# Tough guys don't dance: Intention movements and the evolution of signalling in animal contests

Jason Noble
School of Cognitive and Computing Sciences
University of Sussex
Brighton BN1 9QH
UNITED KINGDOM
`jasonn@cogs.susx.ac.uk`
Tel: +44 1273 678524
Fax: +44 1273 671320

## Abstract

Two animals contesting possession of a resource sometimes fight, but more often engage in aggressive displays until one or the other retreats. Game theory predicts that signals of strength or aggressive intent will not be heeded unless necessity or high cost make them reliable, but some animal threat displays appear to contradict this prediction. An evolutionary simulation model is presented in which two animats of variable fighting ability compete for a resource and have access to their opponent's intention movements, i.e., their real-time behavioural choices. Reliable communication of fighting ability does not evolve. The results of the simulation are used to assess conflicting theories of signalling in competitive interactions.

**Topic Area:** Collective and social behavior.
Submitted as a short paper.

# Tough guys don't dance: Intention movements and the evolution of signalling in animal contests

Jason Noble
School of Cognitive and Computing Sciences
University of Sussex, Brighton BN1 9QH, U.K.
`jasonn@cogs.susx.ac.uk`

### Abstract

Two animals contesting possession of a resource sometimes fight, but more often engage in aggressive displays until one or the other retreats. Game theory predicts that signals of strength or aggressive intent will not be heeded unless necessity or high cost make them reliable, but some animal threat displays appear to contradict this prediction. An evolutionary simulation model is presented in which two animats of variable fighting ability compete for a resource and have access to their opponent's intention movements, i.e., their real-time behavioural choices. Reliable communication of fighting ability does not evolve. The results of the simulation are used to assess conflicting theories of signalling in competitive interactions.

## 1 Introduction

Animals contesting the possession of a resource are often observed to settle the dispute by exchanging signals or threat displays rather than engaging in an all-out fight. For example, mantis shrimps *Gonodactylus bredini* contest the ownership of small cavities in their coral reef habitat. These contests sometimes result in physical combat, but often an opponent is deterred by a claw-spreading threat display (Adams & Caldwell, 1990). Red deer stags *Cervus elaphus* compete for control of groups of females, but unless two stags are closely matched in strength, the weaker will usually retreat after a roaring contest and/or a parallel walk display (Clutton-Brock, Albon, Gibson, & Guinness, 1979).

There is some controversy in the biological literature about how best to explain these observations. The early ethologists suggested that threat displays were honest signals of aggressive intent that served to prevent costly fights, but the group-selectionist overtones of this idea mean that it is no longer taken seriously. Subsequently, game-theoretic models of aggressive interaction have predicted that signals of aggressive intent or high fighting ability will be vulnerable to invasion by bluffers and thus not evolutionarily stable (Maynard Smith, 1982)—receivers of the signal will have no reason to take it seriously. There is widespread agreement that if a signal is for some reason unfakeable (such as the necessary connection between stamina and energetically costly roaring in red deer) then it will be evolutionarily stable for animals to use the signal to settle contests, escalating only when opponents are well-matched. However, the similar use of fakeable signals is not expected to be stable. This latter conclusion is apparently at odds with the ethological data (Hinde, 1981). Sometimes animals do seem to pay attention to threat displays that could be (and are) faked, as when a mantis shrimp, its exoskeleton soft and vulnerable after moulting, successfully drives off an intruder that could in fact defeat it in combat.

Several authors have tried to address this explanatory gap, constructing models of contests in which the strength and intentions of one animal cannot be directly perceived by the other, but might be more or less reliably communicated through a threat display. Adams and Mesterton-Gibbons (1995) argue that the use of a threat by a weak animal involves a "vulnerability cost". This is one version of the handicap principle (Zahavi, 1977), which asserts a necessary relationship between a signal's cost and its reliability. A vulnerability cost exists because, if the threat does not work, a weak animal is likely to be seriously injured by its probably-stronger opponent. On the other hand, weak animals stand to gain proportionately more if the threat is successful. Adams and Mesterton-Gibbons predict that bluffing will be stable: threats will be made and sometimes heeded, but should be expected from the very strong *and* the very weak. Wheeler and de Bourcier (1995) describe a simulation of aggressive territorial signalling. They also invoke the handicap principle, but to different effect: their simulation suggests that if the production costs associated with a threat display are sufficiently high, then the honest signalling of "aggressiveness", i.e., aggressive intent, will be evolutionarily stable.

In contrast, Enquist (1985) presents a game-theoretic model in which contestants are either strong or weak, and cost-free, binary signals are exchanged before the decision to fight or flee is taken. Enquist concludes that, under certain conditions[1], the honest signalling of fighting ability will be evolutionarily stable, even though the signals are cost free. Fights will occur only between evenly-matched opponents, and weak animals will defer to signals denoting strength. In a second model, Enquist goes on to claim that the cost-free signalling of "local strategy", i.e., aggressive intent, can also be evolutionarily stable. Hurd (1997) has recently extended Enquist's first model and also concludes that the cost-free signalling of fighting ability is possible. Furthermore, if only two signals are available, and if they vary in cost, it will be evolutionarily stable for weak animals to use the more costly of the two signals: not as a bluff, but as an honest advertisement of low fighting ability. These paradoxical results clearly run against the grain of most game-theoretic predictions about signalling in cases where the interests of the interacting parties conflict: such signals are either expected to be uninformative (Maynard Smith, 1982), and thus not signals at all, or stabilized by being too costly for poor-quality signallers to produce (Grafen, 1990). Caryl (1987) has criticized Enquist's methodology and argues that reconsideration of his first model leads to the opposite conclusion.

In an attempt to resolve the controversy, the current paper presents a simulation model of contests over an indivisible resource. The results will be compared with the conflicting predictions of the models outlined above. The aim of using a simulation is to avoid oversimplification. In particular, time will be modelled in an approximately continuous fashion: game-theoretic models of aggressive signalling rarely allow for more than two time-steps—an exchange of signals followed by a choice of actions—and thus may fail to capture critical aspects of real-time interactions.

There is a second motivation for the work presented here. Whilst the honest signalling of intentions looks unlikely from a game-theoretic perspective, it has been cogently argued (Tinbergen, 1952; Krebs & Dawkins, 1984) that *intention movements* (i.e., movements necessarily preceding an action, such as a dog baring its teeth to bite) probably function as "seeds" in signal evolution. Rather than incorporating an exchange of artificial, discrete signals, the current model seeks to explore the plausibility of the intention-movements idea by using such movements as the medium of potential information transmission.

## 2 The simulation model

The model involves two simulated animals (henceforth animats) contesting the possession of a resource. Due to a shortage of hard data on the metabolic costs of fighting, the simulation does not model contests in a particular species. However, the mantis shrimp will serve as a useful illustration: assume that two shrimps have discovered a desirable cavity. Each shrimp knows its own fighting ability, but is unable to perceive the ability of its opponent. Each shrimp can perceive movements towards aggression or retreat on the part of its opponent. Based on this information, shrimps can elect to attack or to flee or to do something in between. If the shrimps engage in an all-out fight, both shrimps will suffer costly injuries but the stronger one is likely to win and gain possession of the cavity. Over generational time, the shrimps may or may not evolve a signalling system that allows them to settle contests without fighting.

The model is based on a single behavioural continuum between attacking and fleeing. An animat is located at some point $\Theta$ on this continuum; contests involve two animats randomly selected from a population. Animats begin the contest at $\Theta = 0$ and, each time-step, can move a maximum of $\delta$ units in either direction. An animat with $\Theta > A$ has crossed the "attack line", and is deemed to be physically attacking its opponent. An animat with $\Theta < F$ has crossed the "flee line": the contest ends and the animat has abandoned the contested resource to its opponent. In the runs presented here, $F = -A$ and $\delta = A/4$. The simulation captures intention movements in that animats can neither immediately attack nor immediately flee: from the starting position, it takes at least four time-steps to do either. A plausible threat display might be to "hover" with $\Theta$ just less than $A$. Figure 1 shows a hypothetical contest.

---

[1] $0.5v - c > v - d$, where $v$ is the value of the resource, $c$ is the cost of an escalated fight between two equally matched opponents, and $d$ is the cost to a weak animal of being attacked by a strong one.

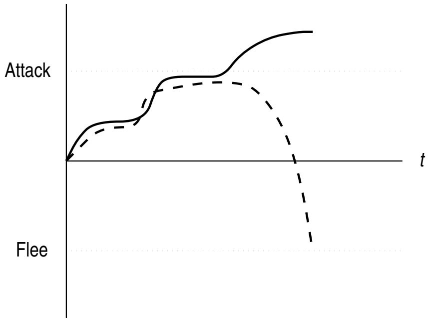

Figure 1: Course of a hypothetical contest: the horizontal axis is time and the vertical axis is the $\Theta$-continuum. The weaker animal (dashed line) flees after the stronger animal shows a greater willingness to back up "threat displays" by attacking.

Animats have associated with them an energy level $e$ that is set to 0 at birth, and a fighting ability $f$ ($5 \leq f \leq 15$) that is randomly assigned at the beginning of each contest.[2] The two costs and one benefit of the contest affect the value of $e$. The first cost is due to being attacked: at each time-step that $\Theta_{opponent} > A$, the animat suffers an injury cost of $-f_{opponent}$. The second cost is an energy cost for aggressive display or attack: at each time-step that $\Theta > 0$, the animat pays a cost of $-k\Theta/A$, with $k = 1$. So, for example, an animat attacking its opponent with $\Theta = 1.1 \times A$ would endure -1.1 units of energy cost per time-step; note that the cost of attacking is always much less than the cost of being attacked, for any $f$. Note also that any activity where $\Theta < 0$ involves no energy cost—this is justified on the basis that backing off, preparatory to running away, is much less energetically expensive than aggressive behaviour. In principle, therefore, cost-free signalling should be possible when $\Theta < 0$. The only benefit in the contest is to gain control of the resource, which is worth $V = 100$ units.

The contest can end in one of three ways. One animat may flee, as discussed above. Secondly, one animat may win the contest through brute force: if an animat loses more than $C = 200$ units of energy *during any one contest*, it has been physically overcome by its opponent. The contest ends immediately and the opponent gains the resource. Regrettably, the value of $C$ puts an artificial cap on the amount of damage an animat can sustain in one contest; however, the values of $C$ and $V$ have been chosen such that, on the face of it, the resource is worth having but not worth suffering serious injury for. In Maynard Smith's (1982) much simpler Hawk-Dove game, these relative cost and benefit values resulted in a mixed-strategy equilibrium where half of the population played Hawk and the other half played Dove. The contest described here can be thought of as a souped-up continuous-time version of the Hawk-Dove game, in which players of varying strength[3] can perceive and respond to their opponent's tendency towards a hawkish or dovish strategy. Finally, the contest can end because a time limit, $t_{max} = 50$ time-steps, has been reached, in which case neither animat gains the resource. The values of $t_{max}$, $C$, $k$, and $\delta$, and the range of values of $f$, have been coordinated such that it is possible for even the weakest animat to overcome an opponent (assuming the opponent does not retaliate) within the time limit.[4]

The animats have as inputs $f_{self}$, $\Theta_{self}$, and $\Theta_{opponent}$. Informally, they know their own strength, they can see what they're doing, and they can see what their opponent is doing. The animats also have access to a random input, to allow for probabilistic strategies. The animats produce a continuous output in the range $\pm\delta$ which is applied to their $\Theta$-position for the next time-step.

The animats were implemented as five-neuron fully inter-connected continuous-time recurrent neural nets (CTRNNs), with the activity of neuron 0 taken as the output. CTRNNs are the among the most general of artificial neural network architectures; they have been treated here as a black-box control system. The recurrent aspect of the nets makes it possible for the animats to evolve some form of short-term memory rather than being purely reactive. All parameter values for the nets were taken from Yamauchi and Beer (1994). The evolutionary engine was a genetic algorithm with a population size of 100, run for 5000 generations. In each generation, animats were randomly selected to play out 500 contests; each animat could thus expect to participate in 10 contests in its lifetime and was guaranteed at least 5. An animat's fitness score was $e_{total}/n_{contests}$. For breeding purposes, these fitness scores were normalized as deviations from the mean: animats with negative scores were

---

[2] Heritable fighting ability would quickly lead to uninteresting fixation. The model selects for responsiveness to varying values of $f_{self}$, as presumably exists in the mantis shrimp which is periodically vulnerable due to moulting.

[3] In the Hawk-Dove game all players are of equal fighting ability.

[4] Although a degree of arbitrariness is inevitable in setting parameter values for an abstract simulation, it is hoped that this coordination of values will at least prevent such disasters as, for example, enforcing honesty amongst the weak by having $C$ too high, or $t_{max}$ too low, for a weak animat to ever win by fighting.

discarded, and roulette-wheel selection was applied to the remainder.

Two control conditions were devised: in the "blind" condition, the animats are denied access to $\Theta_{opponent}$ and therefore any communication is impossible. In the "unfakeable" condition, the animats are given access to $f_{opponent}$, and can be thought of as exchanging unfakeable signals of strength. In order to facilitate comparisons with previous models, both continuous and discrete distributions of fighting ability were used. In the continuous fighting ability (CFA) case, $f$ was uniformly randomly distributed between 5 and 15, while in the discrete fighting ability (DFA) case animats were either weak ($f = 5$) or strong ($f = 15$). The experimental and control conditions were crossed with the two $f$-distributions to make six conditions in all.

# 3 Results

Each of the six experimental conditions was run 10 times with a different random seed. The final 1000 generations of each 5000-generation run were used as a window period for statistical analysis; all of the results below refer to average behaviour within this period. Standard errors refer to error across the 10 trials.

In observing the progress of the contests on a computer monitor, the overwhelming impression was that most of the animats wasted no time in moving rapidly towards either the attack or the flee line. Contests were generally resolved quickly: the mean duration was 10.14 time-steps overall, and less than 0.4% of contests reached $t_{max}$. If threat displays were occurring, they did not involve sustained action.

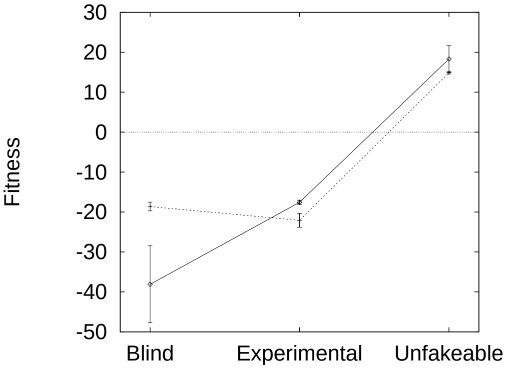

Figure 2: Mean fitness ±1 s.e., for CFA (solid line) and DFA (dashed line), across the three conditions.

Figure 2 shows the mean fitness values across conditions. Several points can be made based on this graph alone. Mean fitness is determined by the value $V$ of the resource, which an animat can expect to obtain about half the time, less the mean costs of aggression and being injured. In an ideal signalling system, where cost-free signals of strength were exchanged and the weaker animat always retreated immediately, the mean fitness would be $V/2 = 50$ (assume that equally matched animats allocate the resource randomly). In the experimental condition, such ideal situations clearly did not evolve—the negative mean fitness implies that the $V/2$ expected payoff was balanced against greater energy and injury costs. The unfakeable condition provides an index of how efficiently the animats can allocate $V$ when they have reliable, cost-free information about their opponent's strength, and in the CFA case their efficiency approaches 60% of the ideal. The fact that mean fitness in the experimental condition was significantly lower than in the unfakeable condition (CFA: $t_{18} = 8.03$, $p < 0.001$; DFA: $t_{18} = 19.30$, $p < 0.001$) is evidence that if fighting ability is being signalled, the signals are costly. Differences in mean fitness between the blind condition and the experimental condition were of marginal statistical significance. Communication is impossible in the blind condition; therefore the similarity in fitness scores suggests that if signals are being exchanged in the experimental condition, their costs cancel out their benefits.

Figure 3 shows the percentage of contests that were resolved by all-out fights, i.e., by one animat overcoming the other (the remainder were almost all resolved by one animat fleeing). In the unfakeable condition, as might be expected, fights were relatively rare. In the experimental condition (CFA) they occurred 24.5% of the time, and of those fights, 41.0% were between animats with closely matched fighting ability ($|f_a - f_b| < 2$). At first glance this might suggest that the animats had evolved a signalling system that allowed them to avoid costly fighting much of the time. However, there was no significant difference in the proportion of fights between the experimental condition and the blind condition, and fights in the blind condition were just as likely to be against closely matched opponents. If a reliable signalling system were in place, it would presumably lead to fewer fights overall and a greater proportion of well-matched opponents when fights did take place.

Overall, stronger animats were more aggressive: they were likely to have a higher Θ-values than weak animats at any particular time-step. This

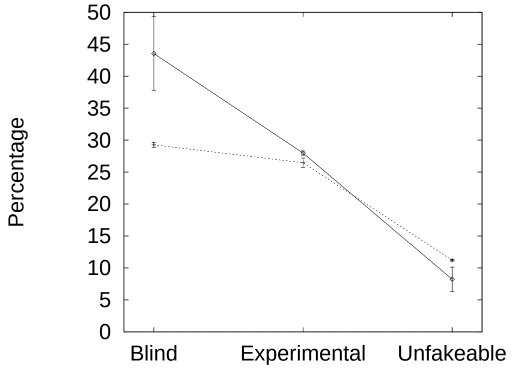

Figure 3: Mean percentage of contests resolved by all-out fights ±1 s.e., for CFA (solid line) and DFA (dashed line), across the three conditions.

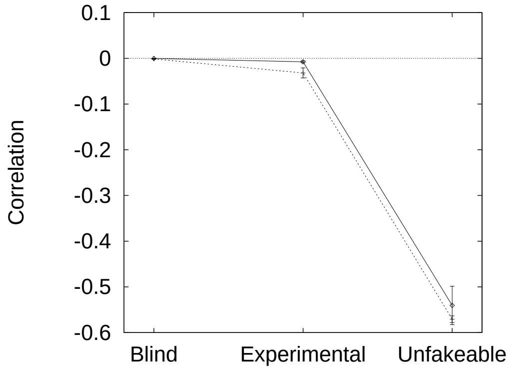

Figure 4: Mean correlation between $\Theta_{self}$ and $f_{opponent}$ ±1 s.e., for CFA (solid line) and DFA (dashed line), across the three conditions.

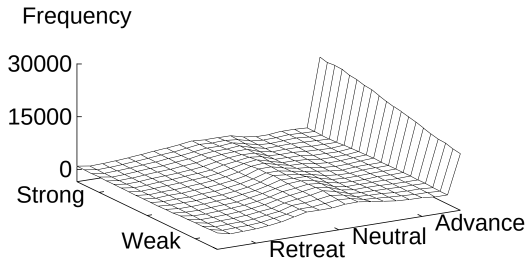

Figure 5: Total frequencies of opening moves by fighting ability, with each variable grouped into 20 bins for plotting. The plot shows data for the window period of run 6 in the experimental condition with CFA.

suggests that weak animats should have been able to infer their opponent's strength from their Θ-value and respond accordingly. Figure 4 shows the mean correlation coefficient between $\Theta_{self}$ and $f_{opponent}$ at time-step 4, when all contests were still in progress (the correlation was calculated for each generation). In the unfakeable condition there is a negative relationship: unsurprisingly, animats that could reliably perceive their opponent's strength were likely to flee from stronger opponents. In the blind condition there is of course no relationship at all. In the experimental condition, there is the suggestion of a negative relationship, but it accounts for less than 1% of the variance in $\Theta_{self}$.

A general link between $f_{self}$ and $\Theta_{self}$ also suggests the possibility of bluffing, i.e., deception by moving to a Θ-value usually characteristic of higher $f$. Given the brevity of the contests and the likely importance of first impressions, bluffing was investigated by tabulating the animats' opening moves (their $\Delta\Theta$ for time-step 1). Figures 5 and 6 show the frequencies of opening moves by fighting ability for CFA and DFA respectively. Substantial variation existed in these relationships across trials, and presenting mean values would obscure the situation—figures 5 and 6 show data from typical runs. In both cases there is evidence of stereotypical, extreme responses. Stronger animats almost always make an aggressive first move. Weaker animats are more likely to play a mixed strategy of sometimes advancing and sometimes retreating, and it seems reasonable to describe the advances as attempted bluffs.

If bluffing in the experimental condition was occurring successfully, then this would presumably result in a relatively high proportion of contests being won by weaker animats: bluffing is impossible in the blind control and pointless in the unfakeable. Figure 7 shows the percentage of contests in which the weaker animat gains the resource (note that the DFA results are not directly comparable as half of the time there *was* no weaker animat). While the experimental condition leads to more "upset wins" than the unfakeable condition (CFA case, $t_{18} = 5.93$, $p < 0.001$), the difference between the experimental and the blind condition is of marginal significance.

# 4 Discussion

Evolved behaviour in the experimental condition was more like the behaviour of "blind" animats

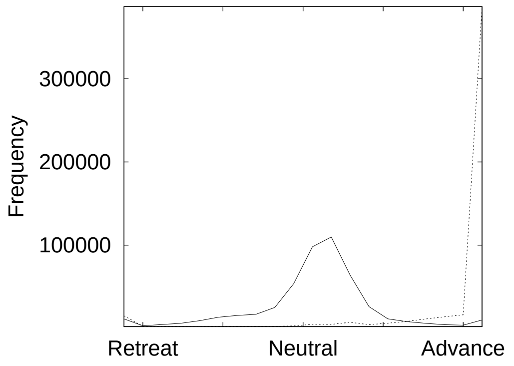

Figure 6: Total frequencies of opening moves by weak (solid line) and strong (dashed line) fighting ability, with the move data grouped into 20 bins for plotting. The plot shows data for the window period of run 5 in the experimental condition with DFA.

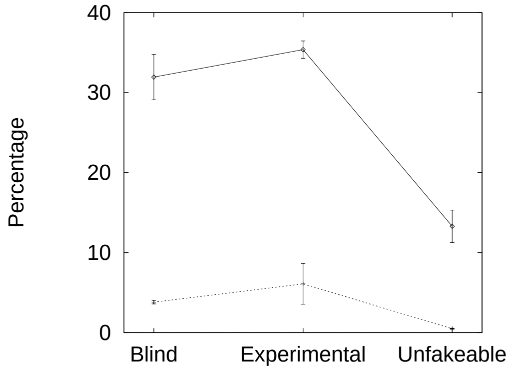

Figure 7: Mean percentage of contests won by the weaker animat ±1 s.e., for CFA (solid line) and DFA (dashed line), across the three conditions.

than that of animats able to perceive the strength of their opponent. The data suggest that if any signalling systems evolved in the experimental condition they were both inherently costly and no more effective in preventing unnecessary fights than blind strategies. The most obvious interpretation is that no signalling systems evolved. The mutual ability to perceive the intention movements of one's opponent in an aggressive interaction is, in the long run, of no fitness benefit.

The predictions of Enquist (1985) and Hurd (1997), that cost-free signalling of fighting ability could be an evolutionarily stable strategy, were certainly not supported. This was true whether or not fighting ability was discretely distributed. Caryl's (1987) charge that Enquist's model simply fails to capture the details of animal conflict may be the best explanation.

Wheeler and de Bourcier's (1995) model suggests that evolution will favour costly, honest signals of aggressive intent. The current simulation allowed for the possibility that stronger animats would use a costly signal to display their greater willingness to engage in all-out fights: e.g., by moving to $\Theta \gg A$. If such a signalling system had evolved, however, and if receivers were paying attention to it, we would expect there to be a reduction in the number of contests resolved by fighting (compared to the blind control). This was not observed, and the conclusion must be that Wheeler and de Bourcier's rather specific simulation does not generalize to the current contest model.

Adams and Mesterton-Gibbons's (1995) prediction, that the strongest animals would signal their strength and the weakest would attempt to bluff, was not borne out either, although high levels of bluffing were observed. Indeed, the data on bluffing strongly support the conventional game-theoretic view (Maynard Smith, 1982; Krebs & Dawkins, 1984) that participants in aggressive interactions will eventually come to pay little attention to each other's manipulative "signals". For example, in the CFA case (figure 5), the weakest animats start with an aggressive move about a third of the time. These bluffs at least occasionally result in the animat gaining the resource (figure 7), but to explain this we do not need to propose that the animat's opponent is paying any attention, because in the blind condition similar results are observed. The blind condition forces animats to develop strategies based only on knowledge of their own fighting ability, and the fact that they manage to avoid all-out conflicts most of the time reminds us that communication is not always necessary for coordination.

The findings also serve as a reminder that accounting for the use of apparently fakeable signals in aggressive interactions between animals stands as an open problem for biological optimization modelling.

As no signalling systems evolved, it is difficult to assess the implications of the work for the idea that intention movements can serve as raw material for signal evolution. However, in nature animals are obviously not given dedicated, artificial communication channels *deus ex machina.* Future simulation models of signalling will, by incorporating phenomena like intention movements, investigate what must happen when evolution co-opts an existing behaviour for the new purpose of signalling.

## References

Adams, E. S., & Caldwell, R. L. (1990). Deceptive communication in asymmetric fights of the stomatopod crustacean *Gonodactylus bredini*. *Animal Behaviour*, *39*, 706–716.

Adams, E. S., & Mesterton-Gibbons, M. (1995). The cost of threat displays and the stability of deceptive communication. *Journal of Theoretical Biology*, *175*, 405–421.

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.

Enquist, M. (1985). Communication during aggressive interactions with particular reference to variation in choice of behaviour. *Animal Behaviour*, *33*, 1152–1161.

Grafen, A. (1990). Biological signals as handicaps. *Journal of Theoretical Biology*, *144*, 517–546.

Hinde, R. A. (1981). Animal signals: Ethological and games-theory approaches are not incompatible. *Animal Behaviour*, *29*, 535–542.

Hurd, P. L. (1997). Is signalling of fighting ability costlier for weaker individuals?. *Journal of Theoretical Biology*, *184*, 83–88.

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.

Tinbergen, N. (1952). “Derived” activities; their causation, biological significance, origin, and emancipation during evolution. *Quarterly Review of Biology*, *27*(1), 1–32.

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.

Yamauchi, B. M., & Beer, R. D. (1994). Sequential behavior and learning in evolved dynamical neural networks. *Adaptive Behavior*, *2*(3), 219–246.

Zahavi, A. (1977). Reliability in communication systems and the evolution of altruism. In Stonehouse, B., & Perrins, C. M. (Eds.), *Evolutionary Ecology*, pp. 253–259. Macmillan, London.