

Talk is cheap: Evolved strategies for communication and action in asymmetrical animal contests

Jason Noble

Center for Adaptive Behavior and Cognition
Max Planck Institute for Human Development
Lentzeallee 94, D-14195 Berlin, Germany
noble@mpib-berlin.mpg.de

Abstract

Animal contests over resources are often settled by displays rather than fighting. Contests may involve asymmetries that cannot be perceived, such as unequal fighting ability. Classical game-theoretic accounts suggest that talk is cheap, and that honest signals conveying information about an asymmetry should not be expected. A model by Enquist (1985) reaches the opposite conclusion, and predicts reliable, cost-free signalling of fighting ability. An evolutionary simulation is presented which tests Enquist's assumption that weak animals will signal honestly because they have so much to lose by bluffing. The status of Enquist's honest signalling strategy as an ESS is questioned, and his conclusions are shown to be dependent on an idiosyncratic way of modelling animal combat.

1. Introduction

In acquiring the resources necessary for survival and reproduction, animals often come into conflict with one another. For example, funnel-web spiders compete for feeding territories (Riechert, 1987), elephants get into disputes over mating opportunities (Poole, 1989), and rival honeybee queens battle for control of a colony (Wilson, 1971). Sometimes these contests consist of a fight to the death, or at least to the serious injury, as in the case of the honeybees. However, much of the time, and across many species, contests stop short of violence and are settled by what appear to be threat displays, signals of strength or aggressive intent, and conventions such as "territory owner always wins."

Why should this be? Ethologists used to believe that the avoidance of fighting in favour of display behaviour served to benefit the species—contests could be settled without the high fitness costs of violent confrontation, such as debilitating injuries or death. With the realization that natural selection produces adaptations for the benefit of the individual (Williams, 1966) this explanation no longer held water. In one of the first applications

of game theory to animal behaviour, Maynard Smith and Price (1973) devised the Hawk-Dove game, which showed that a tendency to stop short of all-out fighting could be explained in terms of evolved self-interest (see also Maynard Smith, 1982).

The Hawk-Dove game considered symmetrical contests, in which each competitor was equally likely to win an escalated fight. But real contests are often asymmetrical. An obvious potential asymmetry is size, which correlates with the ability to inflict injury (Archer, 1987). If one animal is much bigger than another, if this fact is apparent to both parties, and if we assume that large size is a cue which cannot easily be faked, then the tendency of the smaller animal to retreat before the contest escalates is not difficult to explain. Clear size differences in contests between funnel-web spiders, for instance, indeed result in the rapid retreat of the smaller party (Riechert, 1987). Sometimes contests commence with a period of mutual assessment, in which both animals investigate possible asymmetries in size, stamina or fighting ability. For example, the parallel walks and roaring bouts engaged in by red deer stags (Clutton-Brock, Albon, Gibson, & Guinness, 1979) allow each animal to assess the other's size and strength. Two stags will only fight if neither sees the other as significantly stronger.

However, differences in fighting ability may not always be easy to perceive. To take a human example, if a stranger challenged you to a fight, you would not be able to tell just by looking whether or not he or she possessed a black belt in karate. Other relevant asymmetries that may not be available to perception include physiological variables such as hunger (Silbaugh & Ewald, 1987) or body temperature (Hardy, 1998), and differences in the subjective value of the contested resource (Dugatkin & Ohlson, 1990). In cases such as these it seems that communication would be a good way to avoid unnecessary and costly violence. Might there be an evolutionarily stable strategy (ESS) in which a competing animal signals its fighting ability, or some other potential asymmetry, and then retreats if the opponent has signalled a greater level of strength or need than its own?

The classic game-theoretic treatment of this problem

(Caryl, 1979; Maynard Smith, 1982) suggests that talk is cheap, and that any signalling system in which animals are supposed to honestly indicate their fighting ability will collapse due to invasion by bluffers.¹ After all, if an arbitrary signal indicating great strength results in opponents meekly walking away from a contested resource, it is easy to see that there is a temptation for weak animals to falsely signal strength. We can draw a loose analogy to the game of poker: if an opponent told you that they had a great hand, and there was a substantial sum of money in the pot, would you believe them and fold, or would you regard the signal as worthless?

Enquist (1985) presented a game-theoretic model of signalling in contests over a resource which reached precisely the opposite conclusions to those of Caryl and Maynard Smith. In Enquist's model contestants are either strong or weak; they know their own strength but cannot perceive the strength of their opponent. Contests are divided into two stages. In the first stage, contestants produce one of two cost-free signals, A or B. In the second stage, contestants must decide whether to attack unconditionally, to attack if the opponent does not withdraw, or to flee. This decision is based solely on the contestant's own strength and on the signal given by the opponent. Enquist considered the evolutionary stability of an honest signalling strategy, denoted S. Strategy S specifies that strong contestants should signal A, and weak contestants B. In the second stage, strong contestants should unconditionally attack strong opponents (i.e., those who have signalled A) and attack weak opponents (those who signal B) if they do not flee. Weak contestants should flee from opponents signalling strength, and attack opponents who signal weakness. Enquist found that strategy S would be an ESS if $\frac{1}{2}V - C > V - D$, where V is the value of the resource, C is the expected cost of a fight between equally matched opponents, and D is the cost to a weak animal of being attacked by a strong one.

Enquist's (1985) model has been accepted in the theoretical biology literature. For example, Johnstone (1998) endorsed it in a recent review, and Hurd (1997) and Számadó (2000) have both used it as a basis for continued work on signalling during contests. This is despite, or perhaps because of, the fact that Enquist's conclusions directly contradict the earlier game-theoretic idea that cheap talk (i.e., stable conventional signalling) should not be expected in animal contests. Enquist has certainly established that, within the terms of his model, strategy S will be evolutionarily stable under certain conditions. However, accepting that the model has something to tell us about reality requires accepting Enquist's game as a reasonable analogue of real-world animal con-

tests, and accepting the stability conditions for S as biologically plausible. On the latter point, Caryl (1987) has pointed out that $\frac{1}{2}V - C > V - D$ can be rewritten as $D > \frac{1}{2}V + C$, which means that the cost to a weak contestant of being attacked by a stronger one must be significantly greater than the costs of an escalated fight with another weak individual. Thus Enquist's conclusion is driven by the assumption that weak animals cannot afford to risk confronting stronger opponents and must be honest about their shortcomings. But Caryl notes that in real contests weak animals may be able to bluff (i.e., signal that they are strong) and then rapidly retreat if challenged. Even if weak bluffers are briefly attacked as they flee, the expected cost of such attacks may well be lower than the cost of an extended fight with another weak animal; this state of affairs would invalidate Enquist's result.

The purpose of the evolutionary simulation presented here is therefore to put Enquist's (1985) assumptions to the test in a richer modelling context than paper-and-pencil game theory affords. In a semi-realistic model of animal combat, is it really true that weak animals have so much to lose by bluffing that selection will favour the honest signalling of fighting ability?

There is some existing work using simulation methods to look at communication in animal contests, notably Wheeler and de Bourcier's (1995) model of territorial signalling. Wheeler and de Bourcier found evidence for Zahavi's (1975) handicap principle, the idea that high signal costs can enforce honesty in a communication system. More relevant to the current model, however, is my own previous paper (Noble, 1998) which addressed many of the issues discussed above and supported the "no cheap talk" theory, *contra* Enquist (1985). The earlier simulation used continuous-time recurrent neural networks (CTRNNs) as a control architecture for the simulated animal competitors; this design choice turned out to be the source of a serious problem which inspired the current revision. Although CTRNNs provide enormous behavioural flexibility, the evolved networks are extremely difficult to analyse. Thus the previous paper treated the evolved nets as black boxes, and had to rely on a variety of "behavioural" measures in order to determine whether or not signals of fighting ability were being exchanged. The conclusion was that such signals did not evolve, but this was an inference reached by comparing mean fitness and other factors across different experimental conditions, and it was unfortunately not possible to examine the evolved strategies and show directly that the opponent's signals were or were not being attended to.

The current model uses a production system to control each simulated animal's behaviour. This architecture is perhaps less fashionable in the adaptive behaviour community, presumably because of its association with the

¹Note that the fact that honest signalling is not an ESS does not imply that all contests will escalate—the Hawk-Dove game shows that non-violent conventional behaviour can persist without having any communicative function.

“good old-fashioned AI” paradigm, but it has the advantage that the evolved strategies can simply be read off the rule list. In addition, using such a transparent architecture allows the initial population to be programmed with a particular strategy rather than simply beginning with random strategies. This allows invasion studies, in which the evolutionary stability of a strategy such as Enquist’s S can be directly examined. Such techniques aid in developing a good understanding of just what is going on in the simulation, a necessary step in drawing parallels back to the real world (Di Paolo, Noble, & Bullock, submitted). The current simulation is also more ecologically oriented than the previous one: a steady-state genetic algorithm is used, and newborn organisms start out with low fighting ability and become more powerful as they get older, in contrast to being randomly assigned a new fighting ability for each contest. Finally, the actual contest is conducted along more concrete, physical lines than in the previous model: at each time-step, both contestants can choose between retreating, holding their ground, or moving aggressively towards their opponent and thereby attacking them if contact is made.

As argued previously (Noble, 1998), a secondary question that can be examined within a model like this concerns intention movements as the evolutionary seeds of aggressive signalling (Tinbergen, 1952; Krebs & Dawkins, 1984). An important perceptual input in the model is the movement of the opponent, either towards or away from the focal animal, and it is possible that movements themselves could take on a signalling value, rather than signals being exchanged through the arbitrary, cost-free signalling channel supplied.

2. The model

Contests commence with two randomly chosen simulated animals (henceforth animats) facing each other in a simple arena 10 metres across (Fig. 1). Each animat has a fighting ability that ranges between zero and one, and an accumulated energy score—neither of these properties can be perceived by the opponent. The pair are assumed to be competing for possession of a food resource worth 100 energy units. At each timestep, an animat can move up to one metre forwards or backwards, and can also produce an auditory signal ranging in intensity from zero to one.

Contests can end due to one animat fleeing, one animat seriously injuring the other, or because a time limit has been reached. If one animat moves backwards far enough to leave the arena, the other obtains the food resource and the contest is over. If both animats leave the arena simultaneously, neither obtains the food item. If the two animats are within 50cm of each other, any forward movement counts as an attack (and thus attacks may be one-sided affairs). During every timestep that one animat attacks another, it may succeed in in-




Figure 1: Starting positions of the two competing animats. The arena is 10m wide; one animat starts 2.5m from the left edge, the other 2.5m from the right edge.

flicting a serious injury. The probability of injury is proportional to the attacking animat’s fighting ability: $p(\text{injury}) = \text{fighting_ability} \times 0.2$. If a serious injury occurs, the contest ends at once and the injured animat forfeits the resource. If two animats collide head-on, they stop and attack each other at that point—it is not possible for one animat to push the other backwards, nor is it possible for the two to cross over such that the left animat ends up on the right and vice versa. Contests cannot last forever, and at some randomly determined time between 50 and 100 timesteps, the food resource disappears and the contest is terminated.

The animats have access to 11 perceptual inputs, listed below, on which they can base their decisions about movement and signalling. Briefly, each animat has access to privileged information about its own state, and can observe the signals and movements of the other. All inputs are scaled to a value between zero and one.

Own fighting ability: requires no re-scaling.

Own energy level: scaled from no energy (near death), to 1000 units of energy or more, the level at which reproduction is possible.

Opponent signal: the signal emitted by the opponent at the previous timestep.

Max. opponent signal: the maximum signal value heard from the opponent so far in this contest.

Previous opponent move: the distance moved by the opponent during the previous timestep. Scaled from one metre backward (fleeing), to one metre forward (charging).

Distance to opponent: scaled from “in contact”, to 10m, i.e., the width of the arena.

Δ distance to opponent: the change in the distance between the two animats at the last timestep. Scaled from -2m (both charging) to +2m (both fleeing).

Distance to edge: the distance between the animat and its “home” edge of the arena.

Time elapsed: number of timesteps since the start of the contest. Scaled from zero to 100 (i.e., the maximum contest length).

Random constant: a uniform random number between zero and one which stays constant for the duration of the contest.

Random flux: a uniform random number between zero and one which is newly generated each timestep. This value and the previous one could theoretically be used to generate unpredictable, Protean strategies.

The perceptual inputs are translated into outputs (movement and signalling) using a production system. An animat's production system specifies its contest strategy; the system is also the genotype that will be passed on if the animat succeeds in reproducing. There are a maximum of six rules in each production system, and six activity bits serve as switches specifying whether each rule is on or off. Each rule has 11 possible conditions, one for each input, and again there are 11 activity bits specifying which conditions are active and which can be ignored. The structure of a condition is as follows: one bit specifying whether the condition will pass high or low values in the corresponding input, one bit determining whether the input should be compared to a fixed threshold or to another input, and finally the fixed threshold value itself and a pointer to the comparison input. A movement and a signal value are specified as the output to be given if *all* of the active conditions of a rule are satisfied (this means that a rule's conditions are always linked by an AND relation). For example, a rule with two active conditions might be "If (own energy level < opponent's signal) and (time elapsed < 72 timesteps) then advance 0.32m and signal 0.85." At each timestep in a contest, an animat's production system is given fresh inputs, and the movement and signalling outputs are provided by the *first* active rule to have all of its conditions met; the production system thus implements a lexicographic strategy (Fishburn, 1974). In case no rules fire, the genotype includes default movement and signalling values.

Engaging in contests is not the only thing that the animats do. The simulation is organized into days, and the first activity each day is a period of conflict-free foraging in which all animats gain 50 units of energy. This energy bonus is included to ensure that animats with cowardly strategies, such as running away immediately from every contest, have at least a chance of accumulating enough energy to reproduce. Next, each animal in the population (maximum carrying capacity and initial population size are both set to 100) plays out a contest with a random opponent. The result is that each animat participates in at least one contest per day, and expects on average to participate in two.

After all the contests have been resolved, reproduction takes place: any animats with more than 1000 units of energy are randomly paired up and allowed to reproduce sexually. Each pair produces a single offspring. The birth costs each parent 500 units of energy, and the newborn animat begins life with 200 units of energy and a randomly determined fighting ability between 0.0 and 0.1. Newborns take the place of any animats that have recently died from their injuries; if there are no empty places, then the oldest animat is killed in order to make room. Each rule in the newborn's genotype is equally likely to be inherited from either parent, and the mutation rate for each element of the new production system is 0.002. With 380 elements in the system, this means an expected rate of about 0.76 mutations per birth. Mutation of binary values consists of flipping the bit, and the mutation of a pointer to a comparison input consists of randomly choosing a new input. Mutation of a real-valued threshold, movement or signal value is implemented by adding a random gaussian ($\mu = 0$, $\sigma = 0.05$) 95% of the time, and generating a completely new value the other 5% of the time.

At the end of each day there is a round of housekeeping. Any animat whose energy level has fallen below zero (usually because of suffering a serious injury in a contest) is declared dead. All living animats pay an energy cost for the complexity of their production system, proportional to the square of the total number of active conditions. The exact cost is $0.25 \times \text{active_conditions}^2$, which means, for example, that an animat with 10 conditions active would pay a cost of 25 energy units. This "cognition cost" is partly to reflect the general principle that complex cognition is expensive, and partly to foster the evolution of simpler, more comprehensible strategies free from extraneous detail. Finally, animats may grow in strength based on the amount of energy they have accumulated during the day: if an animat has made an energy profit of 100 units or more, then it is automatically awarded 0.1 additional units of fighting ability. Animats that have gained less than 100 energy units, but more than zero, have a *pro rata* probability of gaining the 0.1 unit increment in fighting ability.

The energy budget for the simulation has been set up such that the biggest cost, by far, is due to being seriously injured. This reflects the findings of Riechert (1982) that, in spider contests, the long-term fitness costs of serious injury—and, of course, death—are orders of magnitude greater than other costs such as energetic expenditure associated with threat displays. Table 1 summarizes the energy costs and benefits that the animats experience. Note that the average cost of being seriously injured is about 200 energy units (as well as a 20% loss in one's fighting ability). This is double the value of the contested resource, a ratio that ensures an equal proportion of Hawk and Dove strategists at equi-

Cost or benefit	Magnitude (units of energy)
Gaining resource	+100, if animat wins contest.
Feeding	+50, daily.
Serious injury	$-100 \times (1.5 + f)$, where f is the attacker’s fighting ability. Victim also loses 20% of own fighting ability.
Fighting	-0.2 per timestep; cost of attacking opponent.
Being in a contest	-0.1 per timestep.
Signalling	Cost-free, as stipulated by Enquist (1985).

Table 1: Energy costs and benefits for animats.

librium in the original Hawk-Dove game. Note also that the expected cost to a weak animat (fighting ability 0.2) of an all-out fight with a strong animat (fighting ability 0.8) is about 184 energy units (a 230 unit injury cost, and an 80% chance of being the one that gets injured). Whereas if the weak animat fights another weak animat, the expected cost is only 85 units (a 170 unit injury cost, and a 50% chance of being the one that gets injured). This means that, on the face of it, Enquist’s condition $\frac{1}{2}V - C > V - D$ is satisfied (as $\frac{100}{2} - 85$ is greater than $100 - 184$), and strategy S should be evolutionarily stable in the model.

3. Simulation results

3.1 Basic model

Genotypes in the initial population were generated randomly. Each simulation run continued for 150,000 days—preliminary work showed that this was enough time to achieve a reasonable convergence of strategies. Given that the average lifespan of an animat turned out to be approximately 20 days, this period represents about 7,500 generations. In the basic model and in the other conditions described below, ten evolutionary runs were performed, each with a different random seed value. If some populations went extinct before the 150,000 day limit was reached, as occasionally happened, a new seed value was chosen until ten complete runs had been recorded. The statistics reported below describe means (\pm standard deviations) across these ten runs: some measures summarize activity across the entire run, others report the state of the final population.

Contests lasted, on average, 19.4 (± 12.1) timesteps. This indicates that the animats tended to engage each other in some way, as immediately fleeing the arena would take only three timesteps. Contests were resolved 66.0% of the time through one or both animats fleeing,

If Own fighting ability < 0.53
Distance to opponent $< 1.70\text{m}$
Distance to edge $< 4.57\text{m}$
then retreat 0.93m and signal 0.47.

Otherwise advance 1m and signal 0.80.

Figure 2: The median strategy evolved in run 9. The default movement is maximally aggressive and the one active rule specifies conditions for retreating.

and 24.2% of the time through one animat inflicting a serious injury on the other. The remaining 9.8% of contests reached the maximum time limit.

After 150,000 days of evolution, the “median strategy” was recorded for each population. A rule was active in the median strategy if it was active in more than half of the population. Similarly, a condition was active in the median strategy if it was active for more than half of the population. The details of an active condition were found by checking whether a majority had it set to high- or low-pass, and whether a majority compared the input in question to a threshold or to another input. The identity of the comparison input, or alternatively the value of the threshold, was determined by calculating the median value across the population. Note that there is no guarantee that a median strategy will exist: in a population with high genetic variance, it may be that no rules or conditions are active in more than half of the individuals.

Across the 10 runs, there was significant variation in the median strategies found (and a strategy was indeed found in each case). The strategies were not overly complex, however, and the number of active rules was always between one and three inclusive. A common theme across all 10 strategies was an aggressive default movement, combined with at least one rule spelling out conditions under which the animat would retreat. Figure 2 illustrates this by showing one of the simpler, one-rule strategies in full. An animat following this strategy will start the contest with the default move of advancing one metre: the initial distance to the opponent is always five metres, and so the rule will not be satisfied. If two competing animats are both following this strategy, after two timesteps they will each have moved forward two metres, and they will be only a metre apart. At this point, the rule may fire. If one animat is relatively weak (i.e., its fighting ability is less than 0.53) then it will retreat 93cm, as for this animat the distance to the edge will be 4.5m. If the second animat is strong, it will pursue the first, ensuring that the weaker contestant eventually flees from the arena, although the stronger one will never get quite close enough to attack. Two strong animats will clash head on; neither will ever retreat. The behaviour of two weak animats is more interesting: they face each

Perceptual input	Level of use
Distance to opponent	22.9%(±13.9)
Own fighting ability	19.6%(±10.2)
Distance to the edge	19.0%(±14.6)
Δ distance to opponent	14.9%(±13.3)
Opponent signal	6.4%(±5.9)
Max. opponent signal	6.1%(±5.0)

Table 2: Level of use of the four most popular inputs for making decisions at each timestep, and the level of use of the two signalling inputs, across all 10 simulation runs. Level of use is the percentage of decisions in which this input was implicated in a rule that fired. For comparison, default moves and signals were used 43.5% of the time.

other one metre apart, and then each takes a 93cm step backwards. They then move forward one metre again. Next they will advance yet again and fight, as they will be exactly 4.57m from the edge of the arena and thus the rule will not fire.

What can we make of this strategy? Does it involve the honest signalling of fighting ability? The first point to notice is that users of this strategy pay no attention whatsoever to their opponent’s signals. There is no significance in the fact that a “louder” signal is given when advancing than when retreating, because in a population of animats all playing this strategy, no-one will be listening. However, there is some indication that animats may be signalling, or at least giving away information, through their movements. When weaker animats reach the moment of truth, one metre from their opponent, they reveal their low fighting ability by retreating. The interpretation of this result will be considered in section 4.

Space precludes a detailed analysis of all ten of the evolved strategies. Instead we can look at which inputs the animats actually used in their decision making, and which they were able to ignore. Every timestep in every contest is a decision point: both animats must select a response based on their perceptual input and their evolved strategy. Table 2 shows the percentage of the time that the four most popular inputs were used in making these decisions, averaged across the 10 simulation runs. Overall the animats evolve to attend to inputs such as the distance to the opponent, their own fighting ability, and the distance to the edge of the arena—as used in the example strategy shown in Figure 2.

Table 2 also shows the percentage of the time that the two signalling-related inputs were used. Clearly these two inputs were attended to significantly less often, and it seems likely that the low level of use observed was simply due to the noise introduced by mutant strategies across the course of the run (and the noise of the initial

If Own energy level > Distance to edge
Change in distance to opponent
< Own fighting ability
then advance 1m and signal 0.99.

If Own energy level > Previous opponent move
Maximum opponent signal > 0.31
then advance 1m and signal 0.95.

If Change in distance to opponent < 0.0m
then retreat 0.98m and signal 0.93.

Otherwise advance 0.30m and signal 0.48.

Figure 3: The median strategy evolved in run 2. The default movement is moderately aggressive, the first two rules specify conditions for extreme aggression, and the third rule spells out conditions for a rapid retreat.

random population). Evidence that the use of these inputs was not part of any evolved signalling system can be found by examining the median strategies identified at the end of each run.

Only one of these 10 strategies makes any reference to a signalling input: Figure 3 shows a relatively complicated three-rule strategy in which one rule depends on the maximum opponent signal. The strategy is quite hard to interpret, especially as it involves several points in which one input is compared to another apparently unrelated one. The default behaviour is moderate aggression. Roughly speaking, the first two rules indicate that if the animat’s own energy level and fighting ability are high, if the opponent’s movements are not particularly aggressive, and if the maximum signal heard from the opponent is higher than a threshold value, then an aggressive charge is called for. The third rule specifies that the animat will retreat if the two parties are moving towards each other (assuming that neither of the first two rules has fired).

Does the use of the opponent’s maximum signal value in the second rule in Figure 3 imply the evolution of honest signalling? In fact we can easily see that this is not a signalling strategy, by inspecting the signalling outputs in the different conditions. In every eventuality a signal greater than 0.31 will be given, and so the “Max. opponent signal” condition of the second rule will always be satisfied. The only exception is on the very first move of the contest, when the value for the maximum opponent signal is zero. The adaptive value of attending to this input at all is probably just to ensure that rule two does not fire on the very first timestep of the contest, and that execution will fall through to the default strategy and cause a moderately aggressive advance to be made.

```

If Own fighting ability < 0.4
  Opponent's signal > 0.5
  then retreat 1m and signal 0.0.

If Own fighting ability < 0.4
  then advance 1m and signal 0.0.

Otherwise advance 1m and signal 1.0.

```

Figure 4: An analogue of Enquist’s (1985) strategy S, expressed in the framework of the animats’ production system. The default strategy is an aggressive advance and a loud signal. The first rule specifies that weaker animals will retreat from a loud signal, and the second, that they will advance without signalling if they hear no signal.

3.2 Stability of Enquist’s strategy S

The results presented for the basic model suggest that the exchange of honest signals of fighting ability via the signalling channel is not favoured by selection. However, the 10 simulation runs each began with a randomly generated set of initial strategies. It is possible that stable signalling strategies exist, but that their basins of attraction in genotype space are not large enough for the strategies to emerge given random initial conditions. In this section we will look at what happens when an analogue of Enquist’s signalling strategy S is programmed into the initial population.

Figure 4 shows the way in which strategy S was implemented as a two-rule production system. The cutoff point between weak and strong was set at 0.4 as this was the approximate mean fighting ability observed in the 10 runs described in section 3.1. For stronger animats, the chosen action will always fall through to the default behaviour of aggressively advancing while making a loud signal. This loud signal is equivalent to displaying A in Enquist’s model. For weaker animats, rule one or rule two will always fire. This means that weaker animats will announce their status by always signalling with zero intensity (equivalent to displaying B). If a weak animat detects a signal (i.e., a strong opponent) it will retreat, but if there is no signal it will advance to fight its presumably weak opponent.

Ten brief runs of 1000 days duration were performed with initial populations programmed with strategy S, and a mutation rate of zero, in order to look at the characteristics of a population that engaged in this type of honest signalling. The enforced use of strategy S led to shorter contests (mean duration 7.36 ± 0.02 timesteps), but, perhaps surprisingly, they were bloodier, ending in serious injury 49.0% of the time. As is clear from Figure 4, only two inputs are attended to by strategy S players: 67.0% (± 0.3) of the decisions made were based on

Perceptual input	Level of use
Distance to opponent	30.3% (± 13.8)
Own fighting ability	27.4% (± 4.6)
Previous opponent move	16.5% (± 19.1)
Time elapsed	13.9% (± 7.7)
Opponent signal	4.3% (± 2.2)
Max. opponent signal	3.5% (± 2.9)

Table 3: Level of use of the four most popular inputs for making decisions at each timestep, and the level of use of the two signalling inputs, across all 10 simulation runs beginning with a population of strategy S players. Level of use is the percentage of decisions in which this input was implicated in a rule that fired. For comparison, default moves and signals were used 57.8% of the time.

the animat’s own fighting ability, and 12.0% (± 0.04) of decisions were influenced by the opponent’s signal. The latter figure may seem low, but it is due to the fact that the contests that can be settled by signalling involve a weak and a strong animat, and these contests are over in three timesteps when the weak animal flees. Thus a relatively low percentage of the overall number of animat decisions are based on the opponent’s signal.

The evolutionary stability of strategy S was investigated by conducting ten runs of 150,000 days duration. Animats in the initial populations were set to play strategy S, and the mutation rate was 0.002 as in the basic model. These simulations can therefore show us whether or not a population of strategy S players is resistant to invasion by mutant strategies. Table 3 shows the percentage of the time that the four most popular inputs were used to make decisions, and the usage levels for the two signalling inputs, averaged across all 10 simulation runs. It is immediately clear that strategy S has not been able to resist the invasion of alternative strategies: for example, the “Distance to opponent” input is used most often, despite not being present in the initial population. Inspection of the median strategies shows them to be very much like those that evolved in the basic model. Only two of the median strategies utilize the “Opponent’s signal” input. In one case (run 1) the input is implicated in only 1.0% of decisions. In the other case (run 8) the input is used in 7.1% of decisions, but it is not clear that it has a communicative function—certainly the evolved strategy does not resemble S.

Why is S not stable against invasion? Enquist (1985) shows that it is an ESS under conditions that might appear to be satisfied here: why the inconsistency? As discussed in section 1., Enquist’s argument for the evolutionary stability of S rests on the idea that weak contestants must honestly signal their weakness because they cannot afford the risks of being injured by a stronger

Perceptual input	Level of use
Own fighting ability	27.4% (± 10.3)
Previous opponent move	24.3% (± 11.4)
Distance to opponent	21.8% (± 8.6)
Opponent signal	17.4% (± 23.5)

Table 4: Level of use of the four most popular inputs for making decisions at each timestep, across all 10 simulation runs beginning with a population of strategy S players, and with exit from the arena impossible after 5 timesteps.

contestant. The results discussed so far present a different picture, in which weak animats do not signal their weakness at all, and only give away information about their state by retreating at the last possible moment. It may be that, in the current model, weak animats can afford to behave in this way because the condition $D > \frac{1}{2}V + C$ —identified by Caryl (1987) as unrealistic—is not met. That is, the model lacks a mechanism that would maintain a high value of D (where D is the cost to a weak animal of facing up to a stronger one). Consider the pattern of behaviour outlined for the strategy shown in Figure 2. Clearly, if weak animats can bluff it out against stronger opponents, up to a point, and then retreat without being harmed, then D is not particularly high.

If the above line of reasoning is correct, then it should be possible to change the model in some way, however unrealistic, such that weak animats suffer high costs for not fleeing from strong ones, and thus observe an increase in the success of honest signalling strategies like S. One such change would be to introduce “closing doors” for the arena; these doors start out open but close irrevocably after five timesteps. In this way a weak animal must either flee immediately from a stronger opponent, or be trapped with that opponent in an extended fight. Table 4 shows the levels of use of the four most popular inputs across ten simulation runs performed with this modification in place. Note that the “Opponent’s signal” input is the fourth most popular under these conditions: for the first time there appears to be selection for paying attention what one’s opponent has to say. Actually there is a great deal of variance across runs, and the differences between 17.4% and 6.4% (as observed in the basic model) or 4.3% (as seen when starting with an initial population playing S) are not statistically significant. Nevertheless, they are strongly suggestive.

4. Discussion

Enquist’s (1985) model suggests that weak contestants have so much to lose by bluffing that selection will favour the honest signalling of fighting ability. The simulation reported here shows that this claim is very much depen-

dent on Enquist’s idiosyncratic way of modelling animal combat. Given more realistic signalling and movement over an extended period of time, reliable signalling of fighting ability did not evolve. This result held whether the members of the initial population were allocated random strategies or were programmed to play an analogue of Enquist’s strategy S. Results in the latter condition show that strategy S is not an ESS in the current model, which must detract from Enquist’s claims of generality. Finally, the results in the “closing doors” condition suggest that a highly artificial scenario is necessary to place weaker contestants in a position where it would be advantageous to signal honestly about their weakness.

Although disagreeing with his conclusions, we can sympathise with Enquist’s motivation. Field observations of behaviour in animal contests sometimes do seem to contradict the game-theoretic conclusion that talk is cheap (e.g., Hansen, 1986; Dabelsteen & Pedersen, 1990). There really is a need for explanation in such cases: either the appearance of signalling is an illusion, or our game-theoretic models are leaving something out. But unfortunately Enquist (1985) settled on some questionable assumptions in his attempt to explain apparent honesty in contests, and others have since built upon these dubious foundations (Hurd, 1997; Számadó, 2000).

The available signalling channel was not used by the animats, but there was evidence that they were gaining information about fighting ability based on observations of each other’s movements. Does this count as communication? A poker analogy may again be useful: if you are bluffing with a terrible hand, the other players do not know whether your cards are strong or weak. If someone calls your bluff, by seeing your bet and then raising again, you will probably fold. By doing so, you have given the other players information about your strength (i.e., they now know that you had a poor hand). However, the *reason* you folded was not to provide information to others, but because it was the best way to minimize your expected losses at that point. Similarly, the weak animals using the strategy shown in Figure 2 are giving away information about their weakness when they back off from immediate confrontation at timestep three. But their choice at this point is to retreat or to start fighting against an opponent that may well be stronger than they are. The expected costs of entering such a fight are higher than the costs of retreating, so the animat retreats. Information is conveyed to the opponent by this behaviour, but it is not the function—or in Millikan’s (1984) terminology, the “proper function”—of the retreat to be informative.

All models make simplifying assumptions. Models of animal contests often make assumptions that may seem obtuse to a layman: who would believe that fighting ability cannot be perceived by opponents, or that contestants have no memory for previous interactions, or that

contests occur between random individuals in a population without spatial or social structure? It is more than likely that the evolutionary stability of honest signals in real-world animal contests (if indeed such signals are stable) is due to the effect of one of these neglected factors. For example, if we look at models in which repeated interactions take place, and in which contestants can remember the results of previous meetings, honest signalling ESSs can be found given quite reasonable assumptions (van Rhijn & Vodegel, 1980; Silk, Kaldor, & Boyd, 2000). An important simplifying assumption in the current model is that the animats care about nothing beyond gaining the resource and minimizing their chances of being injured. Thus, once the weaker animat's bluff has been called, it will run away, and the resource will go to the stronger, with no further energy expenditure. Evolution does not select for spitefulness, and there is no reason to expect strong competitors to dish out the punishment that would be necessary to keep the weak ones honest. But it is possible that in a much more ecologically detailed simulation, which modelled territories and repeated interactions with neighbours over time, it would pay for strong animals to attack and even kill younger, weaker intruders, perhaps to prevent them from growing up to be bigger and stronger competitors who live right next door. In this situation Enquist's logic might apply after all.

Acknowledgments

Thanks to Clark Barrett and Peter Todd for comments. Thanks also to Peter Stuer for detailed criticisms of an earlier model.

References

- Archer, J. (1987). *The Behavioural Biology of Aggression*. Cambridge University Press, Cambridge.
- Caryl, P. G. (1979). Communication by agonistic displays: What can games theory contribute to ethology?. *Behaviour*, *67*, 136–169.
- Caryl, P. G. (1987). Acquisition of information in contests: The gulf between theory and biology. Paper presented at the ESS Workshop on Animal Conflicts, Sheffield, UK, July.
- Clutton-Brock, T., Albon, S. D., Gibson, R. M., & Guinness, F. E. (1979). The logical stag: adaptive aspects of fighting in red deer (*Cervus elaphus L.*). *Animal Behaviour*, *27*, 211–225.
- Dabelsteen, T., & Pedersen, S. B. (1990). Song and information about aggressive responses of blackbirds, *Turdus merula*: Evidence from interactive playback experiments with territory owners. *Animal Behaviour*, *40*, 1158–1168.
- Di Paolo, E. A., Noble, J., & Bullock, S. (Submitted). Simulation models as opaque thought experiments. In *Artificial Life VII*.
- Dugatkin, L. A., & Ohlsen, S. (1990). Contrasting asymmetries in value expectations and resource holding power: Effects on attack behaviour and dominance in the pumpkinseed sunfish, *Lepomis gibbosus*. *Animal Behaviour*, *39*, 802–804.
- Enquist, M. (1985). Communication during aggressive interactions with particular reference to variation in choice of behaviour. *Animal Behaviour*, *33*, 1152–1161.
- Fishburn, P. C. (1974). Lexicographic orders, utilities and decision rules: A survey. *Management Science*, *20*, 1442–1471.
- Hansen, A. J. (1986). Fighting behaviour in bald eagles: A test of game theory. *Ecology*, *67*, 787–797.
- Hardy, I. C. W. (1998). Butterfly battles: On conventional contests and hot property. *Trends in Ecology and Evolution*, *13*(10), 385–386.
- Hurd, P. L. (1997). Is signalling of fighting ability costlier for weaker individuals?. *Journal of Theoretical Biology*, *184*, 83–88.
- Johnstone, R. A. (1998). Game theory and communication. In Dugatkin, L. A., & Reeve, H. K. (Eds.), *Game Theory and Animal Behavior*, pp. 94–117. Oxford University Press, New York.
- Krebs, J. R., & Dawkins, R. (1984). Animal signals: Mind reading and manipulation. In Krebs, J. R., & Davies, N. B. (Eds.), *Behavioural Ecology: An Evolutionary Approach* (Second edition), pp. 380–402. Blackwell, Oxford.
- Maynard Smith, J. (1982). *Evolution and the Theory of Games*. Cambridge University Press, Cambridge.
- Maynard Smith, J., & Price, G. R. (1973). The logic of animal conflict. *Nature*, *246*, 15–18.
- Millikan, R. G. (1984). *Language, Thought, and Other Biological Categories*. MIT Press / Bradford Books, Cambridge, MA.
- Noble, J. (1998). Tough guys don't dance: Intention movements and the evolution of signalling in animal contests. In Pfeifer, R., Blumberg, B., Meyer, J.-A., & Wilson, S. W. (Eds.), *From Animals to Animats 5: Proceedings of the Fifth International Conference on Simulation of Adaptive Behavior*, pp. 471–476. MIT Press / Bradford Books, Cambridge, MA.
- Poole, J. H. (1989). Announcing intent: The aggressive state of musth in African elephants. *Animal Behaviour*, *37*, 140–152.
- Riechert, S. E. (1982). Spider interaction strategies: Communication vs. coercion. In Witt, P. N., & Rovner, J. (Eds.), *Spider Communication: Mechanisms and Ecological Significance*, pp. 281–315. Princeton University Press, Princeton, NJ.
- Riechert, S. E. (1987). Between population variation in spider territorial behavior: Hybrid-pure population line comparisons. In Huettel, M. D. (Ed.), *Evolutionary Genetics of Invertebrate Behavior*, pp. 33–42. Plenum Press, New York.
- Silbaugh, J. M., & Ewald, P. W. (1987). Effects of unit payoff asymmetries on aggression and dominance in meadow voles, *Microtus pennsylvanicus*. *Animal Behaviour*, *35*, 606–608.
- Silk, J. B., Kaldor, E., & Boyd, R. (2000). Cheap talk when interests conflict. *Animal Behaviour*, *59*, 423–432.
- Számádó, S. (2000). Cheating as a mixed strategy in a simple model of aggressive communication. *Animal Behaviour*, *59*, 221–230.
- Tinbergen, N. (1952). "Derived" activities; their causation, biological significance, origin, and emancipation during evolution. *Quarterly Review of Biology*, *27*(1), 1–32.
- van Rhijn, J. G., & Vodegel, R. (1980). Being honest about one's intentions: An evolutionary stable strategy for animal conflicts. *Journal of Theoretical Biology*, *85*, 623–641.
- Wheeler, M., & de Bourcier, P. (1995). How not to murder your neighbor: Using synthetic behavioral ecology to study aggressive signaling. *Adaptive Behavior*, *3*(3), 273–309.
- Williams, G. C. (1966). *Adaptation and Natural Selection*. Princeton University Press, Princeton, NJ.
- Wilson, E. O. (1971). *The Insect Societies*. Belknap Press / Harvard University Press, Cambridge, MA.
- Zahavi, A. (1975). Mate selection—a selection for a handicap. *Journal of Theoretical Biology*, *53*, 205–214.