PHILOSOPHICAL TRANSACTIONS OF THE ROYAL SOCIETY B BIOLOGICAL SCIENCES

The Evolution of Exploitation and Honesty in Animal Communication: A Model Using Artificial Neural Networks

David C. Krakauer and Rufus A. Johnstone

Phil. Trans. R. Soc. Lond. B 1995 **348**, 355-361
doi: 10.1098/rstb.1995.0073

The evolution of exploitation and honesty in animal communication: a model using artificial neural networks

DAVID C. KRAKAUER[1] AND RUFUS A. JOHNSTONE[2]

[1] *BBSRC NERC Ecology and Behaviour Group, Department of Zoology, University of Oxford, South Parks Road, Oxford OX1 3PS, U.K.*
[2] *Department of Zoology, University of Cambridge, Downing Street, Cambridge CB2 3EJ, U.K.*

SUMMARY

Conflicts of interest arise between signaller and receiver in most kinds of biological communication. Some authors have argued that this conflict is likely to give rise to deceit and exploitation, as receivers lag behind in the coevolutionary 'arms race' with signallers. Others have argued that such manipulation is likely to be short-lived and that receivers can avoid being deceived by paying attention to signals that are costly and hence 'unfakeable'. These two views have been hard to reconcile. Here, we present results from simulations of signal evolution using artificial neural networks, which demonstrate that honesty can coexist with a degree of exploitation. Signal cost ensures that receivers are able to obtain some honest information, but is unable to prevent exploitative signalling strategies from gaining short-term benefits. Although any one receiver bias that is open to exploitation will subsist for only a short period of time once signallers begin to take advantage of it, new preferences of this kind are constantly regenerated through selection and random drift. Hidden preferences and sensory exploitation are thus likely to have an enduring influence on the evolution of honest, costly signals. At the same time, honesty and cost are prerequisites for the evolution of exploitation. When signalling is cost-free, selection cannot act to maintain honesty, and receivers rapidly evolve to ignore signals. This leads to a reduction in the extent of hidden preference, and a consequent loss of potential for exploitation.

1. INTRODUCTION

Communication has implications for the fitness of both signaller and receiver. Biological signals are thus the product of coevolution between two parties whose interests may conflict. Dawkins & Krebs (1978; Krebs & Dawkins 1984) were among the first to point out that this conflict can give rise to deceit and exploitation. In their terms, the evolution of a signalling system is an 'arms race' between the signaller as manipulator and receiver as mind reader. If receivers lag behind in the race, their outdated response strategies are vulnerable to exploitation. Furthermore, demands are placed on receiver sensory systems from many activities besides communication (e.g. detection of prey and predators). Signallers can thus exploit neural mechanisms of receivers that have evolved for other purposes, a possibility referred to as 'sensory exploitation' (Ryan 1990; Ryan & Rand 1990, 1993; Kirkpatrick & Ryan 1991; Endler 1992, 1993). In the water mite *Neumania papillator*, for instance, males vibrate their first two pairs of legs to draw the attention of females, and Proctor (1991, 1993) has suggested that this behaviour exploits the female's sensory adaptations for detecting the vibrations of copepod prey. This is not how we treat exploitation in the present paper. While 'sensory bias' is the dominant explanation for sensory exploitation, here we stress the role of mutational lag in allowing signallers to gain short-term advantages over receivers *sensu* (Dawkins & Krebs 1978).

Whereas some authors have emphasized the possibilities for deceit and exploitation in animal communication, others have argued that such manipulation is likely to be short-lived in evolutionary terms. Zahavi's handicap principle (1975, 1977*a*, 1987), which is the basis for many of these arguments, suggests that receivers can avoid being fooled by paying attention to signals that are costly and hence 'unfakeable'. Consider a signal that serves to advertise some quality of the signaller: if the cost of advertisement is greater for low-quality individuals, then, at equilibrium, the signal will be expressed in a quality-dependent manner, providing the receiver with useful information (Grafen 1990; Johnstone & Grafen 1992; Pagel 1993; Hutchinson *et al.* 1993; Vega-Redondo & Hasson 1994; Johnstone 1994).

Formal game-theoretical models of signal evolution have shown that Zahavi was right: honesty is stable, while exploitation is not. However, in looking for evolutionarily stable signalling and interpretation strategies, game theoretical models consider only the ultimate outcome of the signalling arms race, and ignore the interactions that occur as signaller and receiver evolve towards equilibrium (Arak & Enquist 1993). Furthermore, they assume that receivers are not constrained by any computational limitations. These models cannot, therefore, determine how quickly and easily honest equilibria can be reached; neither can they predict how often and for how long exploitative strategies are likely to arise and endure in the evolution of real signalling systems. Analysis of exploitation and the signalling arms race requires a different kind of model, which considers the dynamics of coevolution between signallers and receivers of finite capabilities.

In this paper, we present a simulation of signal evolution (applicable to any instance of quality advertisement) based on a genetic algorithm for evolving artificial neural networks (see: Enquist & Arak 1993; Arak & Enquist 1993). Signallers are assumed to vary in some quality that is of interest to recipients, but that recipients cannot directly observe; signallers advertise their individual qualities by employing signals of varying form and varying cost. Recipients, on the other hand, attempt to infer the true quality of signallers from the form of their signals. The distribution of signaller quality does not change over time; rather it is the strategies employed by signallers and receivers that coevolve. Within the world of the simulation, we are able to follow the evolutionary changes in signalling and interpretation strategies over many generations, and observe the interactions that occur as the system approaches equilibrium. In this way, the model allows us to investigate whether, and how often, exploitation can occur as a result of receivers lagging behind in the signalling arms race.

To assess the importance of signal cost for the evolution of honesty, we compare the dynamics of signaller–receiver coevolution when communication is cost-free and when it is costly. In the latter case, both low-cost and high-cost signals are available, but the cost of any particular signal is always greater for a lower-quality signaller (i.e. cost is 'quality-dependent'). To investigate the evolution of complex, multi-component displays (see Møller & Pomiankowski 1993), we manipulate the number of dimensions in which signals can vary, allowing more or less scope for transfer of information.

2. A NETWORK-BASED MODEL OF SIGNAL EVOLUTION

We consider two populations of neural networks, both of size p. One population comprises signaller networks, each of which implements a signalling strategy that maps qualities onto signals. The other comprises receiver networks, each of which implements an interpretation strategy that maps signals onto inferred qualities. Every 'generation' the performance of each network is evaluated as it interacts with every network in the opposing population. Receiver networks that tend to correctly infer signaller quality from the signals produced by networks in the opposing population are assigned a high fitness relative to those that make incorrect deductions. At the same time, signaller networks that tend to be attributed a high quality by networks in the opposing population are assigned a high fitness relative to those that are inferred to be of low quality. The next generation's populations are generated by differential replication of existing networks, favouring those that are of high fitness. Selection is thus assumed to favour accurate information acquisition by receivers, and successful manipulation by signallers. The networks are not, however, replicated exactly. There is a small probability during replication of mutation leading to a change in network structure. In this way, new strategies are constantly emerging and continued evolutionary change is possible.

(*a*) *Strategies and fitness functions*

Individual signalling networks evolve to implement rules for mapping one quality q from a set of positive, real-valued qualities $Q = (q_1, q_2, q_3, q_4)$ onto a signal s from the set $\mathscr{R}^n_{+}$, where n specifies the number of dimensions in which the signal can vary. Individual receiver networks implement rules for transforming a perceived signal s from the set $\mathscr{R}^n_{+}$ into a positive, real-valued quality assessment q_a from the set Q.

Each population seeks to maximize a fitness function. For the receiver population the relevant fitness function is

$$W_{\text{receiver}} = 1/[1+(q-q_a)^2]. \tag{1}$$

Receivers thus seek to maximize the accuracy with which they assess signaller quality, since their fitness is inversely related to the discrepancy between q and q_a. For the signaller population the relevant function is

$$W_{\text{signaller}} = q_a - C(s,q),$$

where $C(s,q)$ specifies the cost of employing a signal s for a signaller of quality q. In other words, the fitness of the signaller depends on the assessment of its quality by receivers and on the cost of its signal.

We consider two different kinds of signal cost. In the first

$$C(s,q) = 0,$$

i.e. signalling is cost-free. In the second

$$C(s,q) = |s|/q,$$

so that signals vary in cost, but a given signal is always more costly for signallers of lower quality.

(*b*) *Structure and function of the networks*

A signaller network comprises an input layer of four cells, an intervening layer of seven hidden cells and an output layer of n cells. Each cell in one layer connects to all cells in the subsequent layer, and with each connection a weight is associated that regulates the strength of the signal passing between the cells. These weights vary from one network to another in the signalling population and undergo evolutionary change during the course of the simulation.

The signaller network implements a strategy for converting qualities into signals. This conversion occurs as follows. First, the cells in the input layer register one of the four possible qualities (drawn from the set Q), which are encoded as four linearly independent vectors of equal magnitude. Each cell receives as input one component of the quality vector and produces an output equal to this value. Next, each cell in the hidden layer receives as input a weighted sum of the output from all cells in the initial layer and produces an output that is a sigmoid function (tanh) of its total input. Finally, in exactly the same way, each cell in the final layer receives as input a weighted sum of the output from all cells in the preceding (hidden) layer and produces an output that is a sigmoid function of its total input. The combined output of the final layer of cells specifies the signal, with the output of each individual cell providing one of the n signal components. The network as a whole thus receives an encoded quality as input, and produces a signal as output.

A receiver network comprises an input layer of n cells, an intervening layer of seven hidden cells, and an output layer of four cells. This network implements a strategy for converting signals into assessments of quality. It functions in a similar way to the signaller networks, except that it receives a signal as input and produces a quality assessment as output. The way in which the network produces a quality assessment is as follows. The combined output of the final layer of cells yields a vector of four components, with the output of each individual cell providing one component. Having encoded the four possible qualities that make up the set Q as four linearly independent vectors of equal magnitude (see above), we proceed to calculate the magnitude of the projection of the final output vector onto each quality vector. The quality for which this value is greatest is the one chosen as the assessment.

(*c*) *Evolving the networks*

The precise strategy implemented by a signaller or a receiver network is determined by that network's configuration of connection weights (see above). At the start of the simulation, these weights are assigned random values (lying between -1 and $+1$). Initially, therefore, the populations each comprise p different nets, each of which implements a different random strategy. The following sequence of steps is then repeatedly cycled through to simulate signal evolution.

(i) *Calculation of fitness values*

Each signaller network, having been assigned a random quality q, displays its signal s to every receiver, and in each case is attributed some quality q_a. The fitness consequences of these interactions for the signallers and receivers are calculated according to expressions (1) and (2). We then assign a single overall fitness to each individual in each population equal to the sum of the fitness values derived from the interactions in which it took part.

(ii) *Differential replication*

New populations are generated by differential replication of existing networks. Each member of a new population is a replica of one individual network from the old population, the probability that a given network will be chosen being given by

$$w_{\text{network}}/p\bar{w},$$

where w_{network} is the fitness of that particular network (as calulated in step (i)) and $\bar{w}$ is the mean fitness of all networks in the old population. Networks that enjoy a high fitness are thus likely to be chosen more than once, so that they are represented at a higher frequency in the next generation, while networks with low fitness may not be drawn at all.

(iii) *Mutation*

Members of the new population undergo mutation and recombination. Mutation influences each weight with a probability of 0.001 and changes its value so that

$$w_{\text{new}} = w_{\text{old}} + (x - 0.5)/c,$$

where c is a constant ($c = 0.7$ for the results given here) that determines the degree of correlation between the new value, w_{new}, and the old value, w_{old}, and x is a uniform random variable in the interval $[0, 1]$. Recombination occurs with a probability of 0.1 per network per generation and involves swapping a random subset of weight values with those of a randomly chosen partner.

This cycle is repeated until average receiver fitness remains within 5% of its preceding value for over 200 generations (we call this the equilibrium) or is terminated automatically at 1500 generations.

3. RESULTS

In figure 1*a*, *b*, we show how mean signaller and receiver fitness change over the course of many generations of simulation, for two different sets of assumptions regarding signal cost. Each graph represents the average result of ten runs, following an initial 200 generations of 'grace' in which receivers were allowed to evolve while the signalling strategy remained static (this enabled an effective interpretative strategy to become established in the population as a foundation for subsequent exploitation or honesty). The total simulation time was varied to determine the level at which fitnesses equilibrated. The results are plotted over the minimum time required to reach this equilibrium (1500 generations). The results shown in figure 1*a* were based on the assumption that signalling was cost-free, so that the fitness of a signalling network depended only on the qualities that were attributed to it. The results shown in figure 1*b* were based on the assumption that signalling was costly, i.e. the fitness of a signalling network depended both on the qualities that were attributed to it and on the form of the signals it produced. Whereas some forms of signal were more costly than others, this cost was assumed to be quality-dependent, i.e. any given signal was less costly for a high-quality than for a low-quality individual.

Where signalling was cost-free (see figure 1*a*, *b*),

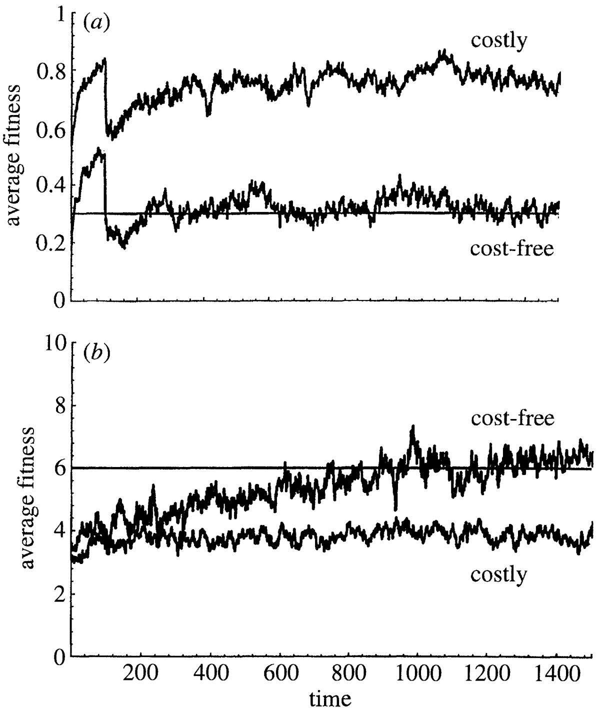

Figure 1. Mean signaller and receiver fitness plotted over 1500 generations (following 200 generations of 'grace' during which signalling strategies remained fixed) for two different sets of assumption about cost. Each plot represents the average ten runs, with a signal dimension of $n = 6$, where the set of possible qualities $Q = \{3, 5, 7, 9\}$. The probability of mutation in each case was 0.001 per connection weight per generation, while the probability of recombination was 0.01 per network per generation (see text for details). Graph (*a*) shows receiver fitness, while graph (*b*) shows signaller fitness.

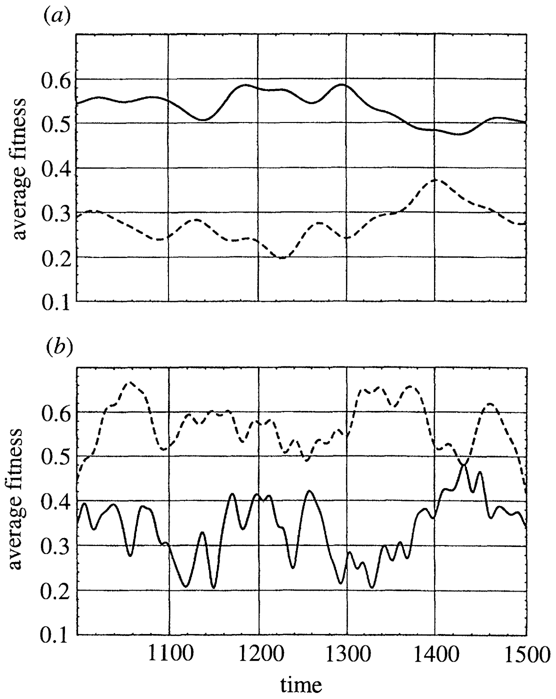

Figure 2. Filtered mean signaller and receiver fitness plotted over 500 generations of equilibrium for (*a*) a sample run with cost-free signalling and (*b*) a sample run with costly signalling. Fourier transforms have been used to remove high frequency oscillations in the data resulting from mutation and recombination. In both cases, the signal dimension was $n = 6$, and the set of possible qualities $Q = \{3, 5, 7, 9\}$. Solid line, signaller fitness; dashed line, receiver fitness.

signaller and receiver fitness rapidly came to fluctuate about the 'non-signalling level'. This represents the expected equilibrium fitness when receivers have to estimate signaller quality without the help of a signal (under these circumstances, with the parameter values listed in the figure legend, receivers maximize their fitness by attributing the mean quality to all signallers). When signalling is cost-free, therefore, it provides receivers with no information that will allow them to improve the accuracy of their estimate of signaller quality.

Where signalling was costly, by contrast, receiver fitness tended to increase to a point considerably higher than the predicted non-signalling level, while signaller fitness remained lower than the predicted non-signalling level (although showing a gradual increase over the course of the simulation). When signalling is costly, therefore, it provides receivers with information that allows them to estimate signaller quality more accurately than they otherwise could; at the same time, the cost involved lowers mean signaller fitness. Signalling is stable despite the fact that mean signaller fitness would be higher at a non-signalling equilibrium, because once receivers have begun to adjust their estimates of signaller quality in relation to the signals they perceive, it no longer pays to refrain from signalling.

Having examined the effects of signal cost on mean signaller and receiver fitness, we now consider in more detail the way in which these values vary over time. In figure 2*a*, *b* we show how mean signaller and receiver fitness fluctuate over 500 generations of evolution at the equilibrium for a single sample run of the cost-free (*a*) and (*b*) cases, respectively. The data have been Fourier transformed to remove high-frequency oscillations, which result from the random assignment of qualities in each generation and from the immediate effects of mutation and recombination. The remaining variation represents the fitness changes that result from the evolution of signaller and receiver strategies. As the figure shows, the two fitness values tend to fluctuate in opposition to one another, so that an increase in signaller fitness is accompanied by a decrease in receiver fitness, and vice versa. This behaviour is explored quantitatively in figure 3, a plot of residuals from untransformed data for the combined results of six runs over the course of 200 generations at equilibrium, revealing that there is a significant negative relation between the deviations of signaller and receiver fitness from the equilibrium mean $r^2 = 0.76$, $P < 0.01$). The oscillations that occur thus represent an evolutionary arms race where fitness gains in one population are paid for by fitness losses in the other (but where these do not have to be zero-sum gains and losses). This is a coevolutionary arms race where the evolution of signaller and receiver strategies has a 'Red Queen' dynamic. As is apparent from figure 2, the amplitude of the fluctuations about mean equilibrium fitness values is lower in the cost-free case (*a*) than in the costly case (*b*).

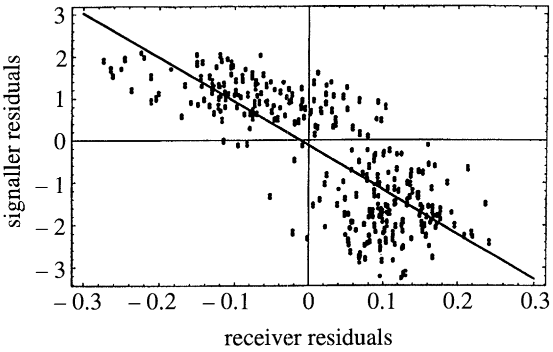

Figure 3. A plot of deviations from the equilibrium population fitnesses for signallers and receivers. The plot is generated through six runs of the simulation assuming signalling to be costly and where the signal dimension was $n = 6$. A regression is fitted to the residuals yielding $r^2 = 0.76$ with $P < 0.01$. Hence 76% of the variation in fitness occurs antagonistically between the signaller and receiver populations.

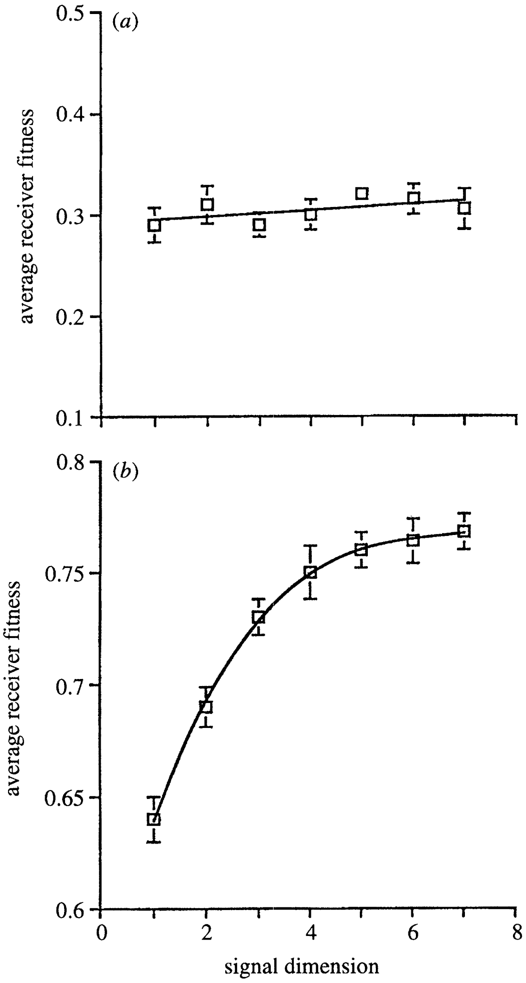

Figure 4. The influence of signal dimension on mean receiver fitness at equilibrium, for the (*a*) cost-free and (*b*) costly cases. Each plot represents the average of ten runs, where the set of possible qualities $Q = \{3, 5, 7, 9\}$.

Finally, figure 4 shows the relation between the number of signal dimensions available and mean receiver fitness at equilibrium. When signalling is cost-free (*a*), there is no relation; receiver fitness is uninfluenced by the number of dimensions. When signalling is costly (*b*), however, increasing the number of dimensions leads to an increase in mean receiver fitness. In the latter, therefore, an increase in the number of channels of communication that signallers can employ leads to an increase in the accuracy with which receivers are able to infer signaller quality.

4. DISCUSSION

Proponents of the view that signals are better understood in terms of selection for effective manipulation than for effective information transfer (Dawkins & Krebs 1978) see communication as an arms race between signallers and receivers. If receivers fall behind in the race, their outdated response strategies become vulnerable to exploitation. From this perspective, deceit appears to be a probable outcome of signal evolution. Zahavi's handicap principle, however, suggests that receivers can cut short the arms race and ensure that they obtain reliable information about signaller quality (or need) by restricting their attention to costly signals that cannot easily be faked (Zahavi 1975; Grafen 1990). Existing game-theoretical models of signalling have established honest, costly signalling as a possibility but have neglected several of the factors that make exploitation an attractive alternative. In particular, they consider only the stable outcome of the signalling arms race, and assume that receivers are not constrained by any computational limitations.

In this paper, we have investigated the same kind of signalling system as a number of previous theoretical studies (Grafen 1990; Johnstone & Grafen 1992; Pagel 1993; Grafen & Johnstone 1993), but we have used a different technique. Rather than employing the methods of evolutionary game theory to calculate a stable signalling equilibrium, we have used a genetic algorithm to simulate the course of signaller–receiver coevolution. Signalling and interpretation strategies are implemented by simple artificial neural networks, the structure of which undergoes evolutionary change during the course of the simulation. How do the results of the simulation compare with the predictions of evolutionary game theory?

In one respect, the results of the simulation closely match game-theoretical predictions: receivers are only able to obtain honest information regarding signaller quality when signalling is costly. When signalling is cost-free, receiver inferences are no more accurate than they would be in the absence of any communication. Models of honest advertisement and the handicap principle have shown that, when there is a conflict of interest between signaller and receiver, quality-dependent signal cost is necessary for the maintenance of honesty, because it prevents low-quality signallers from profitably faking the signals of superior individuals (but see Maynard Smith 1994). The results of the simulation thus provide further support for the handicap principle.

The introduction of quality-dependent signal cost into the simulation does not, however, completely eliminate dishonesty. The costly signalling equilibrium is characterized by successive waves of exploitation, during each of which a manipulative signalling strategy gains temporary success. Ultimately, such strategies lose their selective advantage and are eliminated as receivers evolve a modified interpretation strategy that yields an appropriate response to them. Nevertheless, these episodes are a recurrent feature of the equilibrium (despite the fact that any one exploitative strategy is likely to be short-lived) because new opportunities for manipulation are constantly arising. Over time, therefore, signals continuously change in form, remaining partially honest on average but never perfectly honest (see Andersson 1980; Grafen 1990).

The potential for exploitation arises because an adaptive interpretation strategy can be implemented by many different equally efficient mechanisms. These mechanisms are likely to differ in their response to signal forms that lie outside the currently occurring range, but selection is blind to responses that are provoked by stimuli that receivers never encounter. Consequently, the interpretation mechanism used is subject to change by random drift and the display form that is optimal in stimulating a response will also change over time, giving rise to hidden preferences that are open to exploitation (Arak & Enquist 1993). Furthermore, evolutionary modification of an interpretation strategy in response to one episode of exploitation is likely to alter the response of receiver networks to other stimuli, generating further hidden preferences that can be exploited. Dawkins & Krebs (1984) argued that exploitation would necessitate increasingly exaggerated displays, as receivers evolved increased 'sales resistance'. Our results, however, indicate that increased resistance to one form of exploitative signal only sets the stage for manipulation by a different form of display (Andersson 1980).

The potential for exploitation, however, only arises if receivers show some response to the signals employed. When signalling is cost-free, selection cannot maintain even partial honesty, so that receivers rapidly evolve to ignore any variation in signal form. Under these circumstances, there is little opportunity for exploitation (as shown by the low amplitude of the fluctuations about mean equilibrium fitness in the cost-free signalling simulation). The simulation thus reveals that, while signal cost is necessary for the maintenance of (partial) honesty, it also facilitates the evolution of exploitation.

Sensory bias which has not been included in our model – where pleitropic effects influence the form of the preferred signal – is also likely to give rise to a degree of exploitation (Ryan 1990). This form of exploitation is distinct from mutational lag, which we demonstrate, and may enforce an enduring constraint on signal evolution if the alternative function of a given preference persists through time.

Finally, the relation between mean receiver fitness at equilibrium and the number of signal dimensions (for costly signalling) indicates that more complex multicomponent displays are able to provide receivers with additional information about signaller quality. A possible explanation is that, given the greater range of signal forms available when the number of dimensions is large, signallers of different quality can evolve more divergent displays, which allow receivers to more easily distinguish between them. Contrariwise, increase in the number of signalling dimensions might also reduce the odds of a cheating signaller stumbling across a receiver's preferred signal; higher signal dimensions could therefore be viewed as more robust indicators of quality simply because they are harder to find.

We thank Mark Pagel, Michael Ryan and two anonymous referees for their suggestions and comments on the typescript. D. C. K. is supported by NERC. R. A. J. is a Junior Research Fellow at St Catherines College, Cambridge.

REFERENCES

Andersson, M. 1980 Why are there so many threat displays? *J. theor. Biol.* **86**, 773–781.

Arak, A. & Enquist, M. 1993 Hidden preferences and the evolution of signals. *Phil. Trans. R. Soc. Lond.* B **340**, 207–213.

Dawkins, R. & Krebs, J.R. 1978 Animal signals: information or manipulation. In *Behavioural ecology: an evolutionary approach* (ed. J. R. Krebs & N. B. Davies), pp. 282–309. Oxford: Blackwell Scientific.

Endler, J.A. 1992 Signals, signal conditions, and the direction of evolution. *Am. Nat.* **139**, S125–S153.

Endler, J.A. 1993 Some general comments on the evolution and design of animal communication systems. *Phil. Trans. R. Soc. Lond.* B **340**, 215–225.

Enquist, M. & Arak, A. 1993 Selection of exaggerated male traits by female aesthetic senses. *Nature, Lond.* **361**, 446–448.

Grafen, A. 1990 Biological signals as handicaps. *J. theor. Biol.* **144**, 517–546.

Grafen, A. & Johnstone, R.A. 1993 Why we need ESS signalling theory. *Phil. Trans. R. Soc. Lond.* B **340**, 245–250.

Hutchinson, J.M.C., Macnamara, J.M. & Cuthill, I.C. 1993 Song, sexual selection, starvation and strategic handicaps. *Anim. Behav.* **45**, 1153–1177.

Johnstone, R.A. 1994 Honest signalling, perceptual error and the evolution of 'all-or-nothing' displays. *Proc. R. Soc. Lond.* B **256**, 169–175.

Johnstone, R.A. & Grafen, A. 1992 Error-prone signalling. *Proc. R. Soc. Lond.* B **248**, 229–233.

Kirkpatrick, M. & Ryan, M.J. 1991 The paradox of the lek and the evolution of mating preferences. *Nature, Lond.* **350**, 33–38.

Krebs, J.R. & Dawkins, R. 1984 Animal signals: mind-reading and manipulation. In *Behavioural ecology: an evolutionary approach* (2nd edn ed. J. R. Krebs & N. B. Davies), pp. 380–402. Oxford: Blackwell Scientific.

Maynard Smith, J. Must reliable signals always be costly? *Anim. Behav.* **47**, 1115–1120.

Møller, A.P. & Pomiankowski, A. 1993 Why have birds got multiple sexual ornaments? *Behav. Ecol. Sociobiol.* **32**, 167–176.

Pagel, M. 1993 Honest signalling among gametes. *Nature, Lond.* **363**, 539–541.

Proctor, H.C. 1991 Courtship in the water mite *Neumania papillator*: males capitalize on female adaptations for predation. *Anim. Behav.* **42**, 589–598.

Proctor, H.C. 1993 Sensory exploitation and the evolution of male mating behaviour: a cladistic test using water mites (Acari: *Parasitengona*). *Anim. Behav.* **44**, 745–752.

Ryan, M.J. 1990 Sensory systems, sexual selection, and sensory exploitation. *Oxford Surv. Evol. Biol.* **7**, 157–195.

Ryan, M.J. & Rand, A.S. 1990 The sensory basis of sexual selection for complex calls in the túngara frog, *Physalaemus pustulosus* (sexual selection for sensory exploitation). *Evolution* **44**, 305–314.

Ryan, M.J. & Rand, A.S. 1993 Sexual selection and signal evolution: the ghost of biases past. *Phil. Trans. R. Soc. Lond.* B **340**, 187–195.

Vega-Redondo, F. & Hasson, O. 1993 A game-theoretic model of predator–prey signalling. *J. theor. Biol.* **162**, 309–319.

Zahavi, A. 1975 Mate selection: selection for a handicap. *J. theor. Biol.* **53**, 205–214.

Zahavi, A. 1977 The cost of honesty (further remarks on the handicap principle). *J. theor. Biol.* **67**, 603–605.

Zahavi, A. 1987 The theory of signal selection and some of its implications. In *International symposium of biological evolution* (ed. V. P. Delfino), pp. 305–327. Bari: Adriatica Editrice.

Received 17 November 1994; accepted 13 January 1995