City Research Online

City, University of London Institutional Repository

Citation: Broom, M. & Ruxton, G. D. (2011). Some mistakes go unpunished: the evolution of "all or nothing" signalling.. Evolution, 65(10), pp. 2743-2749. doi: 10.1111/j.1558-5646.2011.01377.x

Title: Some mistakes go unpunished: the evolution of "all or nothing" signalling

Mark Broom[1] & Graeme D Ruxton[2]

[1]Centre for Mathematical Science
City University
Northampton Square
London EC1V 0HB
UK
mark.broom@city.ac.uk

[2]Institute of Biodiversity, Animal Health and Comparative Medicine
College of Medical, Veterinary and Life Sciences
Graham Kerr Building
University of Glasgow
Glasgow G12 8QQ
UK
Graeme.Ruxton@glasgow.ac.uk

Running title: evolution of “all or nothing” signalling

Abstract

Many models of honest signalling, based on Zahavi’s handicap principle, predict that if receivers are interested in a quality that shows continuous variation across the population of signallers, then the distribution of signal intensities will also be continuous. However, it has previously been noted that this prediction does not agree with empirical observation in many signalling systems, where signals are limited to a small number of levels despite continuous variation in the trait being signalled. Typically, there is a critical value of the trait, with all individuals with trait values on one side of the threshold using the same cheap signal, and all those with trait values on the other side of the threshold using the same expensive signal. It has already been demonstrated that these classical models naturally predict such “all-or-nothing signalling” if it is additionally assumed that receivers suffer from perceptual error in evaluating signal strength. We show that such all-or-nothing signalling is also predicted if receivers are limited to responding to the signals in one of two ways. We suggest that many ecological situations (such as the decision to attack the signaller or not, or mate with the signaller or not) involve such binary choices.

Keywords: signaling, signal honesty, Zahavi’s handicap principle, communication, cost of signalling

Introduction

Game theoretical models based on Zahavi's handicap principle have been very influential in offering an explanation for how signalling can remain (on average) honest when there is conflict of interest between signaller and receiver (Maynard Smith & Harper 2003; Searcy & Nowicki 2005) . Johnstone (1994) raised an interesting comparison between the predictions of still-influential models and empirical observation. Models generally predict that the intensity of the signal will vary continuously in relation to the quantity being signalled. For example, in a situation where potential prey individuals vary continuously in the strength of their chemical defences, these models would predict a similar continuous distribution of warning signal intensities to potential predators. To express this another way, these models predict that the signals should provide exact quantitative information about the specific defensive capability of each signaller. In contrast, Johnstone (1994) provides numerous empirical examples of signals where observed variation in signal strength is much less: being confined to a small number (often two) of discrete signal strengths. In the context of our example above, this would suggest that even if there is strong and continuously-distributed between-individual variation in the strength of the defences being signalled, the potential prey only adopt one of two signal intensities. All those individuals with defence levels below some threshold value produce essentially identical signals of the same low intensity; all those with defence values above the threshold signal at the same characteristic high intensity. In comparison to the model predictions then, real signals often seem less quantitatively informative. They inform the receiver not about the specific quality of an individual signaller but only about the range of qualities (either above or below the threshold in the example above) in which the individual falls.

Johnstone (1994) not only drew attention to this apparent tension between model predictions and empirical observations, he also offered a plausible solution. He demonstrated that previous models had assumed that the receiver identifies the intensity of the signal with perfect fidelity. If, however, perceptual errors are introduced into these models, such that the receiver can make errors in their evaluation of the signal intensity, then the predictions of the models change to being much more in line with the "all or nothing" displays often seen in nature. Such perceptual errors are very plausible (Dusenbury 1992; Hailman 2008).

Here we make no criticism of Johnstone's (or any other previous) work but present another modification to previous models which we argue is biologically realistic, very widely applicable and again leads to a prediction of "all of nothing" displays even when no perceptual errors are assumed in the model. Essentially our key modification rests in the evaluation of optimal predator behaviour. Like previous works, Johnstone assumed that the optimal strategy for the receiver was that which minimized the least-square estimate of signaller quality for each perceived advertising level. That is, the receiver is expected to be selected to evaluate the underlying quality of all individuals as accurately as possible, and all deviations from accurate estimation are in some way costly to the receiver. We suggest that there are many biological situations where the challenge facing the receiver is less strict and some mis-evaluations produce no fitness cost.

Consider again the predator that encounters individuals from a prey population that vary continuously in their level of chemical defence. On encountering a potential prey individual, the predator must make a binary decision: to eat the individual or not. If the predator somehow had complete and perfect knowledge of the level of chemical defence in each prey individual then the most rational strategy is to identify the minimum level of defence that makes a prey individual unattractive, then eat all individuals with levels below this threshold and reject all those with levels above it (Skelhorn & Rowe 2007). The problem for most real predators is that they do not have this perfect knowledge, rather they must make their decisions based on each individual's level of signalling (Mappes et al 2005). Let us imagine that the level of defence can vary between zero and one and the threshold value discussed above is denoted by T. The challenge facing the predator is not to evaluate the defence level of each encountered individual as accurately as possible, but rather to make as few misclassifications as possible as it attempts to classify each individual as having a defence level either above or below T. Another way to look at this is that (unlike the formulation of Johnstone 1994 and other models) not all mistakes in the estimation of a prey individual's level of defense incur fitness costs for the predator. If the true level of defence is D and the predator estimates the defence as a different value d, then this error only has fitness consequences for the predator (it only changes its behaviour) if D and d bracket the threshold value T, otherwise the inaccuracy of estimation has no effect. Further, it may be that the cost of a misclassification to the predator depends upon the value of D, but the value of d has no effect on the size of this cost, except in influencing whether or not misclassification occurs (and thus whether or not the cost is paid). Thus, we suggest that models where receivers can only produce a discrete number of responses to the signal might reasonably involve the assumption that fitness is affected not by accurate estimation of the qualitative value of the underlying quality of signallers, but by the less onerous task of correctly classifying prey into a number of distinct categories. We expect that this situation will occur commonly, where a receiver must make a simple binary choice (e.g. to attack or not, to mate or not, to abandon a nest or not). Here we will explore the consequences of this change of fitness function for model predictions.

Model description

For ease of comparison we have attempted to keep our model definition and structure as close to that of Johnstone (1994) as possible.

We suppose that signallers vary in some quantity that is of interest to receivers, but which they cannot directly observe. We denote the value of this quantity held by a specific individual as q (for quality). Signallers can vary in the intensity of some signal that can be directly observed by receivers, with the signal given by a specific individual being denoted a (for advertising). We denote the function $A(q)$ as the signalling strategy, which specifies the signal intensity (the value of a) given by individuals of different qualities (different values of q).

On receipt of the signal from a specific signaller, the receiver can act in one of only two distinct ways (we denote these alternatives "choice 0" and "choice 1"). The receiver strategy is described by $g(a)$, which is the probability of making choice 1 on receipt of a signal of intensity a. By definition, an individual which does not make choice 1 must make choice 0, and vice versa. Unlike Johnstone (1994), we assume perfect fidelity of signal transmission, so if the signaller sends a value a, the receiver receives exactly that same value.

The reward U that a signaller gets from an interaction with the receiver depends on its quality q, the signal strength it used a, and the response of the receiver (either 0 or 1). Thus the reward to the signaller is $U(a,i,q)$, where i is the response of the receiver: $i \in \{0,1\}$.

We assume that choice 1 by the receiver is always more beneficial to the signaller than choice 0. That is $U(a,0,q) < U(a,1,q)$ for all combinations of a and q values. Thus in our previous example, choice 1 is rejection of the signalling prey by the predator. We also assume that the advantage of choice 1 over choice 0 to the receiver does not decrease with q, i.e.

$$\frac{\partial(U(a,1,q) - U(a,0,q))}{\partial q} \geq 0. \qquad (1)$$

For example, a high-quality male will have at least as large a gain from mating over not mating as a lower-quality male. This seems generally likely to be true for mating systems. For our predator-prey example, the difference between choice 1 and choice 0 is between persuading the predator not to attack versus being attacked. In this case, condition (1) means that even very highly defended prey benefit from persuading the predator not to attack at least as much as weakly defended prey do. Whilst it may be that very highly defended prey can survive attacks because the predator discovers the level of defence during the attack and thus aborts the attack, even such abortive attacks can be costly to prey in terms of risk of injury and/or time and energy wasted. Further, in some situations the predator may have already killed the prey before aborting the attack when realizing that the particular prey item is too defended to be eaten. Thus condition (1) seems plausible in a predator-prey context too.

We further assume that signals are expensive to the signaller, and that this expense increases (and so the net reward from an interaction decreases) with increasing signalling intensity. Thus we assume that for all combinations of (a,i,q),

$$\frac{\partial U(a,i,q)}{\partial a} < 0. \qquad (2)$$

We also assume that the cost of higher signal intensity is proportionately greater for a lower quality individual:

$$\frac{\partial^2 U(a,i,q)}{\partial q \partial a} > 0. \qquad (3)$$

These assumptions about the costs of signalling are those generally considered as requirements for honest signalling via the handicap model (Grafen 1990, Bradbury & Vehrencamp 1998, Searcy & Nowak 2005; but see Lachman et al 2001 for an exception).

The reward to a signaller of quality q that signals with intensity a is given by

$$S_q(a) = g(a)U(a,1,q) + (1 - g(a))U(a,0,q) \qquad (4)$$

We assume that there is only a single type of receiver in our model, so that for instance receivers do not vary in quality and hence in their reward functions. We also assume the reward to the receiver from an encounter is a function of the quality of the signaller *q* and the receiver's decision *i*, which we shall denote by *V*(*q*,*i*), and that the higher the quality of the signaller (the higher *q* is) the better it is for the receiver to make choice 1. That is *V*(*q*,1) – V(*q*,0) increases with *q*. In our example, the more defended the prey individual the more advantageous it is for the predator to reject the opportunity to eat it.

Let *f*(*q*) describe the frequency distribution of signallers of different qualities in the local population (which the receiver encounters randomly). The expected receiver reward is a function of its strategy (*g*) and is given by

$$\begin{aligned} R(g) &= \int f(q)\,V(q,0)\,\bigl(1-g(A(q))\bigr)\,dq + \int f(q)\,V(q,1)\,g(A(q))\,dq \\ &= \int f(q)\,V(q,0)\,dq + \int f(q)\,\bigl[V(q,1)-V(q,0)\bigr]\,g(A(q))\,dq \end{aligned} \qquad (5)$$

where integrals are evaluated over all possible values of signaller quality. We shall assume that in the absence of any signal the receiver will always make choice 0 (e.g. predators must always attack some prey to survive, so in the absence of a signal they will attack all prey rather than none), i.e.

$$\int f(q)\,V(q,0)\,dq > \int f(q)\,V(q,1)\,dq \qquad (6)$$

Model evaluation

We know that *V*(*q*,1) – V(*q*,0) increases with *q;* let us suppose in particular that *V*(*q*,1) – V(*q*,0) < 0 if and only if the quality of the signaller is below some critical value q_{crit}, so we have

$$V(q_{crit},0)=V(q_{crit},1). \tag{7}$$

Thus the receiver would benefit from making choice 0 if and only if $q<q_{crit}$.

Any strategy of the receiver must specify how it responds to every possible signal. Denote the set of all signals a for which the receiver actually makes choice 1 as A_1, and the set of all signals for which the receiver makes choice 0 as A_o. A_1 and A_0 are disjoint sets (no possible signal appears in both sets), and all possible signals are a member of either A_0 or A_1.

Since receivers respond to all signals in A_1 identically, but signals are increasingly costly (inequality (2)) to senders as signal intensity increases, the only rational signal in the set A_1 for a signaller to give is the lowest intensity (cheapest) signal in that set: which we denote min(A_1). Similarly since receivers respond to all signals in A_0 identically, but signals are increasingly costly to senders as signal intensity increases, the only rational signal in the set A_o for a signaller to give is the lowest intensity (cheapest) signal in that set: which we denote min(A_o).

Since $U(a,0,q) < U(a,1,q)$ for all combinations of a and q values, for min(A_0) to be optimal for any q, this implies that min(A_0) < min(A_1); that is that the signal associated with the less favourable receiver choice 0 must be of lower cost, and so at a lower intensity, than that associated with the more favourable choice 1. Since all possible signals are in either A_0 or A_1, the signal associated with 0 will be the cheapest signal of all the possible signals that are open to those individuals ($A_1 \cup A_0$). Thus if the lowest cost signal is $a = 0$, then min(A_0) = 0. Let us further define $a_1 \equiv$min(A_1).

Clearly a_1 must be greater than zero. Thus there are at most two distinct signals in any evolutionarily stable signalling system. A necessary qualification at this point is that this is only true when receivers do not vary in quality to a sufficient degree that different receivers would ideally like to respond to many different signallers in different ways. If there is wide receiver variation, our results would no longer be valid. For instance Johnstone & Grafen (1992) consider the Sir Philip Sidney game where the choice to receivers is to donate food to a relative or not. All receivers survive if they do not donate (and all signallers survive if they receive a donation), but some receivers (signallers) are almost guaranteed to survive if they donate (do not receive), and others are almost guaranteed to die. Under such circumstances, assuming high relatedness, different receivers would "want" to make different decisions to a wide range of signallers (equivalent to having very different values of q_{crit} in our model), and consequently their model has a continuous signalling solution.

It should be noted that our argument about the number of distinct signals generalizes to a system where the receiver has any finite number of decisions n. If we denoted the set of all signals for which the receiver would respond with choice i by A_i, then the only potentially consistent signal choices by the signallers would be $\min(A_i)$, and so the maximum number of distinct signals would be n.

Now let us suppose that we have an "honest" signal, namely one that distinguishes the signallers for which the receiver would want to make choice 0, from those for which choice 1 would be best. This would yield

$$g(A(q)) = \begin{cases} 1, & q > q_{crit} \quad (q \in A_1) \\ 0, & q < q_{crit} \quad (q \in A_0) \end{cases} \tag{8}$$

When the receiver plays this strategy then the reward to the signaller simplifies to

$$S_q(a) = \begin{cases} U(a,0,q) & a \in A_0 \quad (q < q_{crit}) \\ U(a,1,q) & a \in A_1 \quad (q > q_{crit}) \end{cases} \tag{9}$$

Thus the optimal signalling strategy associated with an honest signal should be

$$A(q) = \begin{cases} \min(A_0) = 0, & q < q_{crit} \\ \min(A_1) = a_1 > 0, & q > q_{crit} \end{cases} \tag{10}$$

For there to be a stable signalling strategy where all $q < q_{crit}$ individuals pick 0 and all $q > q_{crit}$ individuals pick a_1, for some positive a_1, we need both choices to offer the same reward to the signaller when $q = q_{crit}$ (otherwise individuals of quality either just above or below q_{crit} could do better by switching signal). Thus we need

$$U(a_1,1,q_{crit}) = U(0,0,q_{crit})\,. \tag{11}$$

Since $U(a_1,1,q)$ decreases with increasing a_1, there is at most one value of a_1 that satisfies (11). Such a value will exist provided there is such a critical quality value q_{crit} where the receiver would want to change their strategy, and that the largest signals are sufficiently costly, so that $U(\infty,1,q_{crit})<U(0,0,q_{crit})$. Thus $[0, a_1) \subseteq A_0$ and $a_1 \in A_1$. In fact we shall assume the natural solution of $A_0 = [0, a_1)$ and $A_1 = [a_1,\infty)$.

Inequalities (1) and (2) ensure that for lower quality individuals the relative costs of signalling compared to the benefits of receiving choice 1 are higher, and consequently any individual of quality $q < q_{crit}$ would do worse by changing its signal to a_1 or any other value in A_1, and any individual of quality $q > q_{crit}$ would also do worse by switching signal. Note that the combination of (1) and (2) are sufficient but not necessary, so that the relative costs compared to benefits may decrease with quality even if only one of the two conditions hold.

Note that the exact composition of the sets A_0 and A_1 in such a system depends upon how rogue signals not equal to 0 or a_1 come about. Any individual that uses such a signal is behaving sub-optimally, so we would expect such situations to be rare. The exact solution in these rare cases would depend upon assumptions about the underlying causes of such irrational behaviour (see Discussion).

It should also be noted that only two signals are used at equilibrium, and that if there are no rogue signals as described above, every receiver strategy that responds to these two signals in the same way thus performs equally well at the equilibrium, regardless of how they respond to other signals. We assume that there will be a low level of such “mistakes” which means that all receivers have to play optimally against the "non-played" strategies themselves. This idea is often used in game theoretical modelling, and is known as the “trembling hand” (Selten, 1975).

It is possible to envisage a signalling system that is not entirely honest. For stability all low-quality individuals must play 0, and all high quality individuals must play min(A_1); but perhaps there can be a cut-off point q^* that is different to q_{crit}. If we replace q_{crit} by q^* in (8-11), we would obtain a different equilibrium signalling system with a new level a^* for the higher signal. In the case where $q^*>q_{crit}$, so that $a^*=\min(A_1)> a_1$, such a system could be destabilized by the introduction of a signaller that included $a_1 \in A_1$, which would enable individuals with qualities $q^*>q>q_{crit}$ to signal honestly to the benefit of themselves and the receiver. There will also be a value q_{min} so that if $q^*\leq q_{min}$, (i.e. if q^* is sufficiently small), then (due to inequality 6) the expected reward to the receiver will be at least as high if it changes to make choice 0 against all signals, and so again the system is not stable. This leaves a family of possible "semi-honest" signalling systems with cutoff q^* such that $q_{min} < q^*\leq q_{crit}$ that might be stable in some circumstances (when the "honest" solution also exists). Note that such alternative solutions are "semi-honest" in the sense that every individual giving the higher signal is of better quality than every individual giving the lower signal. However, some individuals with qualities near to (and on one side of) the critical value will gain advantage by using the "wrong" signal from the receiver's viewpoint. Thus it is important to note that we do not claim that the fully honest signal is the one that the population will evolve to. We have shown, however, that such a system is a possible solution, and that all of the other potential solutions have the same all-or-nothing property.

The general solution for our model is that signallers below a defined quality threshold all signal using the lowest-cost signal that is possible, and receivers respond to this signal with the choice that least benefits signallers; signals with quality above this threshold all signal using the same signal, this is a higher cost signal than that used by low-quality individuals and is the signal that leads to the same payoff to individuals of the critical quality regardless of what the receivers do. Receivers respond to the higher-cost signal by adopting the behaviour (from a choice of two) that is more beneficial to signallers.

Thus, although signallers vary continuously in quality, they do not show continuous variation in signal strength at this equilibrium. Rather, the discrete nature of the behavioural responses to signals available to the receiver causes the receiver to be interested in categorizing signallers rather than fully evaluating their quality, and this in turn leads to signalling being restricted to a number of discrete levels, less than or equal in number to the number of behavioural options open to the receiver.

An example

Let us consider a simple example where males of quality q signal to females, who can choose either to mate with a specific male or not.

For the female, there is no reward (or cost) for declining to mate $V(q,0) = 0$. Mating requires a fixed cost (α) and benefits increase linearly with the quality of the male. Thus, at its simplest $V(q,1) = q - \alpha$.

For the male, there is a cost for an individual of quality q to produce a signal of strength a given by a/q. There is an additional payoff of unity if the female chooses to mate and zero otherwise. Thus,

$$U(a,0,q) = -\frac{a}{q},\ U(a,1,q) = 1 - \frac{a}{q}.$$

Substituting these into (7) and (11) yields the solution $a_1 = q_{crit} = \alpha$.

Thus under fully honest signalling we predict that males with quality lower than $q = \alpha$ will signal using the lowest-cost signal available and will always be rejected by females; whereas males with a higher quality than this will signal at level α and will always be mated with by females.

It is easy to see the rationality of this in the very simple case considered. At the equilibrium females always mate with males that offer a net benefit to them, and never mate with males that offer a net loss to them. Given this behaviour by receivers, the minimal-cost signalling of low quality males also seems easy to understand. Since these individuals are destined to be rejected by females, their signal can bring them no rewards and so the best strategy is to minimize the costs of signalling. However, investment in more expensive signalling is rational for the high quality individuals since they can convert this advertising into rewards (mating opportunities). Still they should be selected to invest just enough in advertising to both produce the desired behaviour in the receiver, and to prevent the best of the poor males from cheating. The payoff to low-quality, minimum-cost signallers is zero, the signal level adopted by the high-quality individuals is the cheapest signal that yields a net positive payoff to all individuals that use this signal (except any right on the threshold, who also receive zero).

Discussion

In this paper we have considered a model of signalling behaviour where the receivers have only a discrete number of possible responses to the signal. Our model predicts that even if signallers vary continuously in quality, and signals are received with perfect fidelity, these signals need not show continuous variation in signal strength. Rather, the discrete nature of the behavioural responses to signals available to the receiver causes the receiver to be interested in categorizing signallers rather than fully evaluating their quality, and this in turn leads to signalling being restricted to a number of discrete levels (at most equal in number to the number of behavioural options open to the receiver). Thus we predict that such signals will be commonplace when the behavioural responses of receivers are constrained to take a discrete number of values. Examples of this could include signalling of prey toxicity to predators, where predators can respond either by eating an individual signaller or rejecting the opportunity to eat it. Another example may be mate choice where the choice is again binary: mating with or rejecting the signaller. We thus expect such situations and such all-or-nothing signalling to be commonplace. However, there are other cases where the responses of signal receivers may be more continuously distributed. For example, in response to signal quality of a long-term social partner, a female bird may vary the investment that she makes in the eggs that will become their joint-offspring (Clutton-Brock 1991; Blount et al. 2000). This investment (say in levels of anti-oxidants committed to the eggs) is best seen as a continuously varying response, and so we would predict that the signalling behaviour of the males would not be well represented by the model considered here and (in the absence of perceptual errors) we would consider a continuously distributed signal by the males to be more likely.

Bergstrom & Lachman (1998) present a model that they use to suggest that honest signaling between relatives can be maintained in the absence of substantial costs to signal production. The type of equilibrium that they consider are of the all-or-nothing type discussed here, where signallers of a range of qualities are grouped into a finite number of what the authors term “pools” with all individuals in the same pool producing the same signal. However, a very important difference between our approach and theirs is that a finite number of signal levels is a prediction of our model, whereas the signal being constrained such that only a finite number of signal types are possible is a fundamental assumption of their model. Our methodology does not involve any such constraint on signal production.

The all-or-nothing signalling predicted here may not be seen in situations where there is strong between-individual variation in the receivers in the value of the signallers to them. Consider the example of predators and chemically defended prey. Previously we have considered a critical value of toxins above which the prey becomes unattractive to the predators. There may be some circumstances where individual predators essentially agree on this critical value, in which case we would expect our model to hold. However, there may be other circumstances where there is considerable variation in this value between individual predators. This could be driven by variation between individuals in the need for the nutritional benefits of the prey (with hungrier individuals being prepared to accept higher toxin loads to avoid the risk of starvation) or variation in their ability to cope with the toxins (perhaps through variation in their current toxin burden): see Endler & Mappes (2004) for examples. If this variation in threshold of defence is large then this may cause the all-or-nothing type of signal predicted here to break down and be replaced by a more continuously-varying signal, as in [10].

Johnstone (1994) cited a number of influential papers that predict (in contrast to our model) that signal intensity should vary continuously in relation to the quality or need of the signaller: (Grafen 1990, Godfray 1991, Johnstone & Grafen 1992, Pagel 1993). In each case, it is possible to explain why these models make different predictions to ours. As already discussed, in Johnstone & Grafen (1992) wide receiver variation causes different receivers to wish to respond to many different signallers in different ways, making variation in signalling level viable. In Grafen (1990) and Pagel (1993) this difference is due to the cost function, which they make an explicit function of the error in perception of underlying signaller quality, so that there is a cost which continuously increases as a function to the size of the perceptual error. This is the situation we discussed in the introduction where all errors are considered to be costly. The exact mechanism underlying these costs is not defined in these papers, and choices available to the receivers (on receipt of a particular signal value) are not explicitly given. In Godfray (1991) the choices are explicitly given; these are the possible levels of provisioning by a parent to its offspring. This provisioning effort is considered to vary continuously, so there is a continuum of choices (rather than the binary choice considered here), and thus the scenario is different to ours, and (in the absence of perceptual errors) a continuously varying signal intensity is certainly plausible here.

Notice that the receiver strategy as we have defined it only describes responses to the two types of signal that are expected in the equilibrium situation. There may be occasional aberrant individuals that produce signals that are different from either of the two signals that form the equilibrium. It is likely that the receivers will treat such a signal in a way similar to whichever of the two equilibrium signals it most resembles, with the similarity of response getting stronger as the similarity between aberrant and nearest-equilibrium signals increases. Such generalization across similar signal types is commonly observed empirically (Bradbury & Vehrencamp 1998). However if signals just below the higher signalling level are always treated as the higher signal, the signalling system will be destabilized, so there must be at least some probability of such signals being treated as a low signal for any system to be stable (this would only need to be small for small discrepancies, since the benefit from using a lower-cost signal is greatly outweighed by the cost of being interpreted as a low signal). Overall, the optimal strategy for receivers to deal with aberrant signals will depend on the exact biological mechanism that leads to the production of aberrant signals, since the fine detail of this mechanism will influence the probability distribution of individual signaller qualities (q values) associated with a particular aberrant signal strength. However, we might not expect to see natural receivers closely following this theoretical optimum strategy, since aberrant signals will be rare and so selection pressure shaping responses to such signals will be less that selection on responses to more commonly encountered signals. Rather we might expect to find between-receiver variation in response to aberrant signals (Arak & Enquist 1993), but with all receivers generally showing the rational behaviour of generalization across similar signals such that they treat aberrant signals (in particular high signals) in a way that is like their treatment of the most similar of the signals that makes up the equilibrium set.

In this paper we have been particularly interested in how an honest signalling system could work in our chosen scenario, and this has been our main focus. However, we found that we could not discount the possibility of what we called a semi-honest system, where higher signals mean a better quality individual than lower ones, but where the cut-off is not that of the totally honest signalling system. It may be that such systems can be destabilized through the introduction of signalling errors, as in Johnstone (1994), or alternatively through receiver variation, and this would certainly be worth further investigation.

References

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

Bergstrom, C. T., and M. Lachmann 1998 Signalling among relatives III Talk is cheap. Proc. Nat. Acad. Sci. 95: 5100-5105.

Blount, J. D., D. C. Houston, and A. P. Moller 2000. Why Egg Yolk is Yellow. Trend. Ecol. Evol. 15: 47-49.

Bradbury J. W., and S. L. Vehrencamp 1998. Principles of animal communication. Sinauer, New York, USA.

Clutton-Brock, T. H. 1991. The Evolution of Parental Care. Princeton University Press, Princeton, USA.

Dusenbery, D.B. 1992. Sensory Ecology. Freeman, New York, USA.

Endler, J. A., Mappes, J. 2004. Predator mixes and the conspicuousness of aposematic signals. Am. Nat. 163: 532-547.

Godfray, H. C. J. 1991. Signalling of need by offspring to their parents. Nature 352: 328-330.

Grafen, A. 1990. Biological signals as handicaps. J. theor. Biol. 144: 517-546

Hailman, J. P. 2008 Coding and redundancy. Harvard University Press, Harvard, USA.

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. and A. Grafen 1992. The continuous Sir Philip Sidney game: a simple model of biological signalling. J. theor. Biol. 156: 215-234.

Lachmann, M., S. Szamado, and C. T. Bergstrom 2001. Cost and conflict in animal signals and human language. Proc. Nat. Acad. Sci. 98:13189-13194

Mappes, J., N. Marples, and J. A. Endler 2005. The complex business of survival by aposematism. Trend. Ecol. Evol. 20: 598-603

Maynard Smith, J., and J. Harper 2002. Animal signals. Oxford University Press, Oxford, UK.

Pagel, M. 1993. Honest signalling amount gametes. *Nature* 363, 539-541

Searcy, W.A., and S. Nowicki 2005. The evolution of animal communication. Princeton University Press, Princeton, UK

Selten, R. 1975. A reexamination of the perfectness concept for equilibrium points in extensive games. International Journal of Game Theory 4, 25-55.

Skelhorn J. and C. Rowe 2007. Predators' toxin burdens influence their strategic decisions to eat toxic prey. Curr. Biol. 17: 1479-1483