Can behavioural constraints alter the stability of signalling equilibria?

Stephen R. Proulx

Department of Biology, 1210 University of Oregon, Eugene, OR 97403-1210 (proulx@proulxresearch.org)

In traditional models of signalling one class of individual, the signaller, presents a signal which another class of individual, the receiver, examines. Receivers are typically assumed to have fitness returns that depend on their ability to determine the utility of the signaller to them. Each signaller must decide what level to signal at, which is a function of the quality of the signaller. In addition, a signaller's quality is assumed to be synonymous with the signaller's utility to a receiver. However, there is no reason to believe that signalling costs are incurred in the same currency as the receivers are paid and, thus, no reason to believe that the relationship between signaller quality and utility is linear or even increasing. For instance, in signalling between prey and predators, the utility of a prey item may be its fat reserves, whereas an individual prey pays for signalling (and thus measures quality) in terms of increased risk of capture; quality and utility are synonymous only if a high risk of capture is associated with high fat reserves. In addition, several recent studies have documented increased signalling as utility decreases. If utility and quality are decoupled, so that increasing quality does not always mean increasing utility, then traditional signalling models predict that no signalling equilibrium will exist. I show that if receiver fitness is modelled by a set of behavioural responses, which have both costs and benefits, then a signalling equilibrium can sometimes be recovered. An example of signalling between mates is presented in order to demonstrate this equilibrium.

Keywords: sexual selection; female choice; spatial heterogeneity; joint ESS; evolution

1. INTRODUCTION

In a signalling game, two classes of individuals, signallers and receivers, interact. The fitness of each individual is determined in part by the actions of the individuals in the other class. For example, in one of the most studied signalling games, males signal to females by displaying elaborate ornaments or performing display activities and females choose a mate based in part on these displays. The signalling strategy of the males is a strategic decision to display more or less depending on the males' quality (the amount of their energy reserves, body size or vulnerability to predation). The receiver's strategy is typically thought of as deciding how much value to ascribe to each possible signal (Grafen 1990). The main goal of signalling theory is to identify conditions that promote a stable signalling equilibrium, where signalling reveals information about signallers. At this equilibrium, individual signallers would lose fitness by altering their signalling strategy and individual receivers would lose fitness by altering their perception of signals.

Traditional signalling theory has been built on the notion that receiver fitness depends only on receiver perceptions, while signaller fitness depends on an interaction with the receiver (Grafen 1990; Johnstone & Grafen 1992, 1993; Johnstone 1995, 1998). However, perception is an abstract concept and receiver fitness can only depend on the actions taken in response to observations. By an action, I mean the behaviours that a receiver can use in order to alter the probability of interacting with particular signallers. For example, a female bird may decide to visit a number of leks before making a mating decision, thus increasing her sample size while incurring a time cost. In a game based purely on perception, the female bird's fitness is defined by her ability to identify males

accurately regardless of whom she mates with. Alternatively, in a model that considers actions, the fitness of the female bird is determined by the costs associated with searching for mates, the set of possible mate-searching activities and the probability of finding a high-utility mate given the search strategy.

Clearly, in models where signaller fitness depends on an interaction with receivers an action must have occurred. However, the way in which perception is translated into action can alter the stability of a signalling equilibrium. The development of models which include the translation of perception into action and showing how the form of the translation affects the signalling equilibria is the focus of this paper.

Two assumptions have been required by traditional models of signalling: an increase in the marginal fitness returns of signalling as signaller quality increases and the equivalence between signaller quality and the utility of signallers to receivers. The first of these may be termed the 'handicap criterion' and, although the verbal interpretation has been debated, it must be satisfied for signal strength to be an increasing function of quality (Getty 1998; Szamado 1999). The second assumption, that quality and utility are essentially the same, seems to have been tacitly accepted but can have a substantial impact on signalling equilibria.

The implicit assumption that signaller quality is synonymous with signaller utility probably stems from Zahavi's (1975) seminal paper on the handicap principle. Signaller quality can be thought of as the ability of the signaller to survive in spite of a costly signal (Grafen 1990), while I define signaller utility as the fitness return that a receiver obtains by interacting with a signaller. Thus, quality determines the signalling strategy while utility determines the receiver strategy. In Zahavi's (1975)

original framework, the genetic architecture of handicap traits is deemed irrelevant because 'the handicap it has already imposed is a proof that the [genotype] is above a certain level of quality...' (p. 208). Likewise, this assumption is implicitly carried through in subsequent works, although Grafen (1990) foreshadowed the possibility of a utility/quality curve in saying that 'all the conditions on [fitness] would be unchanged [under a transformation by] increasing invertible differentiable functions' (p. 544). Traditional signalling equilibria rely on a strictly increasing relationship between utility and quality; even a small region of non-monotonicity can destabilize an equilibrium.

It is not difficult to imagine that signaller quality is not exactly what receivers are attempting to evaluate. For example, in signalling of need between offspring and parents, juvenile viability is only one component of total future offspring success. In predator-prey signalling, predators are attempting to assess both prey escape ability and the nutritional benefit that each prey item contains, while prey survivorship with respect to predation depends only on escape ability. In sexual signalling games, females may be interested in male parenting effort, while male fitness may also depend on the ability to procure extra-pair copulations. A male who has high energy reserves and, therefore, a high signalling level may spend less time attending to his own chicks and more time searching for other mates. The biological relationships that change the quality/utility curve to the point where it is no longer strictly increasing might seem less clear at first glance, but are predicted for sexual selection signalling due to life-history considerations (Kokko 1997; Getty 1998) and can be caused by signaller migration in a spatially heterogeneous environment. For example, if males can allocate energy to growth or signalling then high-quality males may be able to reach a larger size earlier and invest less in signalling in order to gain a longer lifespan. Males of lower quality will wait to signal, but then signal at higher levels than high-quality males (Kokko 1997). Moreover, several recent studies have found regions of negative correlation between quality (as measured by viability or lipid stores) and signalling (Lotem 1998; Candolin 1999; Brooks 2000), which is consistent with a non-monotonic relationship between quality and utility. Another possibility is that receivers may prefer a signaller of intermediate quality; parents may not wish to feed the neediest chick or the strongest chick; rather, the future reproductive success of chicks with intermediate reserves may be most sensitive to feeding.

Another source of discrepancy between signaller quality and utility comes from perceptual error (Johnstone & Grafen 1992). The signal level perceived by receivers may be influenced by developmental error on the part of the signaller or errors of assessment by the receiver. In either case, as long as the expected quality of the receiver is increasing in perceived signal level, a signalling equilibrium can occur (Johnstone & Grafen 1992). If perceptual error is more negatively biased for higher-signalling individuals and positively biased for lower-signalling individuals, then traditional signalling theory cannot explain the stability of observed signalling systems.

I first discuss the reasoning behind traditional signalling models and then show how a non-monotonic utility/

quality relationship makes a signalling equilibrium impossible. This result has been alluded to before (Grafen 1990), but is presented here in order to provide a contrast with the results of § 4. I then discuss how explicit inclusion of an action set allows a signalling equilibrium to persist, even when the utility/quality relationship is non-monotonic. In simple terms, whenever the receiver is constrained to respond to signallers based only on relative signal strength, then the signalling game can be stabilized. I then present an example of signalling by males when quality is not the same as utility and females are able to spend more or less time searching for males. In this example, a signalling equilibrium is preserved unless utility and quality are so unrelated that mating at random results in higher returns than mating with males who signal the most. I conclude that signalling equilibria can exist under a wider range of conditions than was previously thought.

2. TRADITIONAL MODEL OF A SIGNALLING EQUILIBRIUM

Models of handicap signalling typically make two assumptions about signaller fitness, which can be described as signal cost and a handicap criterion describing how the costs and benefits of signalling change as the quality of the signaller changes (Grafen 1990; Getty 1998). If we adopt the simplifying assumption that the gains due directly to signalling depend only on the level of signalling (and not quality) then these conditions become

$$\frac{\partial w(a, p, q)}{\partial a} < 0 \quad (2.1)$$

and

$$\frac{\partial^2 w(a, p, q)}{\partial a \partial q} > 0, \quad (2.2)$$

where $w(a, p, q)$ is the fitness of a signaller with signal strength a , perceived at quality p and having quality q . It is important to note that these are partial derivatives; they represent the fitness effects of changing a single component while holding all others constant. Equation (2.1) simply states that signalling has a direct negative effect on fitness (it also has an indirect positive effect through p). Equation (2.2) is more difficult to interpret verbally, but says that increasing signalling has less of a negative effect on the fitness of high-quality individuals. For a further discussion of how equation (2.2) should be interpreted verbally, what the assumptions which lead to equation (2.2) mean and how it corresponds to field measurements of the survivorship costs of display see Hurd (1995), Getty (1998) and Szamado (1999). Regardless of the interpretation, equation (2.2) is generally considered the simplest criterion necessary for handicap signalling and I will refer to it as the handicap criterion.

A purely perception-based model of signalling behaviour assumes that receiver fitness is maximized by accurately assessing signaller utility (Grafen 1990; Johnstone & Grafen 1992, 1993; Johnstone 1995, 1998). It is as if each receiver is shown a signaller and graded on how

Figure 1. The logic of signalling games when utility is a strictly increasing function of quality. (a) The return to signalling as a function of signal strength shown here as an increasing function. This is determined by the receiver strategy. (b) Because returns to signalling are increasing in signal strength, the handicap criterion causes the optimal signal level to be an increasing function of signaller quality. (c) Because utility is strictly increasing in quality, (d) utility will also be increasing in signal. Because utility is increasing in signal, receivers will prefer to interact with individuals with large signals, so that returns to signalling are increasing.

closely they describe the signallers quality. Thus, receiver fitness is written as

$$R(P) = \int_0^\infty -|P(a) - U(a)|g(a)da, \quad (2.3)$$

where $U(a)$ is the expected utility of a signal at a and $g(a)$ is the probability density function for signalling levels. Receiver fitness is a function of the perception function $P(a)$ and is always maximal when $P(a) = U(a)$. In other words, receivers will choose a perception function that maps signals to their expected utility. Again, the perception function just describes how well receivers match signals with quality and only implies a relationship to any behaviour a receiver might perform after observing a signal. Receivers are free to choose any perception function and are assumed to choose the one which maximizes fitness.

Given a receiver preference rule and a relationship between preference and signaller success, we can solve for the equilibrium signalling strategy, which is a function of quality, not utility. An equilibrium will exist when a receiver perception rule and signaller rule make a unilateral departure from the equilibrium disadvantageous (Lachmann & Bergstrom 1998). If signalling has an indirect positive effect on signaller success (figure 1a) then, because of the handicap criterion, the signal is increasing in quality (figure 1b). This is true regardless of the form or strength of receiver preference, as long as increased signalling increases preference. If utility is increasing in quality (figure 1c), then, because the signal is increasing in quality, the expected utility is increasing

Figure 2. The logic of signalling games when utility is a non-monotonic function of quality. (a) As in figure 1 returns to signalling are assumed to be increasing in signal strength. (b) Again, the handicap criterion causes the optimal signal level to be an increasing function of signaller quality. (c) Utility is no longer strictly increasing in quality, (d) so neither is it strictly increasing in signal. However, if the utility of the highest signalling individual is greater than the average utility, then a receiver preference for a high signal can still be advantageous.

in signal (figure 1d). This means that, if a receiver strategy which involves a preference for higher signals were fixed, the signaller strategy would evolve to yield an increasing signal as a function of utility, which would in turn select for receiver preference for an increased signal. Because this model assumes perception maps directly to receiver fitness, receiver perceptions will exactly match signaller utility because of equation (2.3). Thus, the evolutionarily stable strategy (ESS) is for receivers to perceive signallers at their utility, and the signaller strategy is the optimal signal given that receivers have perfect perception. From this logic, it is clear that, as long as the expected utility is increasing in quality, a signalling equilibrium will exist. The result of Johnstone & Grafen (1992) on signal error is that an increase in signal must bring with it an increased perception of quality, which is synonymous with an increasing expected utility/quality curve.

3. SIGNALLING EQUILIBRIA BREAKDOWN UNDER IMPERFECT CORRELATIONS BETWEEN QUALITY AND UTILITY

When the expected utility is not a monotonic function of quality, then some pairs of quality have the same utility. This means that, at an equilibrium, receivers will perceive individuals with different qualities as having the same utility, even if they signal with different strengths (figure 2d). The total derivative of signaller fitness with respect to signal strength must be zero if an equilibrium has been reached, so that

$$\frac{d}{da}(w(a, p, q)) = 0, \quad (3.1)$$

which must hold for all signallers. This means that the total effect of increasing or decreasing a (which includes an indirect effect through p) is not an increase in signaller fitness. For a pair of signallers who have different qualities, but are perceived by receivers to have the same expected utility, the two signallers can either signal at the same level or at different levels. If they signal at different levels, then the one signalling at the higher level is paying a higher cost than necessary and should switch to signalling at the lower level. If they signal at the same level then, because they have different marginal costs (equation (2.2)), but identical returns for signalling, equation (3.1) cannot hold for both (no pooling equilibrium is possible because the signaller strategy is assumed to be continuous). Thus, no signalling equilibrium is possible, no matter how slight the deviation from monotonicity.

In some sense, the breakdown of the signalling equilibrium is due to the perfection with which receivers map a signal to utility. Receivers are capable, by assumption, of matching signals to uses regardless of how complex the relationship between the signal and utility is, and without any behavioural limitations. This view ignores the costs of information gathering and constraints on receiver signal processing. It is important to note that this issue is distinct from that of perceptual error; here, the question is not about how well receivers can identify a signal, but rather how well they can map an identified signal to a utility, and then choose an appropriate action.

4. WHEN RECEIVER ACTIONS ARE CONSTRAINED, THE SIGNALLING EQUILIBRIUM IS PRESERVED

Now I want to consider the set of actions a receiver can take. By considering how perception translates into action, which affects both receiver and signaller fitness, we can identify likely constraints on receiver actions and see that these constraints may increase the domain of stable signalling. For example, the rule 'associate only with the highest signallers' may be impossible for orangutans searching for mates, because individuals meet only occasionally and in small groups. While the rules which govern possible actions are themselves in the evolutionary arena, they are likely to be influenced by many factors, so that the forces imposed by the signalling game cannot override all other selection on perceptual apparatus or behaviour. Other factors which may limit the set of receiver actions are the costs of gathering, storing and analysing information (Luttbeg 1994) and conflicting demands on the cognitive system (Dukas 1998).

The reason why no signalling equilibrium can exist in the purely perception-based signalling model is that receivers are able to ascribe the same value to individuals who signal at different levels. The receiver must adopt a perception rule that requires not just relative comparisons of a signal strength, but also determination of a signal on an absolute scale. This 'perfect pitch' receiver may require additional neurological functions which could entail physiological costs. Furthermore, if the quality/utility curve undergoes temporal shifts, a 'perfect pitch' receiver, even though selected for in one year, would be selected

against the next, possibly leading to canalization of the receiver strategy (Kawecki 2000).

Until now, I have not specified anything about how receiver fitness operates, only saying that increased perceptual accuracy is better. In most signalling games, receivers do best by biasing interactions towards signallers of the highest utility. For example, in parental care systems parents wish to increase their allocation to the offspring with the highest marginal fitness returns to feeding, in predator-prey systems predators do best by attacking prey items that are the most vulnerable, and in mating systems females do best by mating with males who give the most care to her offspring. Consider a receiver that can choose a strategy from a set of possible actions, i.e. behaviours that bias interactions towards signallers with higher signal strengths. This receiver only needs relative information and, by biasing interactions towards higher signals, the receiver may increase the likelihood of interacting with a high-utility signaller. The strategy of receivers may involve costs associated with increased bias in terms of time or energy, so that the optimal receiver strategy will be a balance between increased association with high-utility signallers and costs. The least costly receiver strategy will have some other return associated with it; this could be simply the average of signaller uses. As the bias of the receivers becomes high, they will only interact with signallers of the highest quality. If the utility of the highest quality signaller is higher than the average utility then a signalling equilibrium results (figure 2). Thus, even when the relationship between utility and quality is not strictly increasing, a signalling equilibrium is possible.

In this new framework, the fitness of the receiver can be written as a function of the strategy adopted. If the receiver strategy depends only on one parameter s , which could be the amount of time spent examining signallers, then receiver fitness can be written as

$$R(s) = C(s) \int_0^\infty I(s, a, g(a)) E[U(a)] da, \quad (4.1)$$

where $C(s)$ is the residual fitness after the cost of the receiver strategy s , $I(s, a, g(a))$ is the probability of interacting with a signaller signalling at a given a strategy s and probability density $g(a)$ of signallers, and $E[U(a)]$ is the expected utility of an individual signalling at a . As long as increasing s increases the average returns to a receiver (given that signalling increases in q) and increasing s maintains a signalling advantage for signallers (causing signalling to increase in q), a stable signalling equilibrium can result. Equation (4.1) superficially resembles equation (2.3), but includes the costs for receiver behaviours and links the success of the receiver to fitness through the interaction function (I).

In order to see how this can stabilize a signalling interaction, first consider a situation where all behaviours are equally costly, so $C(s)$ is a constant and has no effect on the maximization of R . If increasing s increases the mean signalling level that a receiver interacts with then, when s becomes very large, the receiver only interacts with individuals near the maximum signalling level. Thus, as long as the maximum signaller's utility is above average, a high s is better than a low (or zero) s . In general, if increasing s

Figure 3. Non-monotonic utility/quality curves. The values of q_0 are 1, 0.75 and 0.0. When $q_0 > 2/3$, the utility of the male with the highest quality is greater than the average utility.

tends to increase associations with high signalling individuals, then a signalling equilibrium will result as long as the cost is not too high and the utility of high signalling individuals is greater than the average utility amongst all signallers.

Of course, by considering only strategies based on the relative signal strength, the receiver strategy of associating each signal with its actual utility is ruled out. In some sense, constraints on receiver perception are what cause the stability of the signalling equilibria. However, when receiver strategies are conceived as being drawn from a set of possible actions, the costs and constraints on those actions become important. If the benefits of increasing returns to 'perfect pitch' receivers outweigh the costs of increasing brain function and signaller quality distributions are stable in time, then 'perfect pitch' receivers may evolve and the signalling equilibrium will break down.

5. AN EXAMPLE: FEMALE CHOICE AND MALE ADVERTISING

In a 'best of n ' mating system (Janetos 1980; Seger 1985), females examine n males and mate with the male that most closely matches her preference. I assume that the number of males examined, n , describes a female strategy and that females prefer males with the highest signalling level. Under a best of n mating strategy, the probability of mating with a male of quality q is independent of the signalling strategy as long as the signal strength is increasing in quality. This is because females are only paying attention to the relative signal strength and, as long as the conditions in equations (2.1) and (2.2) are met, the signalling function will be increasing in q . The mating probability density function is

$$M(Q(a)) = f(Q(a))nF(Q(a))^{n-1}, \quad (5.1)$$

where F and f are the cumulative and probability density functions for the frequency of male qualities, respectively, and $Q(a)$ is the inverse map of signalling to quality. For simplicity, I assume that male quality is uniformly distributed and that the total reproductive output of males is the product of survivorship and mating success.

One survivorship function which meets the conditions in equations (2.1) and (2.2) is given by

$$s(q, a) = qe^{-a/q}. \quad (5.2)$$

One can calculate the male signalling function given the female mating rule by substituting $w = s \times m$ into equation (3.1). Making the replacement of $q = Q(a)$ gives the differential equation

$$Q'(a) = \frac{1}{n-1}, \quad (5.3)$$

which has solution

$$Q(a) = \frac{a}{n-1}, \quad (5.4)$$

which is the inverse of the optimal signalling strategy

$$A(q) = q(n-1). \quad (5.5)$$

Because the mating probabilities do not depend on the signalling rule (as long as it is increasing), they likewise do not depend on the population female mating strategy. While females may receive benefits due to the increase in expected utility by increasing n , they pay a cost in search effort, which is assumed to be exponential in form ($e^{-\lambda(n-1)}$). Thus, female fitness is defined as

$$W(n) = e^{-\lambda(n-1)} \int_0^1 U(q)M(q)dq, \quad (5.6)$$

where $U(q)$ is the utility of males with quality q . Female fitness is the composition of search costs and mating returns. As n approaches infinity the mating return portion of female success in equation (5.6) approaches the utility of the highest quality male ($U(1)$). If $U(1)$ is greater than the utility received by a female who examines only one male, then for a small enough cost parameter λ , a signalling equilibrium exists.

In order to illustrate this with an example, I assume that $U(q) = 1 - (q - q_0)^2$, which could be due to male migration with balancing selection on q , life-history considerations or to the physiological relationships between utility and quality. Migration could induce a non-monotonic relationship between quality and utility if different genotypes are favoured in different locations, so that immigrating males in high condition have genotypes which are maladapted to local conditions. The utility/quality curve can go from strictly increasing through to non-monotonic to strictly decreasing by altering q_0 (figure 3). Substitution into equation (5.6) yields

$$W(n) = e^{-\lambda(n-1)} \{n^2(2 - q_0)q_0 - n(3q_0^2 - 4q_0 - 2) + 2(1 - q_0^2)\} / \{(n+1)(n+2)\} \quad (5.7)$$

The maxima of equation (5.7) can be easily found numerically. As long as the utility of an average male is less than the utility of the highest-quality male, a signalling equilibrium is possible. When $q_0 < 2/3$, the equilibrium for females is to mate randomly no matter how small the cost of being choosy. In this case males do not signal and females do not choose.

Figure 4 shows the fitness of a female who examines n males when male quality is uniformly distributed, $q_0 = 0.8$ and $\lambda = 0.005$. An equilibrium occurs when

Figure 4. Female fitness as a function of choosiness. Females who spend more time searching for mates receive benefits by associating with higher utility males, but also pay increasing costs. The optimal level of female choice is the point where female fitness is maximized. Because female fitness only depends on the relative signal strength in a best of n model, the optimal female strategy is independent of the exact signalling function. The parameters used to generate this figure are $q_0 = 0.8$ and $\lambda = 0.005$.

females examine the number of males which maximize female fitness and males adopt a signalling strategy that is the solution of equation (3.1), given that females have adopted the equilibrium level of female choice. For these parameters, the female ESS is to examine an average of 4.7 males and the equilibrium male signalling function is $A(q) = 3.7q$.

6. CONCLUSIONS

When signallers and receivers value the same quantity, the handicap principle works precisely because signallers cannot lie to themselves. In this case, the handicap criterion guarantees that, as long as higher signalling brings increased returns, individuals of higher quality will signal at higher levels. In the three common biological examples of signalling (predator–prey, parent–offspring and male–female) the signaller is trying to show that it has the most of some quantity (escape ability, need and mate quality). Previous theoretical studies have concluded that stable signalling can only occur when utility and quality are related and the handicap criterion holds for all levels of signaller condition. Any deviation from these assumptions will destabilize the equilibrium. However, these results are based on an exact mapping of perception to receiver fitness and do not consider the costs and constraints on receiver actions.

In previous models, receiver fitness was determined by the receiver's ability to identify the utility of all signals accurately. This must be motivated by some belief of how receiver response depends on perception, but this step is missing from most models of signalling. When an action set is considered, both receiver and signaller fitness depend on the action taken by the receiver. This means that the optimal level of signalling will depend on the receiver strategy, a result which is missing from previous game theoretical analyses (Grafen 1990; Johnstone & Grafen 1993, 1992; Johnstone 1995, 1998). In addition to

making receiver fitness more explicit, this framework allows incorporation of the costs to receivers based on their strategies. Thus, when a constrained set of actions is considered, signalling equilibria can be found under more general conditions. Moreover, the constrained action set can itself be considered evolved under the hypothesis that more complex behavioural repertoires involve the production of costly neuronal machinery, which has been suggested as a cost for increased memory (Dukas 1999).

The simplifying assumption that signalling gains depend only on signal level and not signaller quality eliminates the 'fecundity benefit' discussed by Getty (1998). A fecundity benefit would exist if, for example, male signallers contributed more to parental care and, thus, received greater returns (higher offspring survival) for a given mating attempt. In the example discussed here these types of benefits are not used in calculating male fitness, even though females are assumed to receive greater returns for mating with high-utility males. Including a fecundity benefit will not alter the results of this paper in that non-monotonic utility/quality curves will not allow equilibria under a purely perception-based model.

The notion of expected utility introduced here also serves to simplify the analysis of signalling games with perceptual error. Johnstone & Grafen (1992) constructed a model of sexual selection with a monotonic utility/quality curve in order to show that, when perceptual error occurs, increases in signal must be associated with an increased expected perceived signal in order for a signalling equilibrium to exist. This result can be obtained in a more direct fashion by noting that, as long as females prefer higher signals, male signalling will be an increasing function of quality. Given that signalling is increasing in quality, then an unbiased error function (i.e. an increased signal increases the expected perceived signal) guarantees that the expected utility is increasing in quality. Thus, the condition for female preference to be increasing in signal level is maintained and an equilibrium exists.

This analysis can be taken one step further by considering how the expected utility, as a function of signal, is affected by perceptual error when the utility/quality curve is non-monotonic. Perceptual error will smooth out the utility/signal curve, possibly eliminating small peaks or valleys. Thus, if perceptual error causes the expected utility to be a strictly increasing function of signal then a traditional signalling model will have an equilibrium. However, the magnitude of the averaging effect will depend on the magnitude of the differences in signalling between signallers of different quality so that it is difficult to determine general conditions for an equilibrium. The results discussed here apply even when the expected utility is non-monotonic in signal and any smoothing brought about by perceptual error will only make a signalling equilibrium more likely.

Making models of signalling depend on actions is a step towards realism, and the qualitative change in the range of signalling equilibria bespeaks its utility. However, as signalling is often assumed to reduce survivorship, the distribution of signaller phenotypes should depend on the population signalling strategy; to our knowledge, this factor has received very little attention even though it can have qualitative effects on signalling

equilibria (Siller 1998). In games of sexual signalling the inclusion of frequency dynamics can have a substantial quantitative effect on signalling levels (S. R. Proulx, unpublished data). This brings into question results based on quantitative genetic models where frequency distributions are ignored (i.e. Iwasa *et al.* 1991).

The results presented here suggest that a much wider range of conditions can support signalling equilibria than was previously thought. In particular, signalling equilibria can persist even when there is an imperfect relationship between utility and relative signal strength. Previous models of signalling have taken a purely perception-based approach, which produces equilibria that are vulnerable to even minor perturbations in the utility/quality relationship. Considering the costs of receiver strategy as well as benefits yields equilibria even when the relationship between utility and quality is far from monotonic. This has two major implications for field studies of handicap signalling. First, when receivers benefit the most by maximizing an aspect of signaller state, the observation of non-monotonic signalling strategies should not be seen as incompatible with adaptive signalling and, conversely, when receivers value an intermediate signaller state most, adaptive signalling can still exist.

I thank Fred Adler, Janice Ragsdale, Troy Day and Erik Noonburg for helpful comments on this manuscript. Locke Rowe provided useful suggestions on the relationship between quality and utility. Two anonymous reviewers contributed greatly to clarifying this paper and for suggesting that perceptual error can average out bumps in the utility function.

REFERENCES

- Brooks, R. 2000 Negative genetic correlation between male sexual attractiveness and survival. *Nature* **406**, 67–70.
- Candolin, U. 1999 The relationship between signal quality and physical condition: is sexual signalling honest in the three-spined stickleback? *Anim. Behav.* **58**, 1261–1267.
- Dukas, R. 1998 *Cognitive Ecology: The Evolutionary Ecology of Information Processing and Decision Making*. University of Chicago Press.
- Dukas, R. 1999 Costs of memory: ideas and predictions. *J. Theor. Biol.* **197**, 41–50.
- Getty, T. 1998 Handicap signalling: when fecundity and viability do not add up. *Anim. Behav.* **56**, 127–130.
- Grafen, A. 1990 Biological signals as handicaps. *J. Theor. Biol.* **144**, 517–546.
- Hurd, P. L. 1995 Communication in discrete action response games. *J. Theor. Biol.* **174**, 217–222.
- Iwasa, Y., Pomiankowski, A. & Nee, S. 1991 The evolution of costly mate preferences. II. The 'handicap' principle. *Evolution* **45**, 1431–1442.
- Janetos, A. C. 1980 Strategies of female mate choice: a theoretical analysis. *Behav. Ecol. Sociobiol.* **7**, 107–112.
- Johnstone, R. A. 1995 Honest advertisement of multiple qualities using multiple signals. *J. Theor. Biol.* **177**, 87–94.
- Johnstone, R. A. 1998 Conspiratorial whispers and conspicuous displays: games of signal detection. *Evolution* **52**, 1554–1563.
- Johnstone, R. A. & Grafen, A. 1992 Error-prone signalling. *Proc. R. Soc. Lond. B* **248**, 229–233.
- Johnstone, R. A. & Grafen, A. 1993 Dishonesty and the handicap principle. *Anim. Behav.* **46**, 759–764.
- Kawecki, T. J. 2000 The evolution of genetic canalization under fluctuating selection. *Evolution* **54**, 1–12.
- Kokko, H. 1997 Evolutionarily stable strategies of age-dependent sexual advertisement. *Behav. Ecol. Sociobiol.* **41**, 99–107.
- Lachmann, M. & Bergstrom, C. T. 1998 Signalling among relatives. II. Beyond the Tower of Babel. *Theor. Popul. Biol.* **54**, 146–160.
- Lotem, A. 1998 Higher levels of begging behavior by small nestlings: a case of a negatively correlated handicap. *Israel J. Zool.* **44**, 29–45.
- Luttbeg, B. 1994 A comparative Bayes tactic for mate assessment and choice. *Behav. Ecol.* **7**, 451–460.
- Seger, J. 1985 Unifying genetic models for the evolution of female choice. *Evolution* **39**, 1185–1193.
- Siller, S. 1998 The epistatic handicap principle does work. *J. Theor. Biol.* **191**, 141–161.
- Szamado, S. 1999 The validity of the handicap principle in discrete action-response games. *J. Theor. Biol.* **198**, 593–602.
- Zahavi, A. 1975 Mate selection—a selection for a handicap. *J. Theor. Biol.* **53**, 205–214.

As this paper exceeds the maximum length normally permitted, the author has agreed to contribute to production costs.