, in the absence of information about P_O this can lead to prolonged games, e.g., if both S and O decide to continue, where the "payoff" per round is only the cost C_P with no benefits to either player.

If information about the likely behavior of the other participant is available, then better decisions can be made that

³While the sequence is finite for players with finite budgets, it is not in general finite as new players can join the game any time. For all practical purposes, however, it can be considered finite.

⁴In general, if an expression α (e.g., $(1 - P_S) \cdot B_P + C_P$) is greater than an expression β (e.g., $B_L + C_L$), then $x \cdot \alpha + (1 - x) \cdot \beta$ has a maximum for $x = 1$ for all $x \in [0, 1]$.

ensure short and for some strategies one-round games. Contestants can, for example, *indicate their disposition* to continue the game. Based on these apriori probabilities, contestants can gauge their opponents' willingness to continue and figure it into their decisions. For example, an agent could lower its probability of continuing if the other agent's probability is higher than its own, and otherwise increase it. Such an adjustment pays off, because the likelihood of long encounters is reduced and the more likely outcome—that the agent with the higher aggression level will win—is reinforced. One possible way to define such an adjustment is given by the following equations, we will call the *Social Rule*: S increases its probability P_S to continue the game by a factor of $\frac{1 - P_S}{P_S} \cdot (P_S - P_O)$ if $P_S > P_O$ and decreases it by a factor of $\frac{P_S}{P_O} \cdot (P_O - P_S)$ if $P_S < P_O$, otherwise it remains the same. This rule will increase or decrease S 's probability to continue in proportion to the probability of the other player.

The limit case of the social rule (i.e., of applying the social rule to itself repeatedly) is the following strategy, which we will call the *Rational Rule*: S plays 0 if $P_S \leq P_O$ and 1 otherwise. This strategy is based on the assumption that contestants do not know the actual value of B_P, B_L, C_P , and C_L , hence cannot compute whether $(1 - P_O) \cdot B_P + C_P > B_L + C_L$ (otherwise they could play pure strategies in the first place). Note that this strategy if played by both players guarantees one-round games, which are not guaranteed if any of the Nash equilibria are played: suppose that both agents indicate probabilities p such that $(1 - p) \cdot B_P + C_P > B_L + C_L$, then they both should continue, which will lock them into a prolonged game with only costs. The rational strategy, on the other hand, prevents this by forcing the contestant with the lower probability to "give up" right away and still get at least $B_L + C_L$ (instead of the winner's $(1 - p) \cdot B_P + C_P$).

Of course, this rule is not *fair* in that repeated encounters between the same two individuals will lead to the same outcome, i.e., the same individual will win over and over again. In a group that means that the agent A with highest P_A will have $n \cdot ((1 - p) \cdot B_P + C_P)$ payoff after n encounters and the bottom P_A will have $n \cdot (B_L + C_L)$.

A fair way to distribute resources given that the utility function has a maximum for both $P_S = 1, P_O = 0$ and $P_S = 0, P_O = 1$, is to alternate between getting $B_P + C_P$ and $B_L + C_L$, which allows each player to get the average payoff $(B_P + C_P + B_L + C_L)/2$ every turn. We will call this the *Turn-Taking Rule* (cp. to (Neill, 2003)). Of course, for games with more agents, turn-taking is not as straight-forward, for it is possible that both agents in a conflict took their turn at losing in the previous encounter, so both will expect to win in the current encounter. One way to address this problem is to let each agent keep track of the number of games it won out of all the games it participated in and adapt its strategy depending on this fraction: every time an agent loses a game, the probability to continue games in the future is increased, every time an agent wins a game, the probability to continue

is decreased. Formally, this idea can be defined as follows:

Definition [The *Turn-Taking Rule*] Let m be the mixed strategy (i.e., probability of “continue”) played by agent A (in the absence of turn-taking) and let c be the current strategy ($m = c$ if the turn-taking rule has never been applied). Then the turn-taking rule $TT^+(c)$ (mapping values in the interval $(m/2, m)$ into $(m, (1 - m)/2)$) can be defined for losses:

$$\begin{aligned} TT^+(c) &:= c + (1 - c)/2 \text{ for } c \geq m \\ TT^+(c) &:= 2c \text{ for } c \leq m/2 \\ TT^+(c) &:= m + (2c - m)(1 - m)/2m \text{ otherwise} \end{aligned}$$

Similarly, $TT(c)^-$ (mapping $(m, (1 - m)/2)$ into $(m/2, m)$) can be defined for wins:

$$\begin{aligned} TT^-(c) &:= c - (1 - c) \text{ for } c \geq m + (1 - m)/2 \\ TT^-(c) &:= c/2 \text{ for } c \leq m \\ TT^-(c) &:= m/2 + m(c - m)/(1 - m) \text{ otherwise} \end{aligned}$$

This turn-taking rule has several nice properties, in particular, that the n -fold application in one direction followed by an n -fold application in the other direction yields the original probability:

Corollary 1. Let $TT^{+,n}(c)$ ($TT^{-,n}(c)$) denote the n -fold (recursive) application of TT^+ (TT^-) to c . Then $TT^{-,n}(TT^{+,n}(c)) = c$ and $TT^{+,n}(TT^{-,n}(c)) = c$ (where $TT^{+,1} := TT^+$ and $TT^{-,1} := TT^-$).

More importantly, the turn-taking rule combined with the rational rule is *fair* in a clearly specified sense: the difference between wins and losses in a group of competing agents is bounded by $\text{int}(n/2) + 1$ (where $|P| = n$ is the population size of the competing agents):

Proposition 2. For an agent population P of size n playing both the rational rule and the turn-taking rule, there exists a constant d such that the difference between the number of wins and the number of losses taken from the whole population is bounded by d for an arbitrarily long tournament (i.e., $|\text{wins} - \text{losses}| < d$ for all agents in P).

Hence we expect agents playing the combined turn-taking and rational rule to do well in the survival game, which we describe next.

Survival Games and Foraging Strategies

Conflict games are played as part of *survival games* (Barash, 2003; Maynard-Smith, 1982) in a spatially continuous two-dimensional plane, in which agents are situated. At any time, they can choose among four actions depending on their current context: to *collect* a resource that is present in their vicinity, to *duplicate* and produce an identical copy of themselves if their budgets exceeds C_D (the “cost of duplication”)⁵, to *play* the conflict game if at least one other agent

⁵The exact copy of an agent with budget B_{init} will be added to the survival game in the vicinity of the agent and can be at the earliest produced after D_{min} cycles from the point the agent has joined

is present, or to *forage* for resources (which is the default action).

Resources with a fixed benefit $B_P = 800^6$ are randomly distributed within a given 1800 by 1800 square unit area in the environment. New resources appear at a frequency $R = 1/\text{cycle}$ at a random location within this area and remain there until collected by an agent (agents can empty resources and add their value to their budget by moving over them; empty resources will be removed). Agents move in any direction at a speed $S_A = 4/\text{cycle}$, for which they pay a quadratic “foraging cost” $C_F = S_A^2/\text{cycle}$. They also pay a base cost $C_B = 1/\text{cycle}$.

Agents can perceive resources and other agents within their sensory range $SR = 300 \text{ units}$. They decide the direction to move in based on the sum of all “force vectors”, i.e., each agent A computes the sum of all vectors from its current position to the position of another agent or resource scaled by the square of the distance and multiplied by a gain factor:

$$\sum_n r \cdot R(n) + \sum_m s \cdot S(m) + \sum_k o \cdot O(k)$$

where $R(n)$ is scaled vector to the n -th resource, $S(m)$ the scaled vector to the m -th agent of the same kind as A , $O(k)$ the scaled vector to the k -th agent of a kind different from that of A , and r , s , and o are multiplicative constants called “resource gain,” “same gain,” and “other gain,” respectively (all vectors are scaled by the square of the distance between A and the respective item). This way of determining where to go has been studied extensively in autonomous agent research, in particular, in behavior-based robotics (Arkin, 1989) (but also in biology): agents will follow gradients that lead them towards areas with close-by or high densities of resources.

Agents can engage in conflicts when they are within conflict distance $CD = 35 \text{ units}$, at which point both agents signal their *a priori* probability P_A to continue and then, depending on their conflict strategy, make a decision based on all the inputs they received from other agents (we use $C_P = 50$ and $C_L = 350$, which is the average cost of running away from the site of conflict for a random number of cycles between 5 and 10 at a speed of 7 in order to leave the conflict range and end the game). Typically, this will allow losers to pick up other gradients, and thus to avoid having to compete with the same agents again (e.g., by rejoining an ongoing game or by restarting the previous game), which they otherwise would have given their previous attraction.

While different *strategies* can be defined for the survival game, we will focus on a *fixed strategy* for all agents that depends on only two foraging parameters to be defined in

the survival game (the “reproduction delay”). Note that agents will not be able to identify other agents as their offspring beyond recognizing them as one of their kind.

⁶Note that for space reasons we already add the particular values used in our simulations to the description of the parameters.

Section (i.e., agents use different forage actions or strategies), to reduce the number of free parameters: agents will always collect collectible resources, always procreate whenever possible, and always enter a conflict game if another agents are present in their vicinity.

Assuming that the resource gain (r) is fixed for all agents, we can then define an infinite set of foraging strategies based on the pair of gains (s, o) . Suppose n agents A_1, A_2, \dots, A_n participate in the survival game. Let C be the $n \times n$ matrix given by the probability $c_{i,j}$ that agents A_i and A_j will enter a conflict in the survival game, where the foraging parameters $o_i = s_i = 0$ for all agents. We can then define an $n \times n$ expected utility matrix \mathcal{E} by $e_{i,j} := U_{P_i, P_j}$, where U_{P_i, P_j} is the expected utility for agent A_i in a conflict game with agent A_j and P_i and P_j are the probabilities of continuing in the conflict game for A_i and A_j , respectively. Then the overall expected utility UC of conflict games for agent A_i is given by $UC_i = \sum_{j=1}^n c_{i,j} \cdot e_{i,j}$. Moreover, let $R_i(s_i, o_i)$ be the apriori probability of agent A_i getting a resource with net utility B_F through foraging depending on parameters o_i and s_i without any intervening conflicts (here B_F is obtained from the resource utility B_P minus the average foraging cost to get it, i.e., $k \cdot (C_F + C_B)$ for some $k > 0$). The overall expected utility \bar{U}_i for A_i in the survival game is then given by $R_i \cdot B_F + UC_i$. Note that performance in the survival game depends on the three parameters s_i, o_i , and P_i and that agents can influence their overall expected utility \bar{U}_i both by changing P_i (which will affect only UC_i) and/or their foraging parameters o_i and s_i (which will affect R_i and can affect UC_i , because UC_i depends on C)⁷. The fascinating aspect of this mutual dependence is that good performance in the conflict game does not automatically translate into good performance in the survival game as it is possible for agents to compensate for poor strategies P_i in the conflict game by virtue of changing the o_i and s_i parameters.

Note that while it is true that because all agents play a fixed strategy (only parameterized by o_i and s_i) in the survival game, agents with higher \bar{U} are more likely to have offspring, this does not automatically translates into higher average numbers of survivors in the long run. Whether an agent kind (as determined by the three parameters) is going to be successful will depend on how adding offspring to (or removing agents who ran out of their budgets from) the survival game changes \bar{U}_i for all involved agents. Consider, for example, a survival game that starts with two different, equally sized population of agents A and B , where A has $o_A = s_A = 0$ while B has $s_B > 0$ and $o_B = 0$ (i.e., B agents seek out conflicts with each other). Furthermore, suppose that B agents always play $P_B = 1$, while A agents play $P_A = 0.5$. Then $e_{i,j} < e_{j,i}$ for agents i of type A and agents j of type B . Hence, B agents gain from having conflicts with A agents, but will be locked into long conflicts

among their own, which will only end when one agent has no budget left. Everything else being equal, $\bar{U}_B > \bar{U}_A$ in the beginning, hence B agents will have more offspring than A agents, which changes the agents distribution in favor of type A agents. However, this increases the chances that B will enter conflicts with B agents, which are detrimental and subsequently decreases \bar{U}_B , possibly to a point where $\bar{U}_B < \bar{U}_A$. Performance in the survival game, therefore, depends not only on \bar{U} , but also on how \bar{U} changes over time.

We will in the following experimentally investigate the specific relationships between the conflict and survival games with respect to tradeoffs between *adaptive changes* in strategies in the conflict game with no foraging bias (i.e., $s = o = 0$) compared to fixed strategies in the conflict game with foraging bias (i.e., $s \neq 0$ and/or $o \neq 0$).

Simulation Experiments and Results

We defined 8 basic agent types, which are divided into four *asocial* and four *social* agents, depending on whether or not they take signals of other agents into account in their decision making. We also distinguish *adaptive* from *non-adaptive agents* depending on whether or not they use the turn-taking rule.

Asocial agents consist of three non-adaptive agents called *timid* (playing “always leave”), *Aggressive* (playing “always continue”), and (prototypically) *Asocial* (playing a mixed strategy), as well as *Asocial adaptive* (playing a mixed strategy with the turn-taking rule)⁸. Social agents are either adaptive or non-adaptive (prototypically) *Social* (playing the social rule) or *Rational* (playing the rational rule).

All experimental runs use 50 agents (25 each of two kinds) initially placed in random locations within the 1800 x 1800 resource area together with 50 randomly distributed resources. Their initial energy was set to 2000 and P_A (for continuing a game) for agents with mixed strategies was Gaussian distributed with spread 0.125 around 0.5. Each experiment consists of 40 experimental runs of 10000 cycles each after which the number of survivors is determined. The average number of survivors across all 40 runs (together with its 95% confidence interval) is used as *performance measure*.

Previously, we had determined the performance of all 8 agent kinds based on their strategies for the conflict game, i.e., by putting each kind in direct competition with each other kind with $s = o = 0$ (reference suppressed). The relevant outcome for the following experiments is that pure strategies are, as expected, bad choices: Aggressive agents perform worst (there are no survivors against any other agents) and Timid agents are second worst outperforming only Aggressive agents. Rational Adaptive agents, on the other hand, perform best with no survivors of any competing agent kind. In general, social agents outperform asocial

⁸We use capital letters for names of agent kinds to be able to distinguish them from attributes, e.g., “Asocial” refers to a kind, whereas “asocial” refers to an attribute.

⁷Note that it is possible to change R_i without affecting any $c_{i,j}$.

agents, except for the Asocial adaptive agents, which only perform worse than the Social and Rational adaptive agents.

Asocial Agents with Non-zero S-Gains

To investigate the influence of non-zero s-gains and o-gains of agents with low overall performance in the conflict game, we compare Asocial agents to Asocial adaptive agents and to all social agent kinds (with zero gains). For each of the 5 pairings, we ran experiments varying the s-gain from -10 to -50 and the o-gain from 10 to 50, both in steps of 10, thus yielding a total of 25 experiments (of 40 experimental runs each).⁹ Figure 1 depicts the resulting performance spaces for all 5 pairings (i.e., the average number of survivors of each kind for each of the 25 experiments). As can be seen, non-zero gains have a tremendous influence on the performance of an agent: Asocial agents outperform all other agent kinds (except for Rational adaptive agents) in most points in the gain parameter space except for least negative s-gains and high o-gains. Parameters in these areas cause Asocial agents to exhibit a slight avoidance of their own kind while strongly seeking conflicts with agents of other kinds.

To test the extent to which non-zero s-gains and o-gains could compensate for poor performance in the conflict game, we tested timid and aggressive agents, which have the poorest performance in the conflict games, with non-zero s-gains (while keeping the o-gains at 0). For each pairing, we conducted two experiments where the initial positions of the first agent kind in the first experiment became the initial conditions of the second in the second and vice versa to ensure that the results are balanced and that initial positions do not contribute to the results. Figure 2 shows the results of this comparison for $s = -10$ and $s = -20$. While the Timid agents only show improvement of their performance against the Asocial adaptive (for $s = -10$ they can coexist with Asocial adaptive agents, for $s = -20$ they outperform them), Aggressive agents show already a major improvement against all other kinds for $s = -10$: they have survivors (contrary to their performance for $s = 0$). Yet, for $s = -20$ they manage to outperform all other agents only leaving an insignificant number of surviving Rational adaptive agents in some runs.

At first glance, this enormous performance boost is quite surprising. After reflection, however, it is clear why the aggressive strategy combined with negative s-gains works so well: it ensures that Aggressive agents will mostly compete with other agents, where their chances of winning and getting a benefit are very high (effectively they only lose when they run out of their budget), while avoiding getting locked into long games with their own (compare this to the discussion in the Section “Survival Games and Foraging Strategies”). Note that other agents will still on average compete

⁹The ranges were chosen as positive s-gains and negative o-gains lead to poor strategies (e.g., competitions with one’s own kind and to avoidance of other kinds).

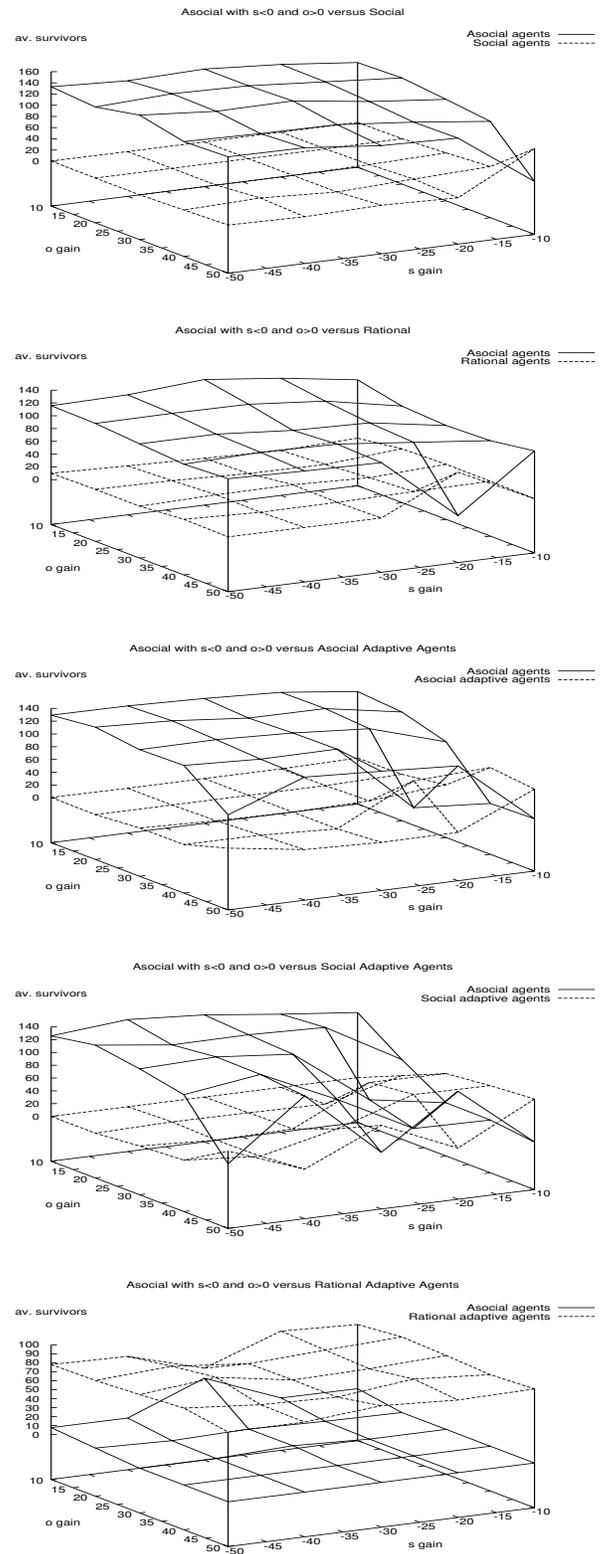


Figure 1: The average number of asocial survivors compared to other agent kinds varying the asocial agents’ s-gains (from -10 to -50) and o-gains (from 10 to 50), both in steps of 10.

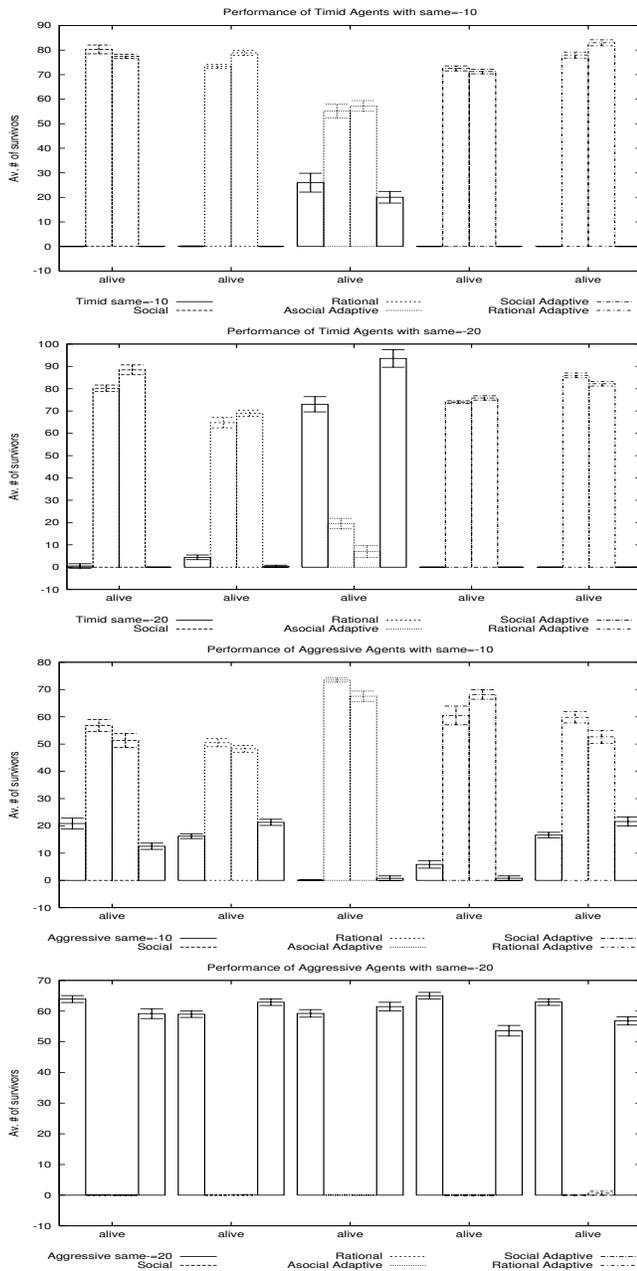


Figure 2: The average number of timid agents and aggressive survivors compared to non-asocial agent kinds and the Asocial adaptive agent (shown for “dual experiments” with initial positions swapped—see the text for an explanation).

the same time with each other as they would have normally. This example demonstrates a case where a strategy that by itself is not beneficial at all can be turned around into a powerful strategy when applied selectively.

Analysis

Negative s-gains of agents i and j reduce $c_{i,j}$ and can thus improve UC_i and UC_j when $e_{i,j} < 0$ (as is the case with Aggressive agents). A large enough improvement can then change an agent kind's A relative overall \overline{U}_A compared to another agent kind's B overall \overline{U}_B for a given population of A and B agents in the survival game such that $\overline{U}_{A_{s < k}} > \overline{U}_B$, while $\overline{U}_{A_{s \geq k}} < \overline{U}_B$ for some $k < 0$ as demonstrated by the above simulation results. Note, however, that a higher expected utility for a given distribution of agents from two kinds by itself does not guarantee that it will be preserved over time as the composition of the population in a survival game changes as a result of duplication and removal of agents. This is where another property of negative s-gains is critical: everything else being equal, they guarantee the relative utility ordering as long as the number of agents of kind A does not get so large that agents are “forced” (based on their summed “force vectors”) into conflicts (because there is not enough space in the environment for them to avoid each other—for space reasons we will not be able to develop the details here). Note, however, that as with Aggressive agents the utility will decrease as a result of the smaller numbers of B agents, from which A agents have to gain most. Eventually, a population of only A agents will be left with $\overline{U}_{A_{s < k}} = R_i \cdot B_F$, where negative s-gains prevent all conflicts.

Discussion

Without modifications of the parameters in the survival game, social strategies are better than asocial strategies, and adaptive strategies that attempt to distribute resources equally among agents are better than non-adaptive ones. The best strategies are combinations of adaptive and social strategies. If, however, parameters in the survival game can be changed that affect the foraging strategy of agents, then the poorest performer in the conflict game can be utilized to beat many social strategies. These results point in an interesting direction: what are the exact relationships between parameters in embedded games and those in larger games and how can these interdependent parameters be utilized to improve performance? Or more generally, what are the interactions between different spatio-temporal levels of organization? E.g., in the present study, we compare individual conflicts and strategies for resolving them to competition of species and strategies for procreation and survival, but other kinds of decision-making or control processes that operate at different spatio-temporal levels could be considered as well (e.g., interactions between the control processes at the level of immune system and the organismal level).

We believe that one particular lesson to be learned from this study is that *the best strategy in an embedded game does not always lead to the best performance in a larger game*, for it is possible that the best strategy implicitly fixes parameters in the larger games that poorer strategies nevertheless can adjust to their advantage. In the context of a survival

game in nature this could mean that there might be situations in which animals never play the optimal strategy in an embedded game, simply because it leads to lower overall performance in the larger game of survival. Consequently, embedded games with such interdependencies on their larger games cannot and should not be considered in isolation.

These kinds of interactions are particularly important in the context of artificial life, where theories of dynamically forming hierarchies of different levels of organizations and of emergence of properties at higher levels of complexities are still in the early stages. A better understanding of the dependencies and influences between parameters of strategies at different levels (e.g., between the embedded and the larger game) can help in the formulation of a viable theory of dependency and possibly *causation* – exactly how upper levels “cause” lower levels to behave in particular ways (e.g., by virtue of constraining their parameter space, state space trajectories, etc.) and how lower-levels, in turn, bring about behaviors and properties at higher levels. The dynamics of such interactions are critical for a thorough analysis of the behavioral repertoire of a complex multi-level system and are critical for determining good if not optimal control principles (e.g., strategies in conflict games).

Conclusion

We have developed a game-theoretic framework for investigating the relationship of strategies for survival games with embedded conflict games. Specifically, we studied the interaction between parameters for foraging strategies and adaptive vs. social strategies in the conflict game, and showed that and why poor strategies in the conflict game can be turned in successful strategies in the survival game if combined with the right set of foraging parameters. We intend to extend this work to isolate general conditions under which suboptimal strategies in embedded games can be improved based on parameters in larger games, which we believe will have important implications for understanding and designing efficient ways of resolving conflicts in embedded conflict games played by artificial agents.

References

Adamo, S. A. and Hanlon, R. T. (1996). Do cuttlefish (cephalopoda) signal their intentions to conspecifics during agonistic encounters? *Anim. Behav.*, 52:73–81.

Adler, M., Davis, A., Weihmayer, R., and Worrest, R. (1990). Conflict-resolution strategies for non-hierarchical distributed agents. In *Distributed artificial intelligence: vol. 2*, 139–161. Morgan Kaufmann.

Arkin, R. C. (1989). Motor schema-based mobile robot navigation. *Int. Journal of Robotic Research*, 8(4):92–112.

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

Axelrod, R. and Hamilton, W. D. (1981). The evolution of cooperation. *Science*, 211:1390–1396.

Barash, D. P. (2003). *The Survival Game: How Game Theory Explains the Biology of Cooperation and Competition*. Henry Holt & Company, Inc.

Enquist, M. and Leimar, O. (1983). Evolution of fighting behavior: decision rules and assessment of relative strength. *Journal of Theoretical Biology*, 102:387–410.

Garfinkel, M. R. and Skaperdas, S. (2000). Conflict without misperceptions or incomplete information: How the future matters. *The Journal of Conflict Resolution*, 44(6):793–807.

Hirshleifer, J. (1988). The analytics of continuing conflict. *Synthese*, 76:201–233.

Hirshleifer, J. (2001). The macrotechnology of conflict. *The Journal of Conflict Resolution*, 45(6):773–792.

Hofmann, H. A. and Schildberger, K. (2001). Assessment of strength and willingness to fight during aggressive encounters in crickets. *Animal Behavior*, 62:337–348.

Klein, M. (1991). Supporting conflict resolution in cooperative design systems. *IEEE Transactions on Systems, Man and Cybernetics*, 21(6):1379–1390.

Lander, S. and Lesser, V. (1990). Conflict resolution strategies for cooperating expert agents. In *International Working Conference on Cooperative Knowledge-based Systems*, Keele University.

Maynard-Smith, J. (1982). *Evolution and the Theory of Games*. Cambridge University Press, Cambridge.

Maynard-Smith, J. and Price, G. (1973). The logic of animal conflict. *Nature*, 146:15–18.

Neill, D. (2003). Cooperation and coordination in the turn-taking dilemma. In *14th International Conference on Game Theory at Stony Brook*.

Reuveny, R. and Maxwell, J. W. (2001). Conflict and renewable resources. *The Journal of Conflict Resolution*, 45(6):719–742.

Shmaya, E., Solan, E., and Vieille, N. (2003). An application of ramsey theorem to stopping games. *Games and Economic Behavior*, 2(42):300–306.

Sycara, K. (1991). Cooperative negotiation in concurrent engineering design. In *Computer-Aided Cooperative Product Development*, 269–297. Springer Verlag.

Touzi, N. and Vieille, N. (2002). Continuous-time dynkin games with mixed strategies. *SIAM J. Control Optim.*, 4(41):1073–1088.