

# Enhancing game theory with coevolutionary simulation models of honest signalling

Dave Harris & Seth Bullock

Informatics Research Institute, University of Leeds, Leeds LS2 9JT, UK.

Tel: +44 (0)113 233 5322, Fax +44 (0) 113 233 5468, Email: dharris@comp.leeds.ac.uk, seth@comp.leeds.ac.uk

**Abstract - Game-theoretic models provide a rigorous mathematical modelling framework, but tractability considerations keep them simple. In contrast, Evolutionary Simulation Models (ESMs) may be complex, but can lack rigour. We demonstrate that careful synthesis of the two techniques provides improved insights into the processes underlying the evolution of cooperative signalling systems.**

## I. INTRODUCTION

Evolutionary systems are of interest to both biologists and computer scientists. For the former, they represent challenging objects of study. For computer scientists, they are both a source of inspiration and increasingly a domain within which computational techniques are being profitably applied. Traditionally, biologists have tended to use mathematical methods to model biological systems. In the field of behavioural ecology, for example, game theory has been the preferred tool since it was first championed by Maynard Smith and Price in the early 1970s [1].

Computer scientists have used simulation models to explore similar systems with some success [2]. However, their uptake within the biological community has been limited by the perceived lack of rigour associated with simulation in comparison to equational modelling techniques [3]. Formal models are expected to explicitly present all of their assumptions and to provide an intelligible account of the moves leading to the model's conclusions. Simulations, on the other hand, are often *opaque*, even to their designers [4]. Unlike the interpretation of mathematical models, understanding why a particular simulation model produces the results that it does is often a significant undertaking. As a result, even when the actual details of the implementation are made clear (which they are often not) by providing the source code, for example, simulation models are difficult to assess and appreciate. This sometimes leads to artefactual claims [5]-[6].

Despite the success of game theory modelling, it suffers from the same tractability limitations as all mathematical modelling techniques. While simple games can be analysed fairly easily, including more detail often renders models effectively insoluble. This has not been a severe problem in behavioural ecology partly because theoreticians have been interested in quite simple games with few equilibria. These simple models have thrown light on many important questions, but further exploration often requires addressing more complicated scenarios.

We believe that this is best accomplished by using game-theoretical models as the basis for evolutionary computer simulation modelling. Wedding the two approaches may serve to overcome the limitations of each—the perceived lack

of rigour in the simulation design, and the tractability constraints on mathematical modelling.

One might consider a game-theoretic model to represent the *pure selective force* driving an evolutionary system. This is because game theory assumes evolution takes place in a world where even the smallest of selective forces will eventually overcome any limitations [7]-[8]. The only constraints on evolutionary change imposed by a game-theoretic model are that strategies must be drawn from the pre-defined strategy set. There are no genetic constraints, no developmental limitations, no noise, no fitness landscapes to traverse, etc., there are only evolutionarily stable strategies (An ESS is defined as a strategy, or mixtures of strategies, that, when prevalent, cannot be invaded by any others).

In contrast, an evolutionary simulation model immediately introduces evolutionary constraints, many of which might be described as “logistic” factors. These factors would include stochastic effects such as noisy fitness functions or sampling errors, and genetic constraints, which govern how far one strategy is from another across a fitness landscape. In addition, population structure, direct vs. indirect costs, life history strategies and other influences are also often brought to bear upon the evolutionary process being modelled. It is, of course, possible to include these kinds of constraints in a game theory model, but only at the cost of greatly complicating the mathematics required to obtain a solution.

Importantly, a simulation modelling approach also enables a researcher to study any *non-ESS*, transitory phenomena. This is particularly important with systems that have non-stable or multiple equilibria, as well as potentially revealing system behaviours that might give an indication of an evolutionary trajectory toward an ESS. Just because a behaviour is not an ESS does not necessarily mean that it will rarely be observed. Commonly occurring transitory phenomena also require explanation.

In the remainder of the paper we will apply both game theory and evolutionary simulation modelling techniques to an important problem in biology. Our primary aim is to demonstrate that the synthesis of both techniques provides a better understanding of the problem than game theory alone, but we also aim to contribute to the biological literature on the evolution of honest signalling systems.

## II. THE EXAMPLE PROBLEM

One problem that has been studied extensively through both game-theoretic [9]-[10] and evolutionary simulation modelling [11] approaches is the discrete action-response game, which models communication between individuals who may or may not have a conflict of interest. In the majority of cases, only one modelling technique was used.

We will take Hurd's game theoretic model of the evolution of honest handicap signalling and by implementing it as an evolutionary simulation, demonstrate that the game theory does not tell the whole story of the system's behaviour. To give an idea of the basis of the model, we present the important equations of the game theory below.

The popularity of the problem is due to the difficulty that modellers have had with the concept of the *handicap principle*, as proposed by Zahavi [12]. He suggested that for signallers to be honest where conflicts of interest exist, signals must be costly, i.e., a handicap. Thus, for example, a high quality peacock can "afford" to bear the cost of a larger tail than one of lower quality. Despite widespread interest in Zahavi's idea, it took around 15 years from its first publication for a successful model to prove its validity. Grafen's game-theoretical proof demonstrated this and also required the condition that signals must be more costly for low-quality individuals to make, where quality is of interest to the signal receiver [13]-[14]. For example, this could be the genetic quality of a displaying male for a female choosing a mate.

A competing view of signal system evolution, which has fallen out of favour since the rise of the handicap principle, proposes that signallers and receivers may settle upon cheap conventions with which to convey information to their mutual advantage, e.g., red insect colouration might indicate unpalatability to birds [15]. There is nothing to prevent these conventions from being abused by free-loaders, mimics and cheats, e.g., red insects which nevertheless are tasty food for birds. As the prevalence of these abusers grows, the value of the signal deteriorates until the signalling system collapses under the weight of its own infidelity. Partly as a result of the inherent instability featured in this account, it has been neglected in favour of the handicap principle's stable equilibria.

A signaller,  $S$ , is given an internal state,  $z$ , which only they know but which is of interest to a receiver,  $R$ . This state will either be  $H$ , representing high quality, or  $L$ , representing low quality. On the basis of this state, the signaller makes a signal,  $s$ , to which the observer makes a response,  $r$ . The signal will either be East ( $E$ ) or West ( $W$ ) and the response will be either Up ( $U$ ) or Down ( $D$ ).

An example of this system might be male blackbirds signalling to females in the hope of eliciting a mating opportunity. Imagine a female blackbird flying about and suddenly observing a male blackbird sitting in a tree. He also observes her and can do one of two things. He could remain inactive (East) and hope she flies over or he can start singing in the hope that she will be attracted by it (West). The female having observed the male's action could either ignore him (Down) or fly over and mate with him (Up).

Note that there is no reason why, for example, the cost of signalling East should not be higher than signalling West or vice versa, or that High state males should find signalling East cheaper than Low state males or *vice versa*.

The game can be played in many ways, but Hurd concentrated upon the version in which the signallers always want to get an Up response from the observer, but the

observer only wishes to give that response to a signaller with a High internal state. This generates a conflict of interest that makes the solution non-trivial.

The participants in the game are given a fitness score determined by how successful they have been. In game theoretical terms these would be:

$$w_S = v(r) - c(z, s) \quad (1)$$

$$w_R = v(z, r) \quad (2)$$

In equation 1  $w_S$  is the fitness score of the signaller,  $v$  is the value of the receiver's response, and  $c$  is the cost of the signaller's action given his internal state. In equation 2  $w_R$  is the fitness score of the receiver, and  $v$  is the value of the receiver's response given the signaller's internal state.

The relative cost of signalling can be defined as:

$$C_H = c(H, W) - c(H, E) \quad (3)$$

$$C_L = c(L, W) - c(L, E) \quad (4)$$

This defines  $C_H$  as the difference in cost for a High quality signaller between signalling West and East.  $C_L$  is the equivalent for the Low quality signaller. Honest signalling is defined arbitrarily as High state signallers making the West signal while Low state signallers make the East signal. Hurd went on to show that honest signalling is only an ESS when:

$$C_L > V > C_H \quad (5)$$

Where  $V$  is the value of the receiver's Up response to the signaller. Other models have explored cases in which the value of the Up response is not the same for High and Low state signallers, but in our model  $V$  is insensitive to signaller state. This is represented graphically in Figure 1.

The importance of equation 5 is that it contradicts Zahavi's claim that signals must be costly in order to remain honest. According to (5), signalling could remain stable despite  $C_H$  being negative. This indicates, that as long as Low state signallers cannot afford to make an advertisement, it can remain honest, even if High state signallers pay nothing to make it, or are even better off making the advertisement than not.

### III. OUR MODEL

In order to explore more fully the evolutionary dynamics of the action-response game, we constructed an evolutionary simulation that faithfully captures the structure of Hurd's model. Due to the exclusivity of the signaller and receiver behaviours, a two population model is appropriate. A population of 100 signallers was coevolved against a population of 100 receivers over 1000 generations. At birth, each signaller was either determined to be High or Low ( $p=0.5$ ). Each signaller's strategy was defined by a two-bit genotype, where the first bit represented which action to




Fig 1. Graph showing where honest signalling is an ESS in Hurd's model

make when Low, and the second which action to make when High. The receiver's two-bit genotype represented which response to give when observing East and which to give when observing West. At the outset of coevolution, all genes were initially set randomly to either 1 or 0.

The available signaller and receiver strategies are named in relation to Hurd's arbitrary honest signalling scenario in which signallers signal West when High and East when Low, while receivers respond Up to West and Down to East. A Cynic never signals West, while a Bluffer always does. Honest signallers signal East when in the Low state and West in the High state, while dishonest signallers signal West in the Low state and East in the High state. A Generous receiver always responds Up, while a Mean receiver never does. A Believer responds Up to West and Down to East, while a Non-Believer behaves in the opposite manner.

Each generation, every signaller was paired at random with a receiver. After playing the action-response game once, they were awarded scores according to the fitness functions defined above. Reproduction was determined by tournament selection with tournament size 2. During reproduction, each gene had a 3% chance of mutation.

For each coevolutionary run, several parameters defined the cost and value of signalling and responding (see Figure 1). Receivers were given a score of 100 if they gave an Up response to a High state signaller or Down to a Low state signaller, otherwise they received a zero score. The value of an Up response for a signaller was set to 100, while a Down response had zero value. The cost of both East and West behaviours for both High and Low state signallers was varied across runs from -100 to +100 in steps of 25. Each of the 6561 combinations of cost parameters was replicated 20 times. We have described the decisions involved in choosing these experimental parameters and the consequent effect on results in another paper [16].

For each run, the number of signallers in the final generation utilising each of the four possible strategies was counted. If the signaller population was behaving randomly, we would expect them to be distributed equally across the

potential strategies with a binomial distribution,  $p=0.25$ ,  $n=100$ . This gives a 95% confidence interval of 10-37. If the number of signallers using a particular strategy lies outside of this interval, we can deduce that the distribution is non-random. For instance, if we observe that more than 37 Cynics in the final generation then we can be confident that there was some selection for this particular strategy. Conversely, if we observe less than 10 Cynics, we can infer that there has been selection *against* this course of action. For each of the 20 runs carried out for each combination of cost parameters, the signaller and receiver populations were scored for each possible strategy (positive selection, negative selection, and no selection).

It is common in signalling theory papers to concentrate on the signaller and to pay little or no attention to receiver strategy. This may be due to the perception that it is the signaller's behaviour that is of interest; it is the phenomenon of honest *signalling* that is under investigation. However, it should be remembered that this is a coevolutionary problem and the receiver's strategy is of equal importance in determining the behaviour of the signalling system. In recognition of this, the receiver strategy will be discussed before the signallers.

#### IV. RESULTS

##### A. Receiver strategy

We found that the mean and generous strategies are only ever used at a low level and there is no region of the parameter space within which they approach an ESS. This is a little surprising given the presence of areas where signallers are predominantly using the Cynic or Bluffer strategy. One might imagine that the Mean or Generous receiver strategies might do as well as any other in these areas.

The believer and non-believer distributions are shown in Figure 2. These distributions are rotationally symmetrical. One item of note is that the area of the parameter space that results in significant numbers of Believer strategists not only coincides with the area that game theory predicts should favour Honest signallers, as one would expect, but is actually more widespread. Likewise, Disbelief is not limited to the area of the parameter space predicted by game theory. Why are receiver strategies sometimes favoured in the absence of the signalling strategy that justifies their existence?

Confronted with, for example, a uniform population of cynics, receivers cannot glean any information concerning signaller state from signaller behaviour. Regardless of whether signallers are High or Low, they always move East. In such a situation each of the four receiver strategies are, on average, equally successful, scoring 100 roughly half the time and zero for the remainder. Given such a scenario, in the perfect world of game theory (where idealizations such as infinite population size and fitness assessment in the limit are assumed to hold) each receiver strategy should be represented equally in the population.

However, in a stochastic simulation with finite population size in which individuals are assessed over a small number of trials, the likelihood that the behaviour of a population of




Figure 2. Graphs showing the distribution of receiver strategies with respect to  $C_H$  (X-axis) and  $C_L$  (Y-axis). Increasingly heavy shading represents an increasing number of runs that display positive selection pressure for a strategy. Left to right: Believer, Non-believer

signallers will convey no information about internal state is very low, since in each generation, mutation will generate a small number of non-ESS signallers. Given that Believers and Non-believers are the only strategists capable of exploiting any correlation between signaller internal state and behaviour, it is perhaps not surprising to see them outperform Mean and Generous strategists.

But why should believers predominate above the line  $C_H=C_L$  with Non-believers below? It must be the case that the cost parameters above the line  $C_H=C_L$  are, to some perhaps quite small extent, biased in favour of Honesty (even outside the area of the parameter space admitting of an Honest signalling ESS). Likewise, below the line  $C_H=C_L$ , Dishonesty is favoured to some extent. Even a small number of either signaller strategy would be enough to tip the receiver population in favour of either Belief or Non-Belief.

### B. Signaller strategy

Figure 3 clearly shows that honesty is observed under the ESS conditions predicted by the game-theoretic model. We can also appreciate which conditions favour the other three strategies. Note that considering the four graphs together, these results are also rotationally symmetrical, with honesty and dishonesty being ESS's on opposite sides of the chart. This demonstrates that given appropriate signalling costs, East and West are effectively interchangeable as one would expect. Near the centre of the parameter space, there exist regions of that are not dominated by any particular strategy. The game-theoretic model suggests that absolute signalling costs are irrelevant to the stability of honest signalling. Rather it is the difference between the cost of East and West that is critical to stability. Each point on the graphs in Figure 3 pools several runs with different cost parameters. Are there significant differences *within* these pooled results? Figure 4 shows that there are not. Varying the cost of signalling East when in the Low state has no effect on the number of signallers using the honest strategy. As predicted, it is the relative cost, ( $C_L$ ) which determines this. Absolute costs are unimportant. This also holds for variation in the other absolute cost parameters. In combination with the rotational symmetry displayed by the graphs in Figure 3, this result allows us to greatly reduce the number of different cost values that we need to explore.




Figure 3. Graphs showing the distribution of signaller strategies with respect to relative signalling costs  $C_H$  (X-axis) and  $C_L$  (Y-axis). Increasingly heavy shading represents an increasing number of runs that display positive selection pressure for a signalling strategy. Clockwise from top-left: Cynic, Honest, Bluffer, Dishonest.




Figure 4. A graph showing the effect of  $C_L$  on strategy frequency whilst varying the cost of signalling East.  $C_H$  was held constant at 0 for all runs.

Figure 4 also highlights an area of particular interest to us. The game theory suggests that we should see a sharp “phase transition” separating honest signalling from no honest signalling. Instead we see a gradual increase in the number of populations significantly using the honest strategy as we cross the predicted boundary in the parameter space. What is responsible for the low level of Honesty?

The distribution of strategies over the graphs is summarised in Figure 5, with regions lacking a dominant strategy labelled NS. A closer examination of these areas suggests that honest signalling occurs under these conditions, but is not stable over time due to the presence of Bluffers.

In order to gain a better idea of what was happening in the areas designated NS, we re-ran the experiment with a finer resolution. For this higher resolution experiment, the cost of signalling East was set to zero and the cost of signalling West varied between 0 and +100 in steps of 5 for each state. Figure

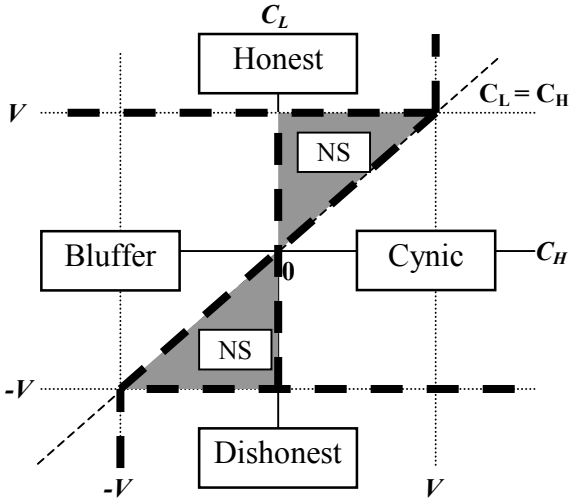


Figure 5. Graph showing the distribution of signaller strategies with respect to  $C_L$  and  $C_H$ .




Figure 6. Graphs showing a higher resolution of the distribution of signaller strategies with respect to  $C_H$  (X-axis) and  $C_L$  (Y-axis). Clockwise from top left: Cynic, Honest, Bluffer. The last graph shows the corresponding area on the original plots




Figure 7. Graph showing the number of individuals using the Honest, Bluffer and Believer strategies during one simulation run in the NS region.

6 shows the higher resolution graphs, which confirm the presence of a non-ESS signalling strategy. This appears to take the form of a mixture of honest and bluffing strategies. This is of particular interest as it has often been theorised that such mixed strategies occur. Despite the fact that they are not ESS strategies, these transitory phenomena still occur frequently in this model.

Figure 7 shows the number of individuals utilising a particular strategy over the course of a simulation run with the signal cost parameters  $C_L=75$ ,  $C_H=25$ , which corresponds to the middle of the area marked NS on Figure 3. One can see that the numbers of Bluffers, Believers and Honest strategists are correlated. The initial increase in the frequency of the Believer strategy is probably due to stochastic effects, since against the background of predominant Cynicism the Believer strategy has the same fitness as the other three response strategies. However, once the frequency of Belief rises slightly above chance, this favours both Honest and Bluffer strategists, since they get an Up response from Believers. The increase in frequency of Honest strategists gives Believers a slight edge over their competitors, boosting the numbers of both. However, Bluffers are able to piggyback on this trend. Eventually, the increasing numbers of Bluffers will devalue the Honest signal, leading to a crash as receivers abandon the believer strategy. This causes a corresponding reduction in the number of signallers using both the Honest and Bluffing strategies. The system returns once more to the non-signalling equilibrium at which signallers are Cynics and receivers are equally well off whatever strategy they pursue.

## V. DISCUSSION

The results clearly show that an area of the parameter space outside of that predicted by game theory contains some honest signalling (Figure 8). The fact that it appears to be evolutionarily unstable explains why. However, it should be pointed out that although it is unstable it does appear in between 50% and 75% of the populations within the area. If one were to rely purely on the game-theoretic results, then it would be predicted that no honest signalling could occur in that area of the parameter space—signals are too cheap for the Low quality signallers, they will invade and destabilise any nascent signalling system. We show that significant periods of honesty will occur in these scenarios, but that it will be mixed with signallers using the Bluffer strategy, and that each evolved signalling system will be relatively short-lived. It is important to realise that this is not just an arbitrary anomaly. As discussed earlier, within behavioural ecology, short-lived signalling systems of this type had been theorised to exist, but the notion of this type of cheap, but fragile signalling system has been displaced by theories of more stable handicap signalling systems.

The extended areas of the Believer and Non-Believer strategies also require some explanation. The most likely reason is that as long as a number of signallers are Honest

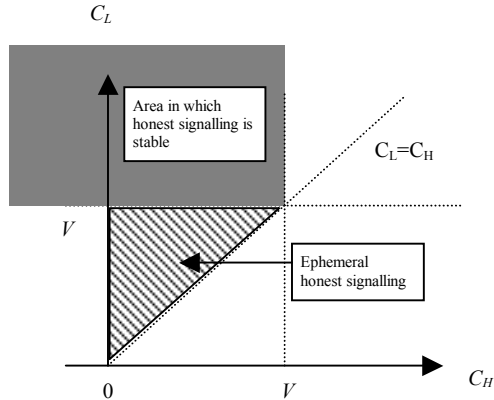


Figure 8. Graph showing where honest signalling occurs in our model, including an area not predicted by game theory.

or Dishonest, signaller behaviour will provide *some* information about signaller state. If a source of this information exists, despite its unreliability, it will allow receivers to make a better than random choice. The question of how reliable such a signal must be remains open.

## VI. CONCLUSION

We have shown the utility of evolutionary simulation models in combination with game theory in relation to discovering potentially important results. The existence of ephemeral honest signalling will be of particular interest to biologists. Current dogma states that signals cannot be considered reliable unless they are costly. Our model shows that not only can relatively cheap signals allow some level of honest signalling, but that receivers are willing to tolerate some level of “cheating” by non-honest signallers.

Biologists have become accustomed to only considering behaviours that have been shown to be an ESS. This is largely due to the success of game theory. However, now that the limits of the ability game theory to shed new light on the mechanisms involved in behavioural ecology may be in sight, it is time to consider additional theoretical methods.

One can consider the discoveries due to the use of game theory to be a framework, we now need to begin filling in the gaps between the supports. Merely knowing the eventual evolutionary outcome of selective forces provides an unbalanced picture, which may be misleading when applied to the real world. By using simulation techniques we can discover both the evolutionary trajectories that may be followed and also the nature of ephemeral phenomena. Such equilibria may be more common than currently assumed. How many seemingly stable signalling systems in the natural world are actually transitory in nature? The use of simulation modelling may help empiricists to find out by both opening their eyes to the possibility and giving them some idea of which conditions favour these phenomena.

Whilst we have concentrated upon the enhancement of game theory using simulation, there is no reason why this process should not work in reverse. This would entail a simulation being reverse engineered to produce a game-

theoretic model, which would show which parts of the simulation result are due to the selective forces acting on the phenomenon under investigation and which are due to the manner in which the simulation has been implemented.

## VII. Further Work

Having demonstrated that logistics are important in signal evolution, there is ample scope to expand upon this work. Potential areas include exploring multiple signalling bouts, different receiver pay-off matrices, the effect of having the value of receiver response vary with signaller state, allowing the signaller’s internal state to vary over a lifetime and an examination of a continuous version of the model.

## VIII. Acknowledgements

The authors would like to thank Jason Noble, John Cartledge, Tom Carden and three anonymous reviewers for their helpful comments during the writing of this paper.

## REFERENCES

- [1] J. Maynard Smith & G. R. Price, “The logic of animal conflict,” *Nature*, Vol. 246, pp. 15-18, 1973.
- [2] E. A. Di Paolo, “Ecological symmetry breaking can favour the evolution of altruism in an action-response game”, *J. Theor. Biol.*, Vol. 203, pp 135-152, 2000.
- [3] M. S. Dawkins & T. Guilford, “An exaggerated preference for simple neural network models of signal evolution?”, *Proc. Roy. Soc. Lon. Ser. B*, Vol. 261, pp 357-360, 1995.
- [4] E. A. Di Paolo, J. Noble & S. Bullock, “Simulation models as opaque thought experiments,” *Artificial Life VII: Proceedings of the Seventh International Conference on Artificial Life*, Portland, Oregon, USA, 1-6 August, pp. 477-486, 2000.
- [5] M. Enquist & A. Arak, “Symmetry, beauty and evolution,” *Nature*, Vol. 372, pp 169-172, 1994.
- [6] S. Bullock & D. Cliff, “The role of ‘hidden preferences’ in the artificial co-evolution of symmetrical signals,” *Proc. Roy. Soc. Lon. Ser. B*, Vol. 264, pp 505-511, 1995.
- [7] P. Hammerstein, “Darwinian adaptation, population genetics and the streetcar theory of evolution,” *J. Math. Biol.*, Vol 34, pp. 511-532, 1996.
- [8] J. Maynard Smith, “*Evolution and the Theory of Games*”, Cambridge University Press, 1982.
- [9] P. L. Hurd, “Communication in discrete action-response games”, *J. Theor. Biol.*, Vol. 174, pp 217-222, 1995.
- [10] S. Szamado, “The validity of the handicap principle in discrete action-response games”, *J. Theor. Biol.*, Vol. 198, pp 593-602, 1999.
- [11] S. Bullock, “An exploration of signalling behaviour by both analytic and simulation means for both discrete and continuous models,” *Fourth European Conference on Artificial Life (ECAL97)*, Brighton, UK, 28-31 July, pp. 454-463, 1997.
- [12] A. Zahavi, “Mate selection – a selection for a handicap,” *J. Theor. Biol.*, Vol. 53, pp 205-214, 1975.
- [13] A. Grafen, “Biological signals as handicaps,” *J. Theor. Biol.*, Vol. 144, pp 517-546.
- [14] A. Grafen, “Sexual selection unhandicapped by the Fisher process,” *J. Theor. Biol.*, Vol. 144, pp 473-516.
- [15] J. R. Krebs & R. Dawkins, “Animal signals: Mind-reading and Manipulation,” *Behavioural Ecology: An Evolutionary Approach*, 2<sup>nd</sup> Edition, Blackwell Scientific Publications, Oxford, pp 380-402, 1984.
- [16] D. Harris & S. Bullock, “Logistic effects influence evolutionarily stable strategies in the discrete action-response game,” In prep.