Adaptive Factors in the Evolution of Signalling Systems

Jason Noble[1], Ezequiel A. Di Paolo[2], and Seth Bullock[1]

[1] Informatics Research Institute, University of Leeds, UK
jasonn@comp.leeds.ac.uk seth@comp.leeds.ac.uk,
www.iri.leeds.ac.uk/biosystems

[2] School of Cognitive and Computing Sciences, Unversity of Sussex, UK
ezequiel@cogs.susx.ac.uk, www.cogs.susx.ac.uk/users/ezequiel/

1 Introduction

Many of the chapters in this book have approached human language with an eye to its unique features, such as recursive syntax, or a large learned lexicon. We propose to take a wider view, seeing human language as one among many animal communication systems, and focusing on the selective pressures affecting the origin and maintenance of such systems. The possibility that human language arose from animal communication through a process of evolutionary change demands that we attend to the conceptual problems at the heart of our current understanding of animal signalling. In doing so we may throw light upon not only the origins of human language, but also its character.

The chief theoretical problem that comes to light when we look at the evolution of communication is accounting for the amount of honesty that is apparently involved (Johnstone, 1997; Noble, 2000a). Let us specify a hypothetical communicative scenario, such as a warning call to alert other animals about an approaching predator, or a display to advertise one's suitability as a mate. We can then construct a game-theoretic model, which allows us to consider the advantages and disadvantages of communicative and non-communicative behaviour in our scenario. There usually turns out to be a tempting payoff for cheats, liars, bluffers, or free-riders, which means that communication should not be evolutionarily stable. It can therefore be difficult to use the model to explain the apparent prevalence of real-world communication in the situation we are modelling.

The problem is sometimes solved by constructing a more subtle game-theoretic model; for example, we might take into account the effects of kin selection, and find that communication will be selected for, as long as signaller and receiver are from the same family group. However, if they are to remain tractable, game-theoretic and other mathematical models can only be made more elaborate up to a point. If we suspect that communication is occurring in a certain natural context, and yet the best game-theoretic model we can construct tells us that communication should not be stable, what are our options?

Moving to individual-based computational modelling lets us test the idea that stable communication may emerge from low-level details of space, time, and interactions between organisms that cannot be captured in a conventional game-theoretic model. We are most interested in evolutionary simulation models (Di Paolo, Noble, & Bullock, 2000), which involve the explicit modelling of individual organisms interacting in a shared environment. Evolution is incorporated in the sense that the more successful organisms (where success is defined by a criterion analogous to energy accumulation) will have a greater likelihood of passing on their genetic material to the next generation. Variation is introduced through mutation, i.e., the occasional random alteration of the transmitted genetic information.

The current chapter looks at simulations of the evolution of communication in the ecological domains of feeding, sexual choice, and contests over resources. We hope to demonstrate the power of individual-based evolutionary simulation modelling to explore more subtle hypotheses about signal evolution than is typically possible using conventional methods.

2 The Role of Ecological Feedback

There is an increasing recognition of the non-trivial effects that many ecological factors can have on evolution in general, and on the evolution of social behaviour such as communication in particular. For instance, it had been suggested that the effects of limited dispersal could enhance the local coefficient of relatedness between interacting individuals (Hamilton, 1964), thus facilitating the evolution of cooperative behaviour. However, contrary to this initial intuition, the effect of increased relatedness due to local interactions may be overcome by the effects of increased local competition when the scale of dispersal and the scale of interactions coincide (Taylor, 1992a, 1992b; Wilson, Pollock, & Dugatkin, 1992), a result that does not contradict the theory of kin selection if relatedness is properly calculated (Queller, 1994). This cancellation of positive and negative effects on fitness is challenged when the scales of dispersal and density regulation do not coincide (Kelly, 1992, 1994). It has also been shown that altruism may evolve in such viscous populations if organisms are modelled as *discrete entities* and the associated stochasticity is taken into account (Goodnight, 1992; van Baalen & Rand, 1998; Krakauer & Pagel, 1995; Nakamaru, Matsuda, & Iwasa, 1997). These cases can be regarded as evidence that ecological dynamics, which can include the effects of spatial situatedness, distribution of resources, mating strategies, and the activity of other species, can transform a simple evolutionary problem into a complex and non-intuitive one.

So far, the best way to approach such problems have proven to be a combination of traditional analytical models and individual-based computer simulations, in which factors such as space, discreteness of individuals, and noise can be included naturally. As an example of particularly non-intuitive ecological effects, we may consider the model introduced by Di Paolo (2000), in which the evolution of altruism in an action-response game is studied via a series of analytical and simulation approaches. Action-response games (e.g., Hurd, 1995) are fairly general models of social interaction which include signalling behaviour as a special case. The evolution of altruism in such games can, for instance, be equated to the evolution of honest signalling systems.

2.1 An action-response game

As in other cases, the game proposed by Di Paolo (2000) starts by assuming a situation of conflict of interest between two actors. The game affords various interpretations but a useful one from the communication point of view is the situation in which an animal (first player) has found a source of food and must decide (by choosing between two possible actions) whether to attract the attention of another conspecific (second player) to this source or to distract it. The second player has a choice of two actions: approaching the first player or ignoring it. Approaching the source of food results in both players sharing the payoff in equal measure, while if the first player manages to distract the attention of the second, it will have access to a larger than half share of the energy contained in it. The degree of conflict c represents the energy proportion that the first player gets in this situation. If $c = 0.5$ there is no conflict from the perspective of the first player; conflict is mild for values of c slightly greater than 0.5 and more significant as c approaches 1.

The word 'coordination' is used to describe the outcome of interactions that lead to the joint exploitation of the source. Without losing generality, it is possible to suppose that this happens in half of the four possibilities that the 2-signal-2-response scheme affords. Signals and responses can be either of types 'O' or 'E' (originally for 'odd' and 'even'), and the outcome of the interaction is denoted by a signal-response pair: OE or EO (in which cases coordination occurs), or OO or EE (in which cases the first player does not share the food).

The signal and response given by each individual player can be described by one of the above four strategies, which are genetically determined. During its lifetime, a given individual will play sometimes the first and sometimes the second role. The cooperation/conflict relationships between the four strategies are described in Figure 1. Each arrow is interpreted as connecting the initiator and responder strategies of those interactions that result in coordination. Thus a player using strategy 'OE' will behave altruistically only towards players using either 'OE' or 'EE' and this is indicated by the two arrows starting at the 'OE' node. An initiator playing 'EE', in contrast, will not behave altruistically towards players using 'OE' since there is no arrow from 'EE' to 'OE', although it will behave altruistically towards individuals playing 'OO' or 'EO'. In a randomly constituted population, the proportion of cooperative coordination will be 50%. In order to say that cooperative interactions have evolved, the proportion should rise above this value.

Notice the cyclic structure of part of the resulting graph indicating a kind of Rock-Paper-Scissors situation which, at first sight, suggests that no single strategy may become dominant because it will always be invaded by its 'neighbour' strategy in the graph. A detailed game-theoretic analysis of this game for the case of infinite, random mating and random playing populations (Di Paolo, 2000) leads precisely to this conclusion for all values of c. The constitution of

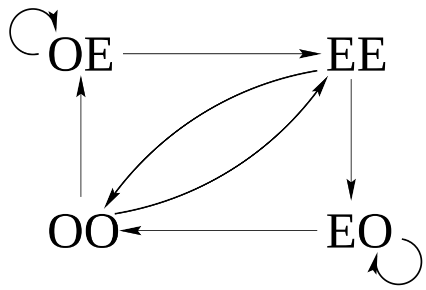

Fig. 1. Relations of cooperation/conflict between the four strategies in the game, see text for details.

the population will oscillate by following the straight arrows in Figure 1. There are no evolutionarily stable strategies in this case.

2.2 Broken symmetries

The above result changes for the case of a finite population and the introduction of noise to the evolutionary dynamics. In this case there is a single stable point attractor in which half of the population play strategy OO while the other half play strategy EE. However, the proportion of cooperative interactions at this equilibrium is only 50%, i.e., equal to the baseline level of cooperation. Thus cooperative coordination cannot evolve under these conditions.

Both these models remain quite abstract and further assumptions could be relaxed. For instance, the population could be considered to be distributed in space so that interactions, as well as reproductive events, are local. The distribution of energy in the food sources could also be described by a local variable, so that a kind of ecological coupling would be introduced resulting in differences in quality between local environments depending on how those environments are exploited. A continuous-time model of this situation, based on partial integro-differential equations, leads to the conclusion that players will tend to aggregate into discrete clusters even if they are uniformly distributed initially, but within those clusters the different strategies will oscillate as in the game-theoretic model with cooperative interactions once again at baseline levels.

This is as far the purely analytical approach can go. If more assumptions are to be relaxed, such as treating individual players as discrete entities instead of 'densities' in the distribution of strategies, an evolutionary simulation modelling scenario must be contemplated. In such an individual-based model each player accumulates energy by interacting with others and thus drawing energy from food sources. This amount of energy must not only cover the energy survival costs (same for all individuals) but must be enough for the player to eventually reproduce. Noise is present at different stages in the model, e.g., in the asynchronous updating scheme, in the choice of co-participants in the interaction, in

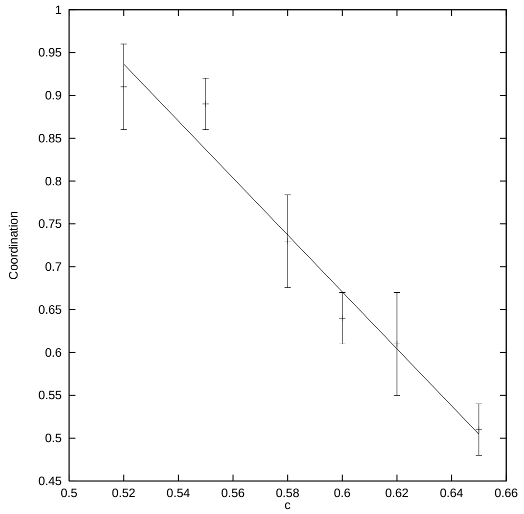

Fig. 2. Average value for the level of cooperative coordination for different values of c. Each point is the average of 5 simulation runs. The line represents a linear regression, (correlation coefficient: -0.982). Error bars indicate standard deviation.

the choice of food sources which may differ in their energy value, etc. (There is, however, no noise in the production of signals and responses.)

The simulation model can be explored both in the spatial and non-spatial cases. In the latter, the result are again oscillations. However, in the spatial case, the level of cooperative interactions depends on the degree of conflict c. For positive, but small, conflict (c slightly greater than 0.5) the population shows a high and stable level of games resulting in cooperative coordination. This level decreases linearly with c until it reaches the baseline level for $c \approx 0.65$, (Fig. 2).

This result is explained by looking at the spatial patterns that form spontaneously in the population and how they break many of the in-built symmetries of the abstract situation by allowing for reciprocal interactions between the evolving population and its environment.

As with the continuous spatial model, stable clusters can also be observed as a consequence of a dynamical equilibrium between two tendencies: the tendency for the population to concentrate in a small region so as to maximize the chances of finding a partner to interact with, and the tendency to move away from dense regions to places where local resources are exploited less frequently. Once a cluster is formed there is an equilibrium in the rate of energy consumption per unit of space. This equilibrium would seem to establish a degree of spatial 'neutrality' in the sense that spatial position does not matter for the rates of energy intake and offspring production. Players at the centre interact more often than those at the periphery but they do so for poorer resources. If the rates of net energy intake were different, the cluster would not be at equilibrium. Such a homogeneous rate of energy consumption (and reproduction) is indeed what is observed in the simulations. However, it is *not* true that spatial position is neutral in evolutionary terms.

If a player is born from a parent near the periphery of the cluster there is a high chance that it will be placed 'outside' the cluster in the sense that it will have a very small number of neighbours. Those players will tend to die before they reproduce. In fact, the chances of originating a lasting genealogy of players diminish as the originating position moves from the centre to the border of a cluster. This is a geometrical consequence of the stochastic and local character of the process of offspring allocation. Given this, we would predict that the position of a cluster's ancestors would tend to be concentrated near the centre of the cluster as one travels backwards in time, and this is what is observed. Thus, it is reasonable to conclude that a player's position within a cluster, although not under genetic control, plays a role with bearing on its fitness.

The above phenomenon is a case of symmetry breaking of the spatial homogeneity. Other symmetries are also broken by the centre/periphery structure such as the frequencies with which individuals play each role; with central individuals playing the role of responder more often.

By analysing the evolutionary stability of different strategies in view of these conditions it is possible to show that for low values of (positive) conflict, even though a cluster composed of altruistic strategists (OE or EO) can indeed be invaded locally by non-altruistic strategies, the environmental conditions, in terms of available energy and rates for role assignment are such that a very small increase in the local density of invaders renders them unviable in the central region. Invasions will occur locally but will be followed by the local disappearance of the invaders, leaving a gap at the centre of the cluster which either allows the altruists to re-invade or causes the cluster to split into two smaller ones, (see Figurer̃efgap). This effect is harsher near the centre of the cluster which is where most lineages originate.

It is important to point out that altruism, in this case in the form of honest signalling, is favoured by a combination of discreteness and ecological organization. Neither of these factors is sufficient on its own, as shown by the results of the continuous spatial model and the non-spatial individual-based model. The rupture of spatial homogeneity is essential for altruism to be favoured in the case of low conflict. But some of the ensuing broken symmetries occur only as a consequence of the discreteness of the players, for example, the dependence of the genealogy length on spatial position within a cluster. Discreteness also plays a role in the local extinction that may occur when non-altruistic players begin to invade the centre of an altruistic cluster. If sufficiently fine-grained density values were permitted within the model, such local extinctions would not occur. Rather, invading strategies would be allowed to take very small, but non-zero, density values corresponding to less than 1 individual in the region of interest. Because of their reduced energy consumption, these 'infinitesimal' individuals would be able to subsist in the unfavourable environment until eventually local energy would have been replenished and they would begin to increase in density.

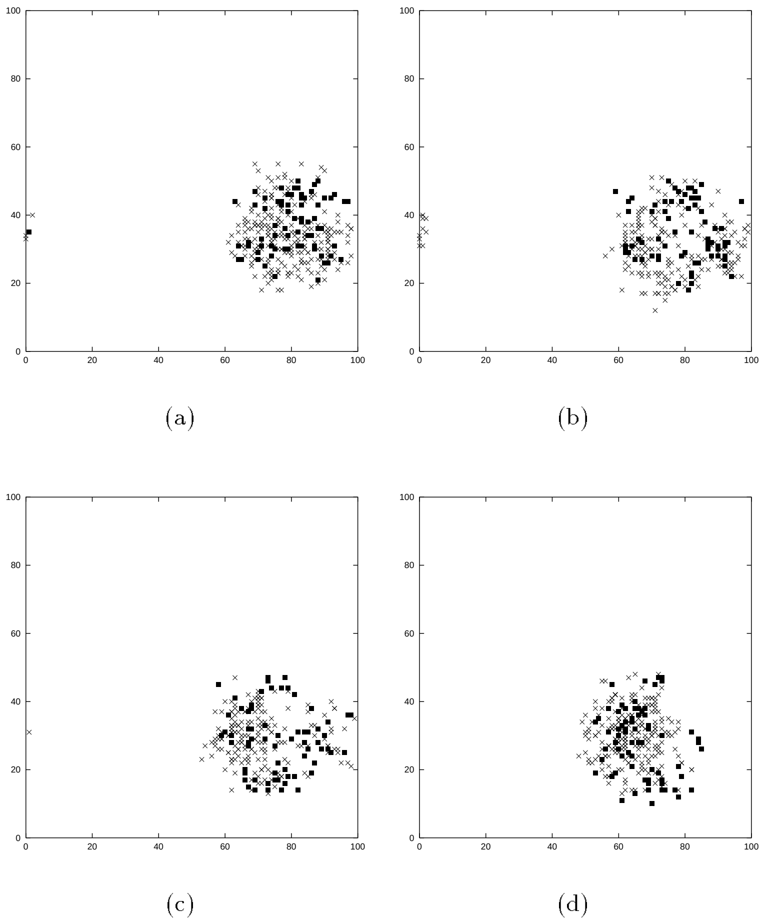

Fig. 3. Example of invasion dynamics of a predominantly altruistic cluster. Each frame shows the cluster constitution at intervals of 3000 time steps (about half an average individual lifetime); $c = 0.55$. Altruists are marked with x's and invaders with squares. The first frame (top-left) shows a high concentration of invaders in the central region. The following frame shows a density gap in the same central region due to the local disappearance of invaders. The bottom-left frame shows an expanded gap leading to a separation into two clusters, and only one of these survives in this case as shown in the last frame.

Even though there is no significant interaction between clusters, and, consequently, interpretations of these results in terms of group selection are not possible at this level, there is a sense in which such an interpretation could be reasonable for units of selection *within* a cluster (van Baalen & Rand, 1998). The local invasions of altruistic players, followed by local extinction, can indeed be interpreted in terms of the viability of two different groups in a specific local environment that one of them sets and the other cannot change fast enough to adapt to. On the flip-side of this interpretation, and also following van Baalen and Rand (1998), we could equally say that an appropriately defined coefficient of relatedness, taking into consideration density-dependent effects, would bring this result within the domain of Hamilton's rule. This is also a viable interpretation of the results, even though a simple estimation of purely genetic coefficients of relatedness (following Queller & Goodnight, 1989; Queller, 1994) was inconclusive in this respect. Finally, it would also be possible to construe these results as a consequence of reciprocal altruism (Trivers, 1971) although there would be little or no difference between this and the kin-selective interpretation since there is no segregation into different species in this model. However, a constant fact in all these possible interpretations remains the two-way coupling between selection and ecological dynamics and the resulting broken ecological symmetries due to the activity of the players. The importance of such couplings have been long noted (Lewontin, 1983) and have been recently highlighted under the label of 'niche-construction', (Laland, Odling-Smee, & Feldman, 2000; Bullock & Noble, 2000).

3 Evolutionary Simulation Modelling and the Handicap Principle

Since Israeli ornithologist Amotz Zahavi first presented his theory a quarter of a century ago (Zahavi, 1975, 1977), the *handicap principle* has been the subject of energetic debate within the evolutionary biology literature. Briefly, Zahavi suggested that extravagant displays such as unwieldy, colourful tail-feathers, or protracted bouts of exhausting bellowing, which are used by creatures throughout the natural world to advertise mate quality, fighting prowess, etc., may not be costly by accident, but because it is only through their extravagance that their trustworthiness is guaranteed.

Zahavi's insight was to suggest that the costs incurred in producing such displays might enforce honesty amongst signallers if these costs were somehow connected to the quality being advertised such that they favoured those signallers with more of whatever was being advertised (the best fighters, the highest quality mates, etc.). For example, an honest advertisement of a predator's ability to efficiently catch prey might be the extent to which the predator deliberately wastes food items. Wasting food is always costly, but it is more costly if you are unlikely to be able to get any more. Since poor predators cannot afford to waste hard-won prey items, a system in which predators demonstrate their ability through wasting as many food items as they can afford to cannot be invaded by cheats who exaggerate their ability, since the increased costs that this would entail ensure that bluffing is simply not worth their while.

This notion of waste as a signal of quality is reminiscent of the concept of "Conspicuous consumption" discussed by Thorstein Veblen (1899), a turn-of-the-century sociologist. Veblen noted that members of the "leisure class" persistently and protractedly overindulged themselves. They left expensive food uneaten, rarely wore their opulent clothes, and spent much of their time and money pursuing costly pastimes for no purpose other than their own entertainment. He suggested that this seemingly senseless hedonism was also a way in which the members of the leisure class demonstrated their class membership. That is, the purchase of prohibitively expensive goods and services could be understood as an indicator of the procurer's wealth. This index of societal status was an effective one because those of lower status could not afford to make the "advertisements" of which wealthier individuals were capable. Indeed, at the lower extreme of the scale, the funds of the poorest individuals were more than accounted for by the demands of simply staying alive, leaving no extra money to "waste" upon "unproductive consumption".

Initially, Zahavi's theory suffered considerable scepticism. Evolution by natural selection was understood to produce *efficient* systems — the opposite of the scenarios Zahavi described. Why would evolution favour wasteful exhibitionism? More specifically, if a peahen were to choose a mate on the basis of an advertisement, surely the advertisement (which her male offspring would likely to inherit) would not be chosen for its ability to seriously handicap its owner, increasing the likelihood that her offspring would die before themselves winning mates? If this was the price of honesty, surely it would make more sense to choose a mate at random and spare one's offspring the handicap? But despite these worries, and only intermittent empirical and theoretical support over the next decade, the handicap principle achieved increasing notoriety.

The rise of the handicap principle in the face of almost continuous criticism (e.g., Davis & O'Donald, 1976; Maynard Smith, 1976, 1978, 1985; Kirkpatrick, 1986) is perhaps attributable to two factors. The first is that Zahavi's theory filled a theoretical vacuum left by the collapse of group-selectionist accounts of signalling. Prior to the reassessment of group selection in the mid-sixties (Hamilton, 1964), the evolutionary function of signalling behaviours could be explained in terms of the benefits that they conferred upon a signalling community as a whole. Mating displays, aggressive posturing, informative dances, begging cries, warning coloration, and danger signals, if honest, enable the efficient distribution of resources (food, sex, shelter, etc.). This efficiency derives from the flow of useful information between the members of an honest signalling system — each member gains much of their information from other members, without having to collect it individually. Contrast a beehive, foraging as a unit on the basis of shared information, with the less efficient behaviour of the same bees obstinately foraging solo, or the difference between settling contests by honest aggressive displays of strength and settling the same disputes through fighting.

However, although the increased efficiency afforded by honest signalling is of benefit to those groups that employ such signals, *individuals* within these groups often stand to gain by freeloading, bluffing, cheating, lying, double-crossing, exaggerating, misleading, or crying wolf. For the individual, then, honesty is not always the best policy. With dishonesty comes mistrust, and eventually the collapse of an honest signalling system, undermined by deceit. But although the selfish actions of individuals were expected to compromise the stability of natural signalling systems, such systems appeared to be the frequent products of evolution. Signalling systems are everywhere in nature. If signalling systems are evolutionarily fragile, why are they so ubiquitous? Zahavi's handicap principle at least offered an explanation, even if it appeared counter-intuitive.

The second factor in Zahavi's favour was the rise of game-theoretic modelling in behavioural ecology (Maynard Smith, 1982). From 1985 onwards, a series of successful, game-theoretic models (Enquist, 1985, being the first, and Grafen, 1990, being the foremost amongst these) demonstrated the soundness of the handicap principle's central tenets, succeeding where population genetic models had previously failed (see Maynard Smith, 1985, for a review). As evolutionary game theory benefited from its success in dealing with ideas which had proven hard to explore using alternative modelling approaches, the handicap principle gained credibility. Although the handicap principle does not yet enjoy the status that Zahavi believes it deserves (having thus far failed to eclipse Darwin's theory of sexual selection), both the vocabulary and explanatory perspective associated with it have attained a central position within current evolutionary thinking.

However, while handicap thinking spreads within biology and beyond (e.g., Miller, 2000), the theoretical biology community face several unanswered questions. In this section, we will try to demonstrate how simulation models of the kind already introduced can help to answer these questions, and reveal new problems that have been neglected up until now. Three issues will be raised in the next sections, before an evolutionary simulation model with which to address them is introduced.

3.1 Balancing the handicap books

While the costs of signalling have clearly been the focus of work on the handicap principle, certain important aspects of these costs remain unclear in Zahavi's writing. Crucially, Zahavi's verbal arguments offer little clue as to the way in which handicap costs are perhaps balanced by the *benefits* to the signaller of achieving whatever goal the signal is intended to bring about. For each individual signaller, must handicap displays reduce their fitness (through loss of time and energy, increased risk of predation, etc.) to a greater extent than these signals on average increase it (through gaining copulations, victories, food, shelter, etc.)?

Zahavi sometimes appears to consider the *net* costs involved in signalling, when, for instance, he asserts that "it is reasonable to expect a population in its optimal fitness to benefit from a handicap" (Zahavi, 1977, , p.604). At least, then, at the population level, the costs of bearing handicaps are assumed to be more than compensated for by the associated benefits. At the level of the individual, matters are not as clear-cut, "so long as the [signaller] ... does not deviate to grow its handicap larger than it can afford, the handicap [may persist] as a marker of honest advertisement" (ibid.), i.e., handicap costs are limited in some way, but how? Compounding this vagueness, when describing natural examples, Zahavi rarely discusses the benefits obtained from signalling, and the manner in which these benefits balance the costs.

Furthermore, Zahavi's terminology is not easy to reconcile with a notion of the handicap principle couched in terms of net costs. For example, as Hurd (1995) and Getty (1998a, 1998b) point out, if the costs involved in signalling must be balanced by consonant benefits, then in what sense are these costs a 'handicap'? However, if these costs are not so balanced, what is the value of signalling? Although the exaggerated costs incurred by a *bluffer* might be characterized as a handicap, since these costs would be larger than the bluffer could afford, this is not the sense in which Zahavi proposed the term. For Zahavi, *honest* signallers suffer a handicap — they must do so in order to demonstrate their honesty.

Not surprisingly, this confusion has led authors to multiple interpretations of the handicap principle. (Wiley, 1983), for example, characterizes Zahavi's (1975) paper as claiming that "signals should evolve to become a *net* handicap to signallers" (p. 176, our emphasis). In contrast, Adams and Mesterton-Gibbons (1995) reach the opposite conclusion, stating that their model *differs* from the handicap principle in that "the net benefit for a given advertisement may not increase monotonically with the signaller's strength" (p. 406), implying that typical handicap thinking proposes that signallers gain a net *benefit* from signalling.

Later we will use an evolutionary simulation model to explore what the costs and benefits are for signallers that are involved in a handicap signalling system.

3.2 Need vs. quality?

A second, separate but related issue concerns the conflicting roles of signalling costs and signalling benefits in stabilising handicap signalling systems. Can, as Zahavi implies, honesty only be ensured by (gross) signalling costs varying such that some signallers stand to lose less from signalling than others and are thus able to signal more? Or might honesty also be maintained by (gross) signaller benefits varying such that some signallers stand to gain more from signalling than others and are thus able to signal more? Johnstone (1997) has usefully divided handicap models into these two kinds. The first attempts to account for the evolutionary stability of the honest advertisement of *quality* as the result of the manner in which the gross *costs* of signalling vary with quality (e.g., Grafen, 1990; Hurd, 1995). The second kind attempts to account for the evolutionary stability of the honest advertisement of *need* as the result of the manner in which gross signaller *benefits* vary with need (e.g., Godfray, 1991; Maynard Smith, 1991).

The latter kind of model includes that used by Godfray (1991) to demonstrate the evolutionary stability of a strategy in which nestlings honestly advertise their hunger (need) by varying the strength of their begging calls. Godfray showed that such a strategy is evolutionarily stable if the costs of begging are the same for all chicks, but the value of any particular parental resource to a begging chick increases with the chick's hunger. In such situations, hungry chicks beg more than satiated chicks because the resources are worth more to them.

The former kind of model include's Grafen's (1990) treatment of a similar scenario, in which a very different stable begging equilibrium was derived. If we assume that the parent wishes to feed the *highest quality* chicks rather than the most needy, Grafen (1990) showed that we can expect chick begging to be an honest indicator of quality if the value of parental resources are the same for all chicks, but the cost of any particular begging display is greater for the lesser quality chicks. In such situations, high quality chicks beg more than lower quality chicks because the signals are more affordable to them.

Are these two scenarios distinct, though complementary, classes of handicap signalling, or two extreme cases from a wider range of possibler signalling systems?

3.3 The attainability of honesty

Until now, we have been concerned with arguments from theoretical biology concerning whether signalling systems can be evolutionarily *stable*. Since evolution has been continuing for billions of years, theoretical biologists assume that the systems we see around us are stable, if not they would most likely have been replaced by some other system that was stable. Since signalling systems are so prevalent and so widespread, it is hard to imagine that each is unstable — in a state of evolutionary flux, poised at the brink of collapse (although some have pursued this idea, Dawkins & Krebs, 1978; Krebs & Dawkins, 1984). However, there is nothing to prevent an evolving system from admitting of several different evolutionary stable situations. In fact, it is becoming clear that many if not most interesting evolutionary systems feature *multiple equilibria* of this kind.

Evolutionary simulation models are well-suited for addressing this issue. Whereas existing formal modelling paradigms (game theoretic models and population genetic models, for example) are able to disentangle the contributions of the various ideas and theories discussed in the previous two sections, evolutionary simulation models are ideally positioned to deal with matters of evolutionary change, modelling as they do the manner in which populations subjected to evolutionary pressures change over evolutionary time. Amongst other things, such models allow us to explore questions of *equilibrium selection* — which of a number of possible equilibrium states will an evolving system reach from some initial ancestral condition?

Here we will compare two different conclusions that may be drawn from the empirical observation that fuelled Zahavi's initial papers introducing the handicap principle, and use an evolutionary simulation model to decide between them.

Observation: many natural signalling systems appear to feature "extravagant" signals.

Conclusion 1: costliness is necessary to ensure the *stability* of honest signalling systems.

Conclusion 2: costliness is necessary to ensure the *attainability* of honest signalling systems.

The first conclusion has been widely explored in the theoretical biology literature. While it has been shown that signal costs are able to stabilise signalling

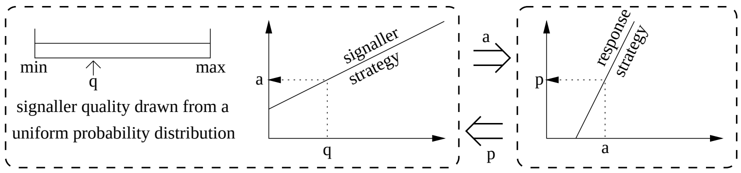

Fig. 4. A simpe signal-response game in which a signaller is allocated an internal state (q, sometimes referred to as quality) at random, and produces an advert (a) with magnitude determined by an inherited signalling strategy (in this case a linear mapping). Adverts may not have negative sign. This advert is passed to the receiver, who produces a response (p) determined by an inherited response strategy (in this case also a linear mapping). Responses are truncated to lie within the range $[q_{min}, q_{max}]$. Receiver fitness is calculated as $\frac{1}{1+|p-q|}$, increasing with the accuracy with which the response matches the signaller's internal state. Signaller fitness is calculated as $pq^R - aq^S$. See text for explanation.

systems, it is unclear whether these costs are "extravagant" or "handicaps" in the sense implied by Zahavi's papers. The second conclusion has been largely unexplored (although Yachi, 1995, has attempted to characterise the conditions under which handicap signalling might evolve).

3.4 An evolutionary simulation model

To recap, the evolutionary simulation model presented here was designed to explore three issues: (i) what is the character of handicap signalling at equilibrium, and how does this character vary with the model's parameters? (ii) what conditions must be met in order for handicap signalling to be evolutionarily *stable* given that both signalling costs and signaller benefits vary with signaller state? (iii) are handicap signalling equilibria *attainable* from appropriate initial conditions, and how does this attainability vary with the model's parameters? For present purposes, a brief description of the model will be given. Full details of the model can be found in Bullock (1997, 1998).

For each evolutionary run, a population of signallers and a population of receivers coevolved for 1000 generations. Fitness scores were determined by pairing up signallers with receivers and allowing them to play a simple signal-response game (see Figure 4). Receiver fitness was awarded proportional to the accuracy with which receiver response, p, matched the internal state of the signaller, q. Signaller fitness was calculated in a slightly more complicated fashion as the benefit of obtaining a response, pq^R, minus the cost of signalling, aq^S, where R and S are model parameters fixed for the course of an evolutionary run. The interests of signallers and receivers conflict, since signallers always benefit from as large a response as possible, whereas receivers benefit from matching their response to a signaller's internal state.

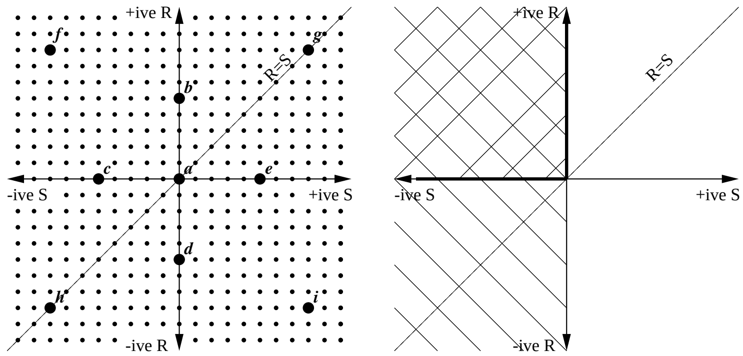

Fig. 5. *Left* — The model's R-S parameter space was extensively sampled. Nine of the $21 * 21$ scenarios explored are labeled. *(a)* $R = 0$, $S = 0$: signaller state does not affect the cost of making an advert, nor the benefit of gaining a response. *(b)* +ive R, $S = 0$: the value of a response increases with singaller quality, the cost of advertising is independent of signaller quality. *(c)* $R = 0$, −ive S: the value of a response is independent of signaller quality, whereas the cost of advertising decreases with signaller quality. *(d)* −ive R, $S = 0$: as *(b)*, but the value of responses decreases with increasing singaller quality. *(e)* $R = 0$, −ive S: as *(c)*, but the cost of advertising increases with signaller quality. *(f)* +ive R ,−ive S: responses are more valuable to high quality signallers, who also pay less for any given advert. *(g)* and *(h)*: as with *(a)*, the effect of signaller quality on signal cost is balanced by its effect on the value of responses. *(i)* −ive R, +ive S: not only is signalling more costly for high quality signallers, but they also gain less benefit from a receiver response. *Right* — Previous models' predictions of parameter values which support honest handicap signalling equilibria: Zahavi (1975, 1977) diagonal hatching; Grafen (1990) cross-hatching; Godfray (1991) and Maynard Smith (1991) bold vertical line; and Hurd (1995) bold horizontal line. The current model predicts honest signalling equilibria will exist in the part of parameter space lying above the line $R = S$.

In order to address the three issues raised above, the evolutionary dynamics of this scenario were explored under a range of different cost-benefit parameters and from a variety of initial conditions. Crucially, we need to manipulate both (i) the manner in which signaller quality influences the negative effect of signal cost on signaller fitness and (ii) the manner in which signaller quality influences the positive effect of receiver response on signaller fitness. The signaller fitness function was designed such that these two manipulations could be achieved by varying two model parameters, S and R, respectively (see Figure 5 Left).

The right-hand panel of Figure 5 depicts areas of the model's parameter space that previous models have predicted will support evolutionarily stable honest signalling. For runs in which $S < 0$ (diagonal hatching), the costs of advertising decrease with signaller quality — this is the condition predicted to guarantee honesty by Zahavi's handicap principle (1975, 1977). Several models have supported Zahavi, in suggesting that portions of this area of the parameter space admit of honest signalling equilibria. However, analysis of the model presented here (Bullock, 1998) suggests that honest signalling will only be stable for scenarios in which $R > S$. While this finding is not incomensurate with previous those of previous models, it contradicts Zahavi's basic premise that the manner in which signalling cost varies with signaller state (i.e., the value of S in this model) is all that determines whether handicap signalling is stable or not.

Evolutionary runs were carried out from three kinds of initial condition. (i) Honest: initially signallers played $a = q$, while receivers played $p = a$, (ii) Random: initially signaller and receiver strategies were determined at random, (iii) Mute-Deaf: initially signallers played $a = 0$, while receivers played $p = 0$. After 1000 generations, each run was terminated and the state of the evolved populations examined. In this way the evolutionary simulation model was used to explore the evolutionary dynamics of a range of scenarios, and, for each scenario, to assess whether stable handicap signalling equilibria could be achieved from a variety of initial conditions.

The simulation results (see Figure 6) supported the analytical results in that no signalling behaviour was observed for scenarios in which $R \not> S$, whereas signalling equilibria were observed for all scenarios in which $R > S$. In addition, where a signalling equilibria were discovered, both the character of the signalling at these equilibria and their attainability, varied with the relationship between R and S.

For scenarios in which R is only slightly larger than S, stable signalling systems exhibited a relatively small range of signals, with the largest signals exhibited themselves being relatively small. These signalling equilibria were also associated with relatively small basins of attraction, which ensured that evolution did not tend to reach them from Random or Mute-Deaf initial conditions.

In contrast, for scenarios in which $R \gg S$, signalling equilibria exhibited a very wide range of signal sizes, with the largest signals being orders of magnitude more massive than the smallest. Furthermore, the basins of attraction for these equilibria were also much larger (and hence more easily attainable from Random or Mute-Deaf initial conditions) than those discovered for scenarios in which R is only slightly larger than S.

3.5 Discussion of the model

These findings have several implications for our understanding of the handicap principle and how it affects the character of natural signalling systems. First, rather than conclude from the existence of seemingly extravagant natural signalling systems that extravagance is necessary in order that such systems remain evolutionarily *stable*, we might now surmise that these observations are due to the relative ease with which such signalling systems are *attained* by evolving populations of signallers and receivers. The simulation model above suggests that although a wide range of stable signalling systems are possible, some featuring relatively restrained signals, while some feature larger and perhaps seemingly extravagant displays, it is only the latter that are easily achieved from non-signalling ancestral scenarios. These results suggest that we may find examples

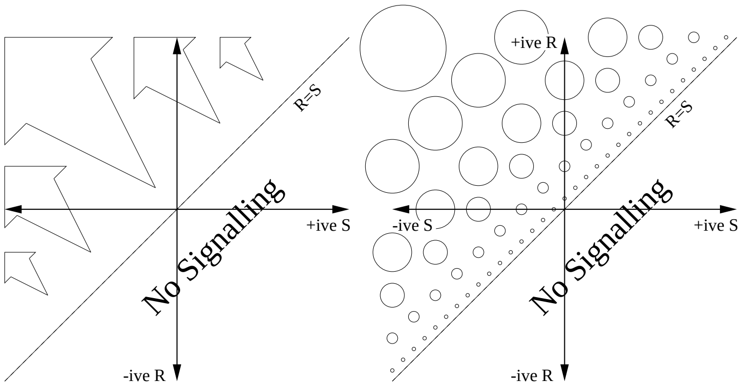

Fig. 6. The evolutionary simulation model reveals that signalling equilibria exist for scenarios in which $R > S$, but not otherwise. *Left* — The range of signalling exhibited at equilibrium, and the magnitude of the largest adverts, increases exponentially as R outstrips S. *Right* — The size of the basins of attraction for signalling equilibria also increase as R outstrips S.

of subtle signalling systems which nevertheless are evolutionarily stable under conditions in which the *net* cost of signalling decreases only slowly with signaller quality (i.e., R is only slightly larger than S). Such systems might be those in which either the resource being signalled for is itself of limited (and relatively constant) value, or in which the cost of advertising do not vary to a great extent with the property being advertised. For instance, the "I-see-you" signals made by a small bird to a stalking cat might be one such system. Despite involving subtle signals that do not appear costly to the casual observer, such a scenario would still be a handicap system, stabilised by the net costs of signalling.

A second implication of the model is that it is *net* costs that must be considered when dealing with handicap signalling systems. Contra Zahavi, signalling equilibria were sometimes exhibited by the model under conditions in which making signals was *more* costly for signallers of high-quality ($S > 0$). Similarly, and again contra Zahavi, signalling equilibria *fail* to exist for some model scenarios which meet Zahavi's handicap criterion, i.e., in which signal costs decrease with signaller quality. Only once consideration is given to the *balance* between the manner in which signal costs vary with quality, and the manner in which signaller benefits also vary with quality, can the distribution of honest signalling equilibria across then model's parameter space be understood.

4 Signalling in Contests

Animal contests — disputes over resources such as food, territory or mates — are good examples of interactions in which the interests of the participants seem to be maximally opposed. This is particularly true of struggles over the control of an indivisible item: one's gain is necessarily another's loss. Nevertheless, 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).

What is happening in these cases? Are the competing animals likely to be exchanging honest signals, informing each other of their fighting ability or their intention to attack? (And if not, what is the function of their aggressive displays?) Intuitively, settling contests by signalling makes sense. We can see that an all-out fight is usually a bad idea: fighting is energetically expensive, and there is always a risk of injury or death. The early ethologists suggested that threat displays were honest signals of aggressive intent that benefited the species by preventing costly fights, but the naive group-selectionist overtones of this idea mean that it is no longer taken seriously. Moreover, standard game-theoretic predictions (Maynard Smith, 1982) suggest that in contest situations, it will not be evolutionarily stable for animals to exchange signals of strength or aggressive intent because would-be honest signallers will always be less fit than bluffers. According to this perspective, there is no room in the arena of animal contests for the co-operative exchange of arbitrary signals; the aggressive displays observed in nature are either unfakeable because of physical constraints, or are the uninformative result of a manipulative arms race (Krebs & Dawkins, 1984).

On the other hand, some theorists have argued that, in effect, competing animals share enough of a common interest in avoiding serious injury that honest signalling can be evolutionarily stable. 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, referred to as strategy S, will be evolutionarily stable. Fights will occur only between evenly-matched opponents, and weak animals will defer to signals denoting strength.

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 (and in this sense Enquist's model can be considered a handicap signalling model). However, Caryl (1987) 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

[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.

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.

We will now look at Enquist's argument in the light of an evolutionary simulation model of contests over an indivisible resource (Noble, 2000b). The aim of using a simulation is to avoid oversimplification. In particular, time will be modelled in an approximately continuous fashion: in Enquist's model there are only two time-steps — an exchange of signals followed by a choice of actions — and thus the model may fail to capture critical aspects of real-time interactions. In a more 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?

4.1 The model

The simulation will be described only briefly; full details are given in Noble (2000b). The contests commence with two players facing each other in a one-dimensional arena (Figure 7). Each player has a fighting ability 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. At each timestep, a player can move forwards or backwards by up to a metre, and can produce an auditory signal of variable intensity.

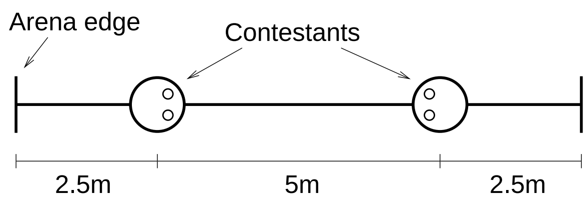

Fig. 7. Starting positions of the two competing players. The arena is 10m wide; one player starts 2.5m from the left edge, the other 2.5m from the right edge.

Contests can end due to one player fleeing, one player seriously injuring the other, or because a time limit has been reached. If one player moves backwards far enough to leave the arena, the other obtains the food resource and the contest is over. If both players leave the arena simultaneously, neither obtains the food item. If the two players 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 player attacks another, it may succeed in inflicting a serious injury with a probability proportional to its fighting ability. If a serious injury occurs, the contest ends at once and and the injured player forfeits the resource.

The players have access to 11 perceptual inputs, on which they can base their decisions about movement and possible signalling. Briefly, each player has access to privileged information about its own state, such as its fighting ability and its energy level, and can observe the signals and movements of the other player. Players are also aware of their position in the arena.

The perceptual inputs are translated into outputs (movement and signalling) using a production system. A player's production system specifies its contest strategy; the system is also the genotype that will be passed on if the player succeeds in reproducing. Production systems include up to six rules; a typical rule might be "If (own energy level $<$ opponent's signal) and (time elapsed $<$ 72 timesteps) then advance 32cm and signal 85%." At each timestep in a contest, a player's production system is given fresh inputs, and the movement and signalling outputs are provided by the *first* rule to have all of its conditions met. In case no rules fire, the genotype includes default movement and signalling values.

The simulation is organized into days. Each day, every player in the population plays out a contest with a random opponent. The result is that each player 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 players with more than a threshold level of energy are randomly paired up and allowed to reproduce sexually. Each pair produces a single offspring.

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. The average cost of being seriously injured is approximately double the value of the contested resource. Furthermore, the cost to a weak player of an extended fight with a strong player is high enough that Enquist's conditions for stable communication of fighting ability are met.

4.2 Simulation results

Basic model Genotypes in the initial population were generated randomly. Each simulation run continued for the equivalent of about 7,500 generations. Ten evolutionary runs were performed, each with a different random seed value.

Contests lasted, on average, 19.4 timesteps. This indicates that the players 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 players fleeing, and 24.2% of the time through one player inflicting a serious injury on the other. The remaining 9.8% of contests reached the maximum time limit.

After 7,500 generations, the median strategy was recorded for each population. A common theme across all 10 strategies was an aggressive default movement, combined with at least one rule spelling out conditions under which the player would retreat. Figure 8 illustrates this by showing one of the simpler strategies in full. A player 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 players 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 player is relatively weak (i.e., its fighting ability is less than 53% of the maximum value) then it will retreat 93cm, as for this player the distance to the edge will be 4.5m. If the second player 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 players will clash head on; neither will ever retreat. The behaviour of two weak players is more interesting: they face each 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.

```
If Own fighting ability < 53%
   Distance to opponent < 1.70m
   Distance to edge < 4.57m
   then retreat 93cm and signal 47%

Otherwise advance 1m and signal 80%
```

Fig. 8. The median strategy evolved in run 9. The default movement is maximally aggressive and the single rule specifies conditions for retreating.

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 players all playing this strategy, no-one will be listening. However, there is some indication that players may be signalling, or at least giving away information, through their movements. When weaker players 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 4.3.

Space precludes a detailed analysis of all 10 of the evolved strategies. However, if we look at the sensory inputs the players actually used in their decision making, we find that the most popular were the distance to the opponent, one's own fighting ability, the distance to the edge of the arena, and the *change* in the distance to the oppponent. The sensory inputs associated with the opponent's signalling activity were attended to only infrequently.

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.

```
If Own fighting ability < 40%
   Opponent's signal > 50%
   then retreat 1m and signal 0%

If Own fighting ability < 40%
   then advance 1m and signal 0%

Otherwise advance 1m and signal 100%
```

Fig. 9. An analogue of Enquist's (1985) strategy S, expressed in the framework of the players' production system. The default strategy is an aggressive advance and a loud signal. The first rule specifies that weaker players will retreat from a loud signal, and the second, that they will advance without signalling if they hear no signal.

Figure 9 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 40% as this was the approximate mean fighting ability implemented in the 10 runs described in section 4.2. For stronger players, the chosen action will always fall through to the default behaviour of aggressively advancing while making a loud signal. For weaker players, rule one or rule two will always fire. This means that weaker players will announce their status by always signalling with zero intensity. If a weak player 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.

The evolutionary stability of strategy S was investigated by conducting another 10 runs, with players in the initial populations set to play strategy S. These simulations can therefore show us whether or not a population of strategy S players is resistant to invasion by mutant strategies. Looking at the proportion of the time that various sensory inputs were used to make decisions in the evolved players, it became clear that strategy S was not able to resist the invasion of alternative strategies: for example, the "Distance to opponent" input was used most often, despite not being present in the initial population. Inspection of the median strategies showed them to be very much like those that evolved in the basic model, with any signalling behaviour on the part of the opponent being largely ignored.

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? 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 contestant. The results discussed so far present a different picture, in which weak players 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 players 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 contestant of facing up to a stronger one). Consider the pattern of behaviour outlined for the strategy shown in Figure 8. Clearly, if weak players can bluff it out against stronger opponents, up to a point, and then retreat without being harmed, then d is not particularly high.

4.3 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 dependent 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. These findings support and extend Caryl's (1987) claim that Enquist's model of animal combat is implausible.

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 models are leaving something out. But unfortunately Enquist (1985) settled on some questionable assumptions in his attempt to explain apparent honesty in contests.

The available signalling channel was not used by the players, 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 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 players using the strategy shown in Figure 8 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 player retreats. Information is conveyed to the opponent by this behaviour, but it is not the function of the retreat to be informative.

5 Overall Conclusion

Conventional game-theoretic models in biology abstract away from the individual organism and incorporate radical simplifying assumptions such as random mating in homogeneous populations, the absence of spatial distribution, and the lack of significant ecological feedback. Evolutionary simulation models are able to highlight the importance of many of these assumptions through exploring their contribution to a model's evolutionary dynamics. Study of the evolution of communication and language is just one domain of enquiry that is crucially concerned with interactions between individuals mediated by an environment. It is hoped that the individual-based evolutionary simulation models presented here demonstrate the methodological value of taking a comparative modelling approach to problems of this kind.

Finally, it is worth stressing that simulation results are no substitute for empirical evidence. If a simulation establishes the plausibility of a hypothesis, this is not the same as establishing its truth. The claim here is only that simulation methods can demonstrate the logical coherence (or indeed incoherence) of a particular model, and that they may suggest new hypotheses for empirical investigation (see Di Paolo et al., 2000, for a more complete treatment of these issues).

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.

Bullock, S. (1997). *Evolutionary Simulation Models: On their Character, and Application to Problems Concerning the Evolution of Natural Signalling Systems*. Ph.D. thesis, School of Cognitive and Computing Sciences, University of Sussex, Brighton, UK.

Bullock, S. (1998). A continuous evolutionary simulation model of the attainability of honest signalling equilibria. In Adami, C., Belew, R., Kitano, H., & Taylor, C. (Eds.), *Artificial Life VI*, pp. 339–348. MIT Press, Cambridge, MA.

Bullock, S., & Noble, J. (2000). Evolutionary simulation modelling clarifies interactions between parallel adaptive processes. *Behavioral and Brain Sciences*, **21**(1), 150–151. Commentary.

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.

Davis, J. W. F., & O'Donald, P. (1976). Sexual selection for a handicap: A critical analysis of Zahavi's model. *Journal of Theoretical Biology*, **57**, 345–354.

Dawkins, R., & Krebs, J. R. (1978). Animal signals: Information or manipulation?. In Krebs, J. R., & Davies, N. B. (Eds.), *Behavioural Ecology: An Evolutionary Approach*, pp. 282–309. Blackwell, Oxford.

Di Paolo, E. A. (2000). Ecological symmetry breaking can favour the evolution of altruism in an action-response game. *Journal of Theoretical Biology*, **203**, 135–152.

Di Paolo, E. A., Noble, J., & Bullock, S. (2000). Simulation models as opaque thought experiments. In Bedau, M. A., McCaskill, J. S., Packard, N. H., & Rasmussen, S. (Eds.), *Artificial Life VII: Proceedings of the Seventh International Conference on Artificial Life*, pp. 497–506. MIT Press, Cambridge, MA.

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

Getty, T. (1998a). Handicap signalling: when fecundity and viability do not add up. *Animal Behaviour*, **56**(1), 127–130.

Getty, T. (1998b). Reliable signalling need not be a handicap — Commentary. *Animal Behaviour*, **56**(1), 253–255.

Godfray, H. C. J. (1991). Signalling of need by offspring to their parents. *Nature*, **352**, 328–330.

Goodnight, K. F. (1992). The effect of stochastic variation on kin selection in a budding-viscous population. *American Naturalist*, **140**, 1028–1040.

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

Hamilton, W. D. (1964). The genetical evolution of social behaviour, I and II. *Journal of Theoretical Biology*, **7**, 1–16; 17–32.

Hansen, A. J. (1986). Fighting behaviour in bald eagles: A test of game theory. *Ecology*, **67**, 787–797.

Hurd, P. L. (1995). Communication in discrete action-response games. *Journal of Theoretical Biology*, **174**, 217–222.

Johnstone, R. A. (1997). The evolution of animal signals. In Krebs, J. R., & Davies, N. B. (Eds.), *Behavioural Ecology: An Evolutionary Approach* (Fourth edition)., pp. 155–178. Blackwell, Oxford.

Kelly, J. K. (1992). Restricted migrations and the evolution of altruism. *Evolution*, **46**, 1492–1495.

Kelly, J. K. (1994). The effects of scale dependent processes on kin selection: Mating and density regulation. *Theoretical Population Biology*, **46**, 32–57.

Kirkpatrick, M. (1986). The handicap mechanism of sexual selection does not work. *American Naturalist*, **127**, 222–240.

Krakauer, D. C., & Pagel, M. (1995). Spatial structure and the evolution of honest cost-free signalling. *Proceedings of the Royal Society of London, Series* B, **260**, 365–372.

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* (2nd edition)., pp. 380–402. Blackwell, Oxford.

Laland, K. N., Odling-Smee, J., & Feldman, M. W. (2000). Niche construction, biological evolution and cultural change. *Behavioral and Brain Sciences*, **21**, 131–146.

Lewontin, R. C. (1983). Gene, organism, and environment. In Bentall, D. (Ed.), *Evolution from molecules to men.* Cambridge University Press.

Maynard Smith, J. (1976). Sexual selection and the handicap principle. *Journal of Theoretical Biology*, **57**, 239–242.

Maynard Smith, J. (1978). The handicap principle — a comment. *Journal of Theoretical Biology*, **70**, 251–252.

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

Maynard Smith, J. (1985). Mini review: Sexual selection, handicaps and true fitness. *Journal of Theoretical Biology*, **115**, 1–8.

Maynard Smith, J. (1991). Honest signalling: The Philip Sidney game. *Animal Behaviour*, **42**, 1034–1035.

Miller, G. F. (2000). *The Mating Game: How sexual choice shaped the evolution of human nature.* William Heinemann, London.

Nakamaru, M., Matsuda, H., & Iwasa, Y. (1997). The evolution of cooperation in a lattice-structured population. *Journal of Theoretical Biology*, **184**, 65–81.

Noble, J. (2000a). Cooperation, competition and the evolution of prelinguistic communication. In Knight, C., Studdert-Kennedy, M., & Hurford, J. (Eds.), *The Emergence of Language*, pp. 40–61. Cambridge University Press.

Noble, J. (2000b). Talk is cheap: Evolved strategies for communication and action in asymmetrical animal contests. In Meyer, J.-A., Berthoz, A., Floreano, D., Roitblat, H., & Wilson, S. W. (Eds.), *From Animals to Animats 6: Proceedings of the Sixth International Conference on Simulation of Adaptive Behavior*, pp. 481–490. MIT Press, Cambridge, MA.

Queller, D. C. (1994). Genetic relatedness in viscous populations. *Evolutionary Ecology*, **8**, 70–73.

Queller, D. C., & Goodnight, K. F. (1989). Estimation of genetic relatedness using genetic markers. *Evolution*, **43**, 258–275.

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.

Taylor, P. D. (1992a). Altruism in viscous populations - an inclusive fitness approach. *Evolutionary Ecology*, **6**, 352–356.

Taylor, P. D. (1992b). Inclusive fitness in a homogeneous environment. *Proceedings of the Royal Society of London, Series* B, **249**, 299–302.

Trivers, R. L. (1971). The evolution of reciprocal altruism. *Quarterly Review of Biology*, **46**, 35–57.

van Baalen, M., & Rand, D. A. (1998). The unit of selection in viscous populations and the evolution of altruism. *Journal of Theoretical Biology*, **193**, 631–648.

Veblen, T. (1899). The theory of the leisure class. In Lerner, M. (Ed.), *The Portable Veblen*, pp. 53–214. Viking Press, New York. Collection published 1948.

Wiley, R. H. (1983). The evolution of communication: Information and manipulation. In Halliday, T. R., & Slater, P. J. B. (Eds.), *Communication*, Vol. **2** of *Animal Behaviour*, pp. 156–189. Blackwell, Oxford.

Wilson, D., Pollock, G. B., & Dugatkin, L. (1992). Can altruism evolve in purely viscous populations?. *Evolutionary Ecology*, **6**, 331–341.

Yachi, S. (1995). How can honest signalling evolve? The role of handicap principle. *Proceedings of the Royal Society of London, Series* B, **262**, 283–288.

Zahavi, A. (1977). The cost of honesty (further remarks on the handicap principle). *Journal of Theoretical Biology*, **67**, 603–605.

Zahavi, A. (1975). Mate selection — A selection for a handicap. *Journal of Theoretical Biology*, **53**, 205–214.