

Chapter 7

Co-operative and competitive signalling

This chapter has a threefold purpose. Firstly, it introduces a simple signalling game that can be used to model situations such as food and alarm calls, in which one animal informs another about some state of the world. Secondly, it is an attempt to test Krebs and Dawkins's (1984) theory that two kinds of signal co-evolution should be expected in nature (see section 2.5.2): expensive signals resulting from manipulative arms races when participants have conflicting interests, and conspiratorial whispers that evolve when the interests of the participants are congruent. Finally, it is an attempt to position some of the previous artificial-life work on the evolution of communication in a broader theoretical context.

7.1 Background

7.1.1 Explaining food and alarm calls

In many social species, an individual that has discovered a supply of food may, under some circumstances, produce a signal that serves to alert conspecifics to the presence of the resource. For example, chimpanzees *Pan troglodytes*, on discovering a fruit tree, will make loud hooting sounds that attract others (Reynolds & Reynolds, 1965; Sugiyama, 1969). Male domestic chickens *Gallus gallus* give a distinctive call in response to food; they are more likely to produce the call if a hen is present, and the calls attract other chickens (Evans & Marler, 1994). The elaborate dances of bees (von Frisch, 1967) can be considered a particularly sophisticated food signal. Some social animals also produce alarm calls, in which an individual that has detected a predator alerts other group members: the calls of vervet monkeys are an excellent example and have already been much discussed. Alarm calls are also given by other mammalian species (see e.g., Sherman, 1977) and by many birds (Klump & Shalter, 1984; Hauser, 1996). Sometimes alarm calls even serve to recruit conspecifics to mob (i.e., to jointly attack or distract) the approaching predator.

The function of these kinds of signalling systems seems transparent: the signal serves to alert others, and the response of a receiver (i.e., approaching the food or running away) is likely to have positive fitness implications given the context. Barring misidentification, as could occur when what appears to be a food call turns out to be an aspect of sexual advertisement signalling for instance, the adaptive significance of food and alarm calls looks obvious. However, as discussed

in sections 1.2 and 2.3.3, the problem of altruism means that food and alarm call systems are not so easily explained. It is easy to see where the benefit lies for receivers of the signal; being informed of the approach of a predator or the location of food is clearly useful. It is not so easy, however, to determine why the signaller should share the relevant information. In many contexts there will either be no benefit in doing so, or, more likely, costs involved. These costs may be due to, for example, energy expenditure in the production of the signal, an increase in personal risk for the signaller, or the loss of food that might have been consumed alone. There is thus a degree of altruism in such signalling, and a conflict of interests between the signaller and the receiver. With mobbing calls, the altruism runs in the other direction: why should receivers of the signal risk their own lives by assisting in a group attack on the predator?

The problem of accounting for honesty becomes even more acute when we consider communication that occurs with a more explicit conflict of interests between signallers and receivers. For example, in aggressive or territorial signals, each animal would prefer that the other respond by retreating. In many sexual advertisement signals, it is in the interests of the average male to convince any female he meets to copulate with him, but it is in the average female's interests to be difficult to persuade. In these cases and in the apparently co-operative context of alarm and food calls, what prevents the invasion of free-riders who gain the benefit of others' honest signals, but do not pay the costs of honesty themselves? How can honest signalling be an ESS? Furthermore, how might communication have evolved in the first place—why, against an initial background of non-communication, would the first proto-signallers have been selected for their behaviour?

Reciprocal altruism (Trivers, 1971), kin selection (Hamilton, 1964), and the handicap principle (Zahavi, 1975, 1987) are among the mechanisms that have been proposed to explain the evolution of stable, honest signalling, and each of these ideas has spawned a vast literature of its own—particularly that on reciprocal altruism and the Prisoner's Dilemma. However, these three mechanisms will only be treated briefly if at all in this chapter. Our goal is instead to consider a prediction arising from Krebs and Dawkins's (1984) account of animal signalling.

7.1.2 Expensive hype and conspiratorial whispers

Krebs and Dawkins (1984) provide another way of looking at the problem of honesty. As we have seen in section 2.5.2, Krebs and Dawkins challenge the default notion that animal communication is about information transmission; they suggest that propaganda and advertising make better metaphors for animal communication than does the co-operative use of language to share information. They predict two distinct varieties of signal co-evolution. On the one hand there will be evolutionary arms races between manipulative, exploitative signallers and sceptical receivers. This will occur when there is a conflict of interests between the two parties, and the result will be increasingly costly signals. On the other hand, there are some situations in which—to use Krebs and Dawkins's terminology—it is to the receiver's advantage to be manipulated by the signaller. When the two parties share a common interest in this way, there will be selection for signals that are as cheap as possible while still being detectable: "conspiratorial whispers".

The aim of this chapter is to construct a model of food and alarm call situations and then to ascertain whether, given appropriate manipulation of the degree to which the participants have common or conflicting interests, these two types of signal evolution in fact take place. If so, Krebs

and Dawkins's theory may turn out to be a sufficient explanation for "honest" signalling in nature: signalling systems in contexts of common interest are not subject to invasion by dishonest, free-riding mutants, while signalling systems that exist despite conflicting interests are likely to involve much more costly signals and to be ultimately unstable.

In contrast to the handicap principle, few mathematical or simulation models of Krebs and Dawkins's theory have ever been constructed. Presumably, their ideas were accepted without detailed modelling because the argument followed so naturally from the dominant selfish-gene paradigm—models of the handicap principle were constructed because there was fierce debate over whether it would or would not work. In order to test Krebs and Dawkins's prediction, it will first be necessary to determine whether communication should be expected *at all* when signallers and receivers have a genuine conflict of interests.

7.1.3 Putting artificial-life models of communication in perspective

We have seen in section 5.1.3 that previous artificial-life work on the evolution of communication has considered situations in which signallers and receivers have common interests (Werner & Dyer, 1991; MacLennan & Burghardt, 1994), conflicting interests (de Bourcier & Wheeler, 1994; Bullock, 1997b), and intermediate cases in which signallers are ambivalent about the response of receivers (Ackley & Littman, 1994; Oliphant, 1996). A secondary goal of the current chapter is to position this earlier simulation work in an over-arching context. Section 7.1.4 below describes a classification scheme for common and conflicting interests between signallers and receivers; investigating the course of signal evolution across a range of contexts will allow us to incorporate the earlier findings in a unified picture.

While on the subject of previous artificial-life work, it should be noted that the model developed in this chapter postulates a single environmental variable that animals might come to communicate about. This state can take on one of two values, corresponding to, for example, the presence or absence of food. Earlier work—notably MacLennan and Burghardt (1994)—considered "multiple meaning" situations in which a number of environmental states came to be paired up with a number of potential signals. However, if MacLennan and Burghardt's simulation shows us anything, it shows by existence proof that positive payoffs all round for successful communication can transform initially random token-meaning relationships into a workable communication system. The same can be said for work by Steels (1995). The current chapter is limited to the simple one-meaning case in order to more clearly study the effects of different payoff and signal-cost values on the evolution of signalling.

7.1.4 Conflicts of interest

The first requirement in constructing a general model of communication is a classification scheme for determining when a conflict of interests exists between signallers and receivers—Figure 7.1 shows such a scheme, adapted from Hamilton (1964). Assume that a successful instance of communication in a particular scenario has fitness implications for both participants. The fitness effect on signallers, P_S , and the fitness effect on receivers, P_R , together define a point on the plane in Figure 7.1. For example, consider a hypothetical food call, by which one animal alerts another to the presence of a rich but limited food source. By calling and thus sharing the food, the signaller in-




Figure 7.1: Possible communication scenarios classified by their effects on the fitness of each participant.

curs a fitness cost; by responding to the call, the receiver benefits through obtaining food it would otherwise have missed. Thus, the call would be located in the “altruism” quadrant. The situations modelled by Ackley and Littman (1994) and Oliphant (1996)—discussed in section 5.1.3—in which receivers benefit but signallers are ambivalent, can be thought of as points on the positive vertical axis, i.e., where $P_S = 0$ and $P_R > 0$.

Conflicts of interest can be defined as interactions in which natural selection favours different outcomes for each participant (Trivers, 1974), or in which participants place the possible outcomes in a different rank order (Maynard Smith & Harper, 1995). Conflicts of interest therefore exist when P_S and P_R are of opposite sign, i.e., in the upper-left and lower-right quadrants. Selection will, by definition, favour actions that have positive fitness effects. In the upper-left and lower-right quadrants, one individual but not the other will be selected to participate in the communication system: their interests conflict. The “spite” quadrant does *not* represent a conflict of interests because agents will be mutually selected not to communicate.

If the specified fitness effects of participating in a communicative interaction are truly *net* values, and already include such factors as the cost of signalling and the cost of making a response (as well as inclusive fitness considerations and costs due to exploitation of the signal by predators, etc.), then predicting the evolution of the communication system is trivial. Proper signalling requires that it be in the interests of both signallers and receivers for the communication system to exist, and so presumably will only develop when $P_S > 0$ and $P_R > 0$, i.e., when individuals in both roles are selected to participate. However, real animals sometimes appear to communicate despite conflicts of interest, as in signalling during contests (chapter 8) and sexual signalling (chapter 9). Recent models (Grafen, 1990a; Bullock, 1997b) have established that, in certain situations where communication would otherwise be unstable, increasing the production costs of the signal can lead to evolutionarily stable signalling. The costs of signalling (and responding) have therefore been separated from the cost or benefit associated with the outcome of the interaction. In other words, P_S and P_R refer to gross fitness effects before the specific costs of producing the signal, C_S , and making the response, C_R , have been taken into account. Assuming for the sake of the




Figure 7.2: Extended form of the simple signalling game. The shaded cell in each chart icon indexes the relevant payoff value in Table 7.1.

argument that Krebs and Dawkins are correct in predicting two kinds of signal co-evolution, this separation makes it possible to identify the two regimes based on variations in C_S , the cost of signal production.

7.2 A simple signalling game

If the signalling interaction is to involve information transmission, and allow for the possibilities of proper signalling, deception, and manipulation, it must be modelled as a game of imperfect information, in which the signaller knows something that the receiver does not. Figure 7.2 shows the extended form of a simple action-response game that captures the structure of the alarm- or food-call context, and arguably other contexts besides. The game begins with a chance move (the central square) in which some state is randomly determined to be either “high” or “low”. The signaller has access to this state, and we can suppose that it represents either a feature of the environment that only the signaller has detected (e.g., noticing an approaching predator), or a hidden internal state of the signaller (e.g., ovulation). Based on this state, the signaller (player I) must decide whether or not to send an arbitrary signal of cost C_S . The receiver (player II) is ignorant of the hidden state and only knows whether or not a signal was sent—the dashed rectangles show the receiver’s information sets. The receiver can respond either positively, i.e., perform some action “appropriate” to the high state, or negatively, i.e., not respond at all. Positive responses incur a cost, C_R . If and only if the hidden state is high, a positive response results in the payoffs P_S and P_R to the signaller and receiver respectively. Table 7.1 specifies the payoff matrix.

Hurd (1995), Oliphant (1996), and Bullock (1997b) used similar games with different payoff structures. In each of these earlier games, the receiver was explicitly rewarded for accuracy in determining the hidden state. In contrast, in the current game accuracy is not a goal of the receiver *per se*; the receiver simply wants to maximize its average payoff. This is in keeping with Bullock’s point about the information requirements of receivers, discussed in section 2.6.2. Depending on the precise payoff values, the best way to maximize one’s payoff might be to respond in a blanket way, i.e., responding negatively or positively whatever the signal. This is meant to reflect the fact

	State of environment	
	Low	High
No signal		
Neg. response	0 , 0	0 , 0
Pos. response	0 , $-C_R$	P_S , $P_R - C_R$
Signal		
Neg. response	$-C_S$, 0	$-C_S$, 0
Pos. response	$-C_S$, $-C_R$	$P_S - C_S$, $P_R - C_R$

Table 7.1: Payoff matrix for the simple game. Entries in the table represent the payoff to the sender and receiver respectively.

that receivers in natural contexts can presumably opt out of the communication system if it is to their advantage to do so; there is no force compelling them to pay attention to the signaller.

The game models a range of possible communicative interactions. For example, suppose that the high state represents the signaller's discovery of food. Sending a signal might involve emitting a characteristic sound, while not sending a signal is to remain silent. For the receiver, a positive response means approaching the signaller and sharing the food, whereas a negative response means doing nothing. Various possibilities exist besides honest signalling of the high state: the receiver might *always* approach the signaller in the hope of obtaining food, regardless of whether a signal was sent. The signaller might be uninformative and never signal, or only signal when food was *not* present. One important feature of the game is that the signaller is ambivalent about the receiver's response in the low state—in terms of the example, this represents the assumption that when no food has been discovered, the signalling animal does not care about whether the receiver approaches or not.

The strategies favoured at any one time will depend on the relative values of P_S , P_R , C_S and C_R , as well as on what the other members of the population are doing. (Another parameter of interest is the relative frequency of high and low states; in the models presented here each state occurred 50% of the time.) Allowing the base fitness effects P_S and P_R to vary across positive and negative values will allow the payoff space of Figure 7.1 to be explored, and thus determine whether changes in signal and response cost can produce stable signalling in situations that would otherwise involve conflicts of interest. Note that in the simple game, there is no potential for signals of varying costs, and thus no room for costly signalling arms races. Variable-cost signalling will be considered later on in the chapter; this initial game is only a first step towards assessing Krebs and Dawkins's conspiratorial whispers theory.

7.2.1 Stable strategies in the simple game

A signalling strategy in the simple game specifies whether to respond with no signal (NS) or a signal (Sig) to low and high states respectively. Likewise, a response strategy specifies whether to respond negatively (Neg) or positively (Pos) when faced with no signal and when faced with a signal. A complete strategy is the conjunction of a signalling and a response strategy; e.g.,

(NS/NS, Pos/Pos) is the strategy that specifies never signalling and always responding positively.

The strategy (NS/Sig, Neg/Pos) specifies signalling only in the high state, and responding positively only to signals—call this the “honest and trusting” strategy. Evolutionary stability depends on a strategy being the best response to itself; i.e., a strategy must be uninvadable in order to be an ESS. Honest and trusting players meeting each other can expect an average payoff per interaction of:

$$\frac{P_S - C_S + P_R - C_R}{4}$$

This will be higher than the expected payoff for any possible invading strategy (i.e., honesty and trust will be an ESS) if:

$$\begin{aligned} P_S &> C_S > 0 \\ P_R &> C_R > 0. \end{aligned}$$

That is, honest signalling is stable if the costs of signalling and responding are both positive, and if the payoffs in each case outweigh the costs. The requirement that P_S and P_R must both be positive means that the honest strategy is only expected to be stable when the interests of the parties do not conflict: positive values of P_S and P_R place the interaction in the upper right “mutualism” quadrant of Figure 7.1. For the derivations of these results and others presented in this chapter, the reader is referred to appendix A.

Of the 16 possible strategies, there are three besides the honest strategy that involve the transmission of information, in that the receiver responds differently to different hidden states. None of these three are ESSs if C_S and C_R are both positive; these two values represent energetic costs and so cannot sensibly be negative. If $C_S = 0$, i.e., if giving a signal is of negligible cost, then the reverse honesty strategy (Sig/NS, Pos/Neg) can be stable, although P_S and P_R must still be positive. It is also worth noting that a population consisting entirely of individuals playing (NS/NS, Pos/Pos) or (NS/NS, Pos/Neg), both non-signalling strategies where the receiver always responds positively, cannot be invaded by any other strategy if the payoff to the receiver is large enough, i.e., if:

$$\begin{aligned} C_S &> 0 \\ P_S &> -C_S \\ P_R &> 2C_R > 0. \end{aligned}$$

The analysis indicates that while the cost of signalling plays some role in stabilizing the honest strategy, there are no circumstances in which stable communication is predicted when a conflict of interests exists. This is despite the fact that we have separated the costs of signalling and responding from the base fitness payoffs of a communicative interaction.

7.2.2 Evolutionary simulation model

An evolutionary simulation model of the simple game was also constructed in order to determine whether communicative behaviour might sometimes be found outside the range of identified ESSs. A straightforward genetic algorithm (GA) was used. Each individual could play both signalling and receiving roles; a strategy pair was specified by a four-bit genotype as shown in table 7.2.

	Bit value	
	0	1
If low state...	No signal	Signal
If high state...	No signal	Signal
Response to no signal	Negative	Positive
Response to signal	Negative	Positive

Table 7.2: Genetic specification of strategies.

The population size was 100, the mutation rate was 0.01 per locus, and, due to the trivially small genome, crossover was not used. Each generation, 500 games were played between randomly selected opponents. An individual could therefore expect to play 5 games as a signaller and 5 as a receiver. The fitness score was the total payoff from these games. For breeding purposes, the fitness scores were normalized by subtracting the minimum score from each. Proportionate selection was then applied to the normalized scores. The genetic algorithm was run in this manner for 500 generations. In the results presented below, the games played in the final, i.e., 500th, generation have been used as a snapshot of the evolved signalling strategies.

An attempt was made to investigate evolutionary dynamics, in that the initial populations were not determined randomly but started as either “honest” or “non-signalling”. Honest initial populations were made up entirely of individuals who played the honest and trusting strategy, i.e., a genome of “0101”. Non-signalling populations underwent 100 generations of preliminary evolution in which their receiving strategies were free to evolve but their signalling strategies were clamped at “00”, i.e., no signalling. For each class of initial conditions, a simulation run was performed for all combinations of integer values of P_S and P_R between -5 and +5, making 121 runs in all. Each run was repeated 25 times with different random seeds. The values of C_S and C_R were fixed at 1.

Communication was indexed by cross-tabulating the hidden state value with the receiver’s response and calculating a chi-squared statistic. The receiver has no direct access to the hidden state, so any reliable correspondence between state and response indicates that information has been transmitted and acted upon. Values of the χ^2 statistic close to zero indicate no communication, and values close to the maximum (in this case $\chi_{max}^2 = 500$, due to the 500 games played in the final, snapshot generation) indicate near-perfect communication.

Figure 7.3 shows the average values of the communication index for honest initial conditions. Seeding the population with honesty tests the stability of honest signalling given a particular payoff pair, much as a game-theoretic analysis does. The results are compatible with the conditions outlined in the previous section: honesty is stable when the payoffs to signaller and receiver are positive and greater than their respective costs. However, there is some suggestion of intermittent or imperfect communication when $P_R = C_R = 1$, indicating that ambivalent receivers may occasionally co-operate.

Figure 7.4 shows the average values of the communication index for non-signalling initial conditions. Starting the GA with a non-signalling population tests the likelihood that communication




Figure 7.3: Mean communication index by P_S and P_R ; honest initial conditions. Each point is a mean calculated over 25 runs. Mean standard error = 2.96.

will emerge, given a particular payoff pair. Clearly the conditions for emergence and stability-once-present are not the same. If $P_S > 1$ and $P_R = 2$ communication develops but when $P_S > 1$ and $P_R > 2$ it does not.

In the latter region $P_R > 2C_R$ and the population remains at the non-signalling equilibrium described in section 7.2.1. Despite the fact that communication would result in a higher average fitness, the high value of P_R keeps the receivers responding positively all the time, removing any incentive for the signallers to bother signalling. This response strategy could be called “blind optimism”, as receivers always respond positively. It should be noted, however, that the condition $P_R > 2C_R$ is dependent on the 50% frequency of high states; if high states occurred 10% of the time, for instance, then $P_R > 10C_R$ would be required to make blind optimism a stable strategy.

The difference in results between the two classes of initial conditions is interesting, but should not obscure the fact that no communication was observed under conditions of conflicting interests. We must conclude that, at least in the simple model discussed so far, stable communication is only to be expected when it is in the interests of both parties.

7.3 A game with variable signalling costs

In the simple signalling game, signallers can choose between a costly signal or no signal at all. The model does not allow for a range of possible signals with differing costs, and in this respect it is unrealistic. It may be that Krebs and Dawkins’s implicit prediction, that signalling can occur when a conflict of interests exists, is in fact true, but can only be demonstrated in a more complex game with a range of signal costs. The simple signalling game (see Figure 7.2) was therefore extended to incorporate signals of differing costs.

7.3.1 Stable strategies in the variable-signal-cost game

In the extended game, the signalling player has three options: not signalling, which costs nothing; using the “soft” signal, which costs C_S , and using the “loud” signal, which costs $2C_S$. Strategies in the extended game require specifying the signal to give when the hidden state is low, the signal




Figure 7.4: Mean communication index by P_S and P_R ; non-signalling initial conditions. Each point is a mean calculated over 25 runs. Mean standard error = 1.67. Graph rotated for clarity—co-operative quadrant appears at top left.

to give when it is high, and the response to give to each of no-signal, soft and loud. The two strategies representing conspiratorial whispers or cheap signalling are (NS/Soft, Neg/Pos/Pos) and (NS/Soft, Neg/Pos/Neg). Both strategies call for the soft signal to be used in the high state, and for positive responses to the soft signal; the strategies differ only in the response to loud signals. Neither of these strategies can strictly be considered an ESS on its own (because neutral drift can take the population from one to the other) but it is shown in appendix A that the set of all mixed strategies involving these two is an ESS under the familiar conditions:

$$P_S > C_S > 0$$

$$P_R > C_R > 0.$$

Costly signalling would involve the use of the loud signal for the high state, and either the soft signal or no signal to denote the low state, with a corresponding response strategy. None of the four strategies in this category can be an ESS. For example, (NS/Loud, Neg/Pos/Pos) cannot be an ESS assuming positive costs of signalling and responding. The similar strategy (NS/Loud, Neg/Neg/Pos) is almost stable if $P_S > 2C_S$, but can drift back to the previous strategy which can in turn be invaded by the cheap strategy (NS/Soft, Neg/Pos/Pos).

Analysis of the extended game indicates that if signalling is favoured at all, then at equilibrium the signallers will always use the cheapest and the second-cheapest signal available (i.e., no signal and the soft signal). Extending the game by adding ever more costly signalling options, until we have approximated a continuous range of signal costs, does not alter this conclusion. None of the costly signalling strategies can even be an ESS, let alone support communication in the face of a conflict of interests. The possibility of expensive signalling arms races starts to look remote. However, it may be that an evolutionary simulation model will reveal signalling strategies that, while unstable in the long term, nevertheless lead to transient communication under conditions of conflicting interest.

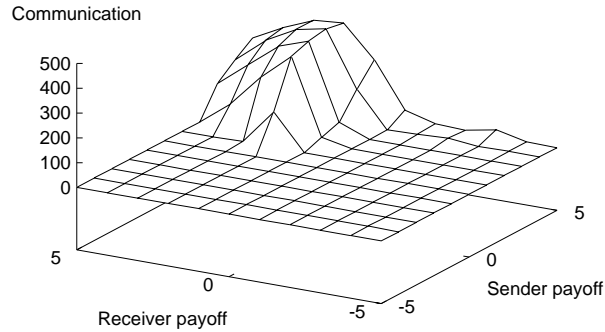


Figure 7.5: Mean communication index by P_S and P_R in the continuous simulation; honest initial conditions. Each point is a mean calculated over 25 runs. Mean standard error = 4.22. Graph rotated for clarity—co-operative quadrant appears at top.

7.3.2 Evolutionary simulation model

A second evolutionary simulation was constructed, in which the cost of signalling was continuously variable. Signalling strategies were represented by two positive real numbers C_{low} and C_{high} : the cost of the signals given in the low state and in the high state respectively. Response strategies were represented by a real-valued threshold T ; positive responses were given to signals with costs greater than the receiver’s threshold value. Note that threshold value could be negative, indicating a positive response to any signal.

A real-valued GA was used to simulate the evolution of strategies over time. Generally, the same parameters were used as in the previous simulation model, e.g., a population of 100. Mutation was necessarily a different matter: each real-valued gene in each newborn individual was always perturbed by a random gaussian value, $\mu = 0$, $\sigma = 0.05$. If a perturbation resulted in a negative cost value the result was replaced by zero. In addition, 1% of the time (i.e., a mutation rate of 0.01) a gene would be randomly set to a value between 0 and 5 for signal costs, or between -5 and +5 for the threshold value. This two-part mutation regime ensured that offspring were always slightly different from their parent, and occasionally very different.

The C_S parameter was no longer relevant, but C_R , the cost of responding, remained fixed at 1. Honest initial conditions were implemented by setting $C_{low} = 0$, $C_{high} = 1.0$ and $T = 0.5$. Non-signalling initial conditions were implemented by setting T to a random gaussian ($\mu = 0$, $\sigma = 1$) and then clamping $C_{low} = C_{high} = 0$ for 100 generations of preliminary evolution.

Figures 7.5 and 7.6 show the average values of the communication index for honest and non-signalling initial conditions respectively. The results are qualitatively similar to those of the discrete simulation model: communication occurs in both cases, but in a more limited range of the payoff space for non-signalling conditions. In neither case does communication occur outside the “co-operative” quadrant.

However, there is some evidence that transient communication can occur when the conflict of interests between the two agents is not too extreme. For example, consider the payoff pair




Figure 7.6: Mean communication index by P_S and P_R in the continuous simulation; non-signalling initial conditions. Each point is a mean calculated over 25 runs. Mean standard error = 3.61. Graph rotated for clarity—co-operative quadrant appears at top left.

$P_S = 5$ and $P_R = 0$. This defines a point on the boundary between mutualism and selfishness, although when the constant cost of responding ($C_R = 1$) is taken into account, the net payoffs indicate that communication under these circumstances would be selfish (from the point of view of the signaller). Nevertheless, as Figure 7.7 shows, unstable communication evolves, even from non-signalling initial conditions.

The continuous model also allows investigation of the cost and threshold values over the payoff space. C_{low} , the cost of the signal given in response to the low state, always remained close to zero—this was unsurprising as signallers are ambivalent about the receiver’s response to the low state. However, the value of C_{high} varied both inside and outside the region where communication was established: Figure 7.8 shows the mean values of C_{high} for honest initial conditions. The signals given in response to the high state are most costly when P_S , the payoff to the sender, is high and when the receiver’s net payoff is marginal, i.e., $P_R \approx 1$. In order to study this effect more closely, additional simulation runs were performed, with P_S fixed at 5 and P_R varied between -5 and +5 in increments of 0.1. These runs can be thought of as exploring the cross section through $P_S = 5$ in Figure 7.8. Figure 7.9 shows the cross-sectional mean values of C_{high} . Note that the “energy” devoted to signalling is at a maximum around $P_R = 1$ and drops off as P_R increases—it can be seen from Figure 7.5 that $P_R = 1$ is approximately the point where significant communication is established. The same pattern was observed for non-signalling initial conditions (not shown for reasons of space).

The threshold values showed corresponding variation. Figure 7.10 shows the mean value of T across the payoff space. The threshold values are typically very high (a “never respond” strategy) or very low (an “always respond” strategy), but in the region where communication evolved, receivers become progressively less demanding, i.e., T gets lower, as P_R increases. Figure 7.11 shows the cross-sectional results for $P_S = 5$.

Figure 7.12 plots the mean cost of high and low signals and the mean reception threshold all on one graph. This makes the relationship between costs and threshold clear: at approximately




Figure 7.7: Mean communication index plotted over generational time. A typical run with $P_S = 5$, $P_R = 0$ and non-signalling initial conditions.




Figure 7.8: Mean cost of high-state signals by P_S and P_R ; honest initial conditions. Each point is a mean calculated over 25 runs. Mean standard error = 0.032. Graph rotated for clarity—co-operative quadrant appears at top.




Figure 7.9: Cross-sectional means (± 1 s.e.) for high-state signal costs with $P_S = 5$; honest initial conditions. Each point is a mean calculated over 25 runs.




Figure 7.10: Mean threshold value by P_S and P_R ; honest initial conditions. Each point is a mean calculated over 25 runs. Mean standard error = 0.19. Graph rotated for clarity—co-operative quadrant appears at top left.




Figure 7.11: Cross-sectional mean threshold values (± 1 s.e.) with $P_S = 5$; honest initial conditions. Each point is a mean calculated over 25 runs.




Figure 7.12: Cross-sectional means: cost of high and low signals, and reception threshold. $P_S = 5$, honest initial conditions. Each point is a mean calculated over 25 runs.

$P_R = 1$, the threshold falls to a level where the mean high-state signal will generate a positive response. As P_R increases, i.e., as the two players' payoffs approach each other, the signallers become less extravagant and the receivers less "sceptical". This is *contra* the game-theoretic result of the previous section, which implies that when signals of varying costs are available, either the cheapest pair of signals will be used, or no signalling will occur—something like Figure 7.13 would be expected if the soft-loud signalling game accurately modelled the continuous case.

Note that the initial values of C_{high} and T under honest initial conditions were 1.0 and 0.5 respectively. For all but the highest values of P_R , C_{high} has increased on average over the 500-generation run. This rules out any explanation of the results of Figure 7.12 in terms of there having been insufficient evolutionary time for a cheaper signalling equilibrium to have been reached when the profit for receivers ($P_R - C_R$) was marginal. Evolution has taken the populations *away* from the cheap signalling solution.




Figure 7.13: Approximate predicted results for Figure 7.12 according to discrete-cost game-theoretic model.

7.3.3 Discussion

In all of the models presented so far, stable communication evolved or was predicted to evolve only within the co-operative region of the signaller-receiver payoff space. This means that no signalling at all (costly or otherwise) was observed when the signaller and the receiver were experiencing a conflict of interests, apart from transitory communication on the boundaries of the co-operative region as shown in Figure 7.7.

The second game-theoretic model, in which discrete signals of varying costs are available, suggests that communication, if selected for, will involve the cheapest pair of signals available. However, the second simulation model, incorporating the more realistic assumption that signals can vary continuously in cost, implies that cheap signals will only be used when both parties stand to gain a high payoff from effective communication. When the net payoff to the receiver is marginal, evolved signals will be more costly than strictly necessary to convey the information. The relationship is not symmetrical: when the net payoff to the signaller is marginal, a non-signalling equilibrium, in which the receiver always responds positively, is likely to occur.

Krebs and Dawkins (1984) predicted that signalling would be costly if a conflict of interests existed; strictly speaking the results do not support nor contradict their prediction, as no signalling occurred in the conflict-of-interest cases. It might be the case that conflicts of interest in the context of a different signalling game would indeed result in costly signals. However, the failure to evolve communication given conflicts of interest in this simple game strongly suggests that in many natural contexts (e.g., food calls, alarm calls) reliable signalling should not be expected unless it is in the interests of both parties. This conclusion is not altered by separate consideration of the specific costs of producing a signal and of making an appropriate response to that signal.

The results from the second simulation model do not confirm Krebs and Dawkins's conspiratorial whispers theory, but they definitely suggest a modification of it. As Figure 7.12 shows, when the net payoff to the receiver is marginal, receivers will be sceptical and express "sales-resistance" by responding only to costly signals; signallers in turn will be prepared to invest more energy in "convincing" receivers to respond positively. When communication is unambiguously good for

both parties, signals are cheaper and response thresholds lower. Therefore both expensive hype and conspiratorial whispers are expected to evolve, but in a much smaller region of the payoff space than Krebs and Dawkins's theory suggests, i.e., within the co-operative region. Expensive hype is what happens when honest signalling is highly profitable to the signaller, but only marginally so to the receiver. For example, suppose that a juvenile benefits by honestly signalling extreme hunger to its parent, because the parent responds by feeding it. If the net inclusive-fitness payoff to the parent is only slight, perhaps because the parent is the ostensible father and the species has a high ratio of extra-pair copulations, then costly signals by the juvenile are expected. Thus the model predicts that chicks should beg more loudly to their fathers than to their mothers, for instance.

7.4 Variations on the continuous-signal-cost game

In line with the reasoning presented in section 6.3, a number of variations of the evolutionary simulation model with continuous signal and threshold values will now be presented. In order to avoid any further profusion of graphs, the variants will incorporate only non-signalling initial conditions. Rather than requiring the reader to constantly compare each figure with Figure 7.6—the mean communication index data for the continuous-signal-cost game with non-signalling initial conditions—the communication index results in each variant will be presented as *differences* from that graph. That is, Figure 7.6 will be used as a reference level of communication; positive results for a variant will indicate a greater relative level of communication and not an absolute measure.

7.4.1 Noise and uncertainty

The use of continuous values for the cost of signals and for the response threshold suggests the possibility of random noise in the signalling channel. In the real world signals will not always be accurately perceived, and Johnstone (1994) found that modelling noise or perceptual error in a signalling game in fact altered the predictions about which strategies were expected to be stable. It was thought that perhaps the inclusion of noise would alter the region of the payoff space in which communication evolved.

Noise was implemented by adding a random gaussian value ($\mu = 0$) to the energy level of the signal before it was perceived by the receiver. Thus, signals will sometimes be heard as “louder” or “softer” than they in fact are. When the random gaussian value had a standard deviation of 0.2, noise made very little difference to the communication index data, i.e., communication evolved much as in Figure 7.6. When the standard deviation was set to 2.0, on the other hand, communication was entirely disrupted. Presumably intermediate levels of noise would have led to a progressive degradation of communication. However, there was no evidence that the addition of noise could lead to honest signalling in regions of the payoff space where it would otherwise not have occurred.

Randomness was also applied to the payoff values P_S and P_R in order to investigate the effects of realistic uncertainty. The payoff values, as in all game-theoretic accounts, are intended to be average expected payoffs. However, computer simulation allows us to assign payoffs in a particular interaction that are drawn from a random gaussian distribution. Thus the long term mean will be as specified, e.g., $\overline{P_S} = 2$ and $\overline{P_R} = 2$, but the rewards for successful communication in any one game will be somewhat unpredictable. When the standard deviation of the random gaussian was 0.2, the




Figure 7.14: Difference in mean communication index between uncertain payoff variant ($\sigma = 2.0$) and standard continuous-signal-cost game; non-signalling initial conditions. Each point is the difference between two means, each calculated over 25 runs. Graph rotated for clarity—co-operative quadrant appears at top right.

evolution of stable communication was unaffected. When the standard deviation was increased to 2.0, communication started to degrade as shown in Figure 7.14. However, there was again no suggestion that the modelling of uncertainty in payoff values could lead to communication where it would not have otherwise evolved.

7.4.2 Exploitation of sensory biases and mutational lag

The simple games and simulations described here are in one sense an unfair way to test Krebs and Dawkins's (1984) conspiratorial whispers hypothesis. Krebs and Dawkins discuss the likely evolution of signals in complex real-world cases, and can therefore appeal to the exploitation of response patterns that had originally been selected for other purposes, the effects of differing mutation rates in signallers and receivers, etc. Communication in their predicted costly signalling arms races was not necessarily expected to be stable. For example, in a real-world situation where it was not in the interests of receivers to respond positively to a particular signal from a predator, they might nevertheless continue to do so for some time if the signal was structurally similar to a mating signal made by members of the same species. The manipulative signalling system would break down as soon as an appropriate sequence of mutations resulted in organisms that could distinguish between the predator's signal and the conspecific mating signal. In the signalling models presented all this complexity is abstracted into the base fitness payoffs for signallers and receivers.

In an attempt to investigate these issues, two simple modifications were made to the standard continuous-signal-cost game. In the first of these, we suppose that the receivers have some other ecological reason for having a low threshold value, e.g., that the same sensory mechanisms are involved in food detection. This opens up an opportunity for signallers to exploit a “sensory bias” (Guilford & Dawkins, 1991; Ryan & Rand, 1993) in the receivers. Selection pressure for low thresholds (T) was implemented by giving receivers in each game an energy bonus (b) as follows:




Figure 7.15: Difference in mean communication index between sensory bias variant and standard continuous-signal-cost game; non-signalling initial conditions. Each point is the difference between two means, each calculated over 25 runs. Graph rotated for clarity—co-operative quadrant appears at top left.

$$b = \begin{cases} 0 & \text{if } T > 5 \\ 1 & \text{if } T < 0 \\ \frac{5-T}{5} & \text{if } 0 \leq T \leq 5 \end{cases}$$

The results of simulation runs of this variant are shown in Figure 7.15 (using Figure 7.6 as a baseline). When receivers have other reasons for maintaining a low response threshold, communication evolves much more reliably in the usual co-operative region of the payoff space, and also occurs in the selfish region. That is, signallers are able to manipulate receivers to their own (the signallers’) advantage. Furthermore, as predicted by Krebs and Dawkins (1984), the most costly signals indeed occurred when communication had been established despite a conflict of interests.

In another variant, it is supposed that response strategies might evolve more slowly than signalling strategies, i.e., there is a mutational lag on reception thresholds relative to signal cost values. Such a state of affairs could come about in the real world if the sensory equipment used to detect signals was older and affected by a larger network of genes than the organs used for signalling. It would then be possible that signallers might “out-evolve” receivers, and succeed in getting them to respond to selfish, manipulative signals. The idea was implemented by reducing both of the mutation rates for reception thresholds by a factor of 10. That is, the real-valued threshold gene in a newborn individual was perturbed by a random gaussian value, $\mu = 0$, $\sigma = 0.005$, and 0.1% of the time (i.e., a mutation rate of 0.001) a completely new threshold value was generated in the range ± 5 . The results are shown in Figure 7.16.

As with the sensory bias variant, communication is established more strongly in part of the co-operative region, but it also evolves in the selfish region for high values of P_S . Again, the most costly signals were also found when selfish communication had evolved. A puzzling feature of the result is that it does not appear to have come about simply because the low rate of mutation for threshold values meant that 500 generations was insufficient time for the optimal value to be reached. Mean threshold values when $P_S = 5$ and $P_R < 0$ were approximately 4 in both the




Figure 7.16: Difference in mean communication index between mutational lag variant and standard continuous-signal-cost game; non-signalling initial conditions. Each point is the difference between two means, each calculated over 25 runs. Graph rotated for clarity—co-operative quadrant appears at top left.

mutational lag variant and the original simulation data.

7.4.3 The effects of spatial arrangement

Ackley and Littman (1994) and Oliphant (1996) both found that arranging signalling populations in space led to a greater degree of altruistic signalling. In Ackley and Littman’s model individuals lived in small groups, communicating and breeding only with their group-mates, but occasionally migrating to another nearby group. There was no spatial arrangement within each group, but the groups themselves were laid out on a grid. In Oliphant’s model individuals were arranged in a ring, and were likely to communicate and to breed with their neighbours.

A spatial variant was implemented by arranging the population of 100 individuals on a toroidal 10×10 grid. Individuals interacted only with their 8 neighbours: in each game, a signaller was chosen at random from the population and a receiver was chosen at random from the signaller’s neighbours. Breeding was also local. When one generation replaced another, the parent of the individual who would occupy a particular square was chosen, using roulette-wheel selection¹ according to fitness, from among the nine local candidates from the previous generation. That is, the parent of the occupant of a given square would either be the previous occupant or one of the previous occupant’s neighbours. The results for the spatial variant are shown in Figure 7.17.

Arranging the population in space leads to an increase in the reliability of communication, but only in that section of the co-operative region where honesty has already been observed to evolve. The agents have clearly not been induced to participate in altruistic communication with their neighbours. There is no communication even when signallers are merely ambivalent ($P_S = 0$). However, it can be shown that altruism of a sort has occurred. Figure 7.18 shows the difference in

¹Roulette-wheel selection refers to a process whereby any one individual’s probability of being selected is proportional to its fitness score. The probabilities of selection can be envisaged as sectors of varying size on a roulette wheel.




Figure 7.17: Difference in mean communication index between spatial variant and standard continuous-signal-cost game; non-signalling initial conditions. Each point is the difference between two means, each calculated over 25 runs. Graph rotated for clarity—co-operative quadrant appears at top left.

mean fitness between the spatial variant and the original simulation. There is a spike of increased fitness in the altruistic quadrant at the front of the graph: this occurs because *receivers* are refraining from constant positive responses, and thus being altruistic towards their signalling neighbours who would be penalized by a positive response because of the negative value of P_S in this area.

7.4.4 Insistent signallers

The signalling game used is not likely to be a universal model of all possible communicative interactions. In particular, and despite having the same basic structure with two signals possibly used to transmit information about a binary hidden state, the signalling game is different from those employed by Hurd (1995), Oliphant (1996) and Bullock (1997b). Hurd’s game, for instance, models sexual signalling, and the male signaller is *not* ambivalent about the female receiver’s response when the hidden state is low; the signaller always prefers a positive response. A low hidden state maps to low male quality, a positive response represents a copulative episode, and even low-quality males want mating opportunities. The current signalling game, in contrast, cannot model so-called “handicap” signalling, because low-state signallers do not care about what the receiver does. Furthermore, in previous games, receivers are explicitly rewarded for accuracy in discerning the hidden state, but the game presented here allows the ecologically plausible outcome that receivers simply become disinterested in the signal. The current game is a reasonable model of situations such as alarm calls and food calls, in which potential signallers have no reason to care about what receivers do when no predator has been sighted or no food source has been found. Whereas Hurd’s game serves as a (discrete) model of situations where signallers vary on some dimension, the current game models situations where signallers fall into two groups, only one of which is relevant to the potential response.

However, it is a simple matter to alter the present game such that signallers are always interested in gaining a positive response. The payoff matrix is altered such that P_S , the payoff to




Figure 7.18: Difference in mean fitness between spatial variant and standard continuous-signal-cost game; non-signalling initial conditions. Each point is the difference between two means, each calculated over 25 runs. Graph rotated for clarity—co-operative quadrant appears at top left.

the signaller, is awarded whenever the receiver responds positively, regardless of the value of the hidden state. On the other hand, receivers are still only awarded their payoff, P_R , when they respond positively and the hidden state is high. There is thus a different kind of conflict of interests between the signaller and receiver.

Making signallers want positive replies all the time in this way almost completely breaks down communication—see Figure 7.19. There are no circumstances in which receivers can trust signallers, and extreme response strategies (always responding positively or always responding negatively) are formulated purely on the basis of the payoff to the receiver. Interestingly, communication can be salvaged if the conditions of the handicap principle are applied: that is, if the unit cost of giving a signal in the low state is greater than for the high state. The results for a run in which signals in the low state cost 5 times their normal value are shown in Figure 7.20; relative to the standard game, communication levels are only somewhat degraded.

7.5 General discussion

The results from simulations of the simple and continuous-cost signalling games suggest that communication will not evolve when there is a conflict of interests between signallers and receivers. Even when signallers and receivers share a common interest, the evolution of communication is not straightforward. Firstly, receivers may fall into blindly optimistic strategies (i.e., always responding positively) that are less efficient than the communicative equilibrium but nevertheless stable. This is particularly likely to occur when the net payoff to the receiver is high. (The expected payoff for always responding positively will of course depend on the relative frequency of high and low hidden states, a factor that was not varied in the models presented). Secondly, communication may evolve but the signals involved will be more or less costly depending on the marginal payoff of the receiver, as discussed in section 7.3.3.

Variations on the continuous-cost signalling game, while only briefly explored, suggest that




Figure 7.19: Difference in mean communication index between insistent signallers variant and standard continuous-signal-cost game; non-signalling initial conditions. Each point is the difference between two means, each calculated over 25 runs. Graph rotated for clarity—co-operative quadrant appears at top left.

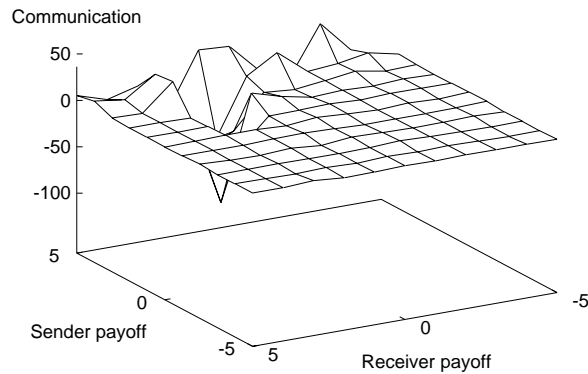


Figure 7.20: Difference in mean communication index between handicap principle variant and standard continuous-signal-cost game; non-signalling initial conditions. Each point is the difference between two means, each calculated over 25 runs. Graph rotated for clarity—co-operative quadrant appears at top left.

communication can in fact evolve under conditions of conflicting interest if receivers have a sensory bias that maintains low response thresholds, or if response strategies do not evolve as quickly as signalling strategies. In these two cases, manipulative or selfish communication can occur. Of course, in the case of a sensory bias communication that evolves is not *really* occurring under a conflict of interests, because receivers are choosing the strategy that maximizes their two sources of fitness: the communication game and the independent response bias. However, an observer unaware of the receivers' response bias would observe agents responding to signals in a way that was not in their immediate interests.

Altruistic communication (considered from the point of view of signallers) was not observed under any circumstances, including the spatial variant simulation. Spatial arrangement of the population would seem not to be a guarantee of kin-selected altruism. The occurrence of apparently altruistic food and alarm calls in nature, in circumstances where reciprocal altruism and kin selection cannot be invoked, therefore remains to be explained. In other words, this model alone cannot tell us why an alarm-calling monkey resists the temptation to quietly slip away and save itself; if an empirical study was to show that some animal gives alarm calls to non-relatives without hope of reciprocation, then we would have a genuine conundrum on our hands.

However, the model may be a step towards understanding the evolution of a different kind of non-kin, non-reciprocal altruism. Mobbing calls seem to involve a benefit for the signaller, who recruits allies to help drive off a predator, and a cost for receivers, who sustain a risk of being injured in the attack. Mobbing calls would therefore be classified as selfish under the scheme presented in Figure 7.1. In the sensory bias and mutational lag variations, this sort of selfish communication was in fact observed. From the point of view of *receivers*, this represents altruism directed towards the signaller. It might be the case that some manipulative mobbing calls are maintained despite a real cost to those who respond, because, for example, the call-production behaviour can evolve faster than the ability to distinguish between the calls of relatives and non-relatives.

The evolutionary simulation models presented are unusual in their use of non-random initial conditions. The use of non-signalling initial conditions in particular can be seen as an attempt to get at the origin or emergence of communication rather than just studying the conditions for its stability, as does orthodox game theory. Non-signalling initial conditions embody the assumption that communication must emerge from a non-communicative context—the un-clamping of signalling strategies after a period of preliminary evolution can be seen as the introduction of a mutation that allows the *possibility* of signalling. The simulation results have certainly demonstrated that the conditions for stability can be very different from those for emergence.

A final qualification must be made concerning the results: the way that conflicting and congruent interests have been defined may be too simplistic. In the simple signalling game, it is true that with positive net payoffs to the signaller and the receiver, and if the hidden state is high, both agents will benefit from a positive response, and they therefore have congruent interests. However, if we consider the moment before the hidden state has been determined, it is not clear whether the interests of the two agents conflict or not. If the signaller, for example, could somehow choose the strategy of its opponent, the receiver, it would want the opponent to play an “always respond positively” strategy—that way the signaller would always receive the payoff and would not have to ex-

pend energy in signalling. However, the receiver, if similarly allowed to determine the signaller's strategy, would prefer that the signaller used an honest strategy, precisely so that the receiver could avoid the costs of responding positively to the low hidden state. Recall that Trivers (1974) defined a conflict of interests as an interaction in which natural selection favours a different outcome for each participant. It seems that the signaller and receiver in this situation favour different strategies in their opponent, and thus have a conflict of interests, even though a high value of the hidden state would mean that their interests became congruent. If this strategy-based definition of conflicting interests were adopted, any situation in the co-operative payoff region, assuming signalling had a positive cost, would involve a conflict of interests—this would in turn mean that *all* of the signalling observed in the simulation models evolved despite a conflict of interests. The problem is perhaps that Trivers's (1974) and Maynard Smith and Harper's (1995) definitions are not specific enough about just what constitutes an “outcome” of the signalling game. The simpler definition of conflicting interests, as used in the body of the paper, is useful in isolating the co-operative region of payoff space as the place to expect signalling. It is not yet clear how the results should be interpreted if the strategy-based definition of conflicting interests was pursued.