

Co-operation, Competition and the Evolution of Pre-Linguistic Communication

Jason Noble

Center for Adaptive Behavior and Cognition
Max-Planck-Institut für Bildungsforschung
Lentzeallee 94, D-14195 Berlin, GERMANY
Email: noble@mpib-berlin.mpg.de

1 Language origins and Darwinian thought

Theories of the origin of language are necessarily speculative. Calvin (1983) suggests that the development of language involved a transfer of the skills involved in stone-throwing; Knight (1998) puts the roots of language in ritual; Bickerton (1998) argues that language arose from proto-language in a single catastrophic mutation. Any one of these accounts might be true, but it is difficult or impossible to gather direct evidence that would allow us to decide between them. An unkind observer might conclude that anything goes, and that one foundation myth is as good as another.

However, such cynicism would be misplaced. In recent years the range of acceptable speculation has been greatly narrowed by the recognition that any account of language origins must be consistent with the principles of evolution by natural selection. For instance, modern Darwinism tells us that complex traits do not evolve without their having some function, that *all* of the intermediate stages in the evolution of modern linguistic capacity must themselves have had adaptive value, and that gradual development is more plausible than catastrophic change. These sorts of constraints immediately rule out many stories of language origin, such as the suggestion by Gould (1987) that language is a mere by-product of having a large and complex brain.

The chief problem for a Darwinian account of human speech, however, is the apparent level of altruism involved. The orthodox position in evolutionary biology (Dawkins, 1976) suggests that organisms are best understood as products of their selfish genes: they do not do things for the good of the group or the species, but in order to propagate copies of their own genetic material. Given this perspective, speech (and many other forms of co-operative behaviour) can be difficult to account for. Why do speakers freely exchange valuable information when the theory of natural selection predicts selfishness? In a hypothetical proto-linguistic community, what would prevent the rise of a selfish mutant strain that listened but did not speak? Speaking or signalling always costs something in terms of time and energy, and may involve more indirect costs such as exposing the signaller to greater predation risk. Why not reap the benefits of the informative signals of others, without paying the costs of signalling oneself? Or worse, why not use the communication system to lie, misinforming others for one's own benefit?

Possible answers to this dilemma are usually phrased in terms of kin selection (Hamilton, 1964) or reciprocal altruism (Trivers, 1971): speakers or signallers remain honest because they are helping their relatives or because they want others to do the same for them, respectively. There are alternative explanations: Dessalles (1998, this volume) presents

the intriguing suggestion that honest information is given freely because it is a way of competing for status within the group. Knight (1998) argues that the co-operative exchange of information that characterizes speech involves a great risk of deception, and therefore that speech-like communication could only be evolutionarily stable if there was some mechanism that made it strategically sound to trust other members of the group. Knight believes that this mechanism is ritual; group members demonstrate their allegiance to the common cause by performing a costly ritual act, and this allows the rest of the group to believe their potentially fakeable signals in future.

Knight's argument relies in part on a view of communication presented in the behavioural ecology literature by Krebs and Dawkins (1984). Krebs and Dawkins do not define animal communication in terms of information transmission but as a method whereby one animal exploits the muscle power of another. They outline two possibilities for the co-evolution of signalling and response behaviour (see section 3): one that leads to costly, manipulative signals, and another that leads to quiet, efficient and honest communication. The latter occurs when it is in the interest of both animals that the signaller successfully "manipulate" the receiver; in other words, co-operative signalling. Knight argues that human speech is the sort of system that one would expect to have resulted from the second process, and that motivates his hypothesis that ritual was the key to creating the necessary co-operative context.

2 Simulating the evolution of communication

Interesting as Knight's work is, the goal of this chapter is to explore not his theory but that of Krebs and Dawkins (1984), using game theory and computer simulations of evolution. The sceptical reader may need convincing, however, that Krebs and Dawkins's ideas are relevant to the evolution of language. It is true that their work is most easily applied to simple animal signalling systems, and it is certainly true that the communication systems presented here (in simulation) will be much simpler than language. Nevertheless, Krebs and Dawkins's theory is important and relevant because it forces us to recognize the Darwinian truth that animals, including ourselves, must be expected to be manipulative rather than informative, all things being equal. This fact must be constantly borne in mind in trying to account for the anomalous levels of altruism in speech. Furthermore, modelling simple pre-linguistic communication is useful because it puts further constraints on theories of how language itself evolved—as things stand there is room for far too many plausible possibilities.

Mathematical and simulation modelling are necessary steps if we are to go beyond an impasse in which the proponents of competing theories merely trade rhetoric. Formal models can produce counter-intuitive results, and show, for instance, that of two apparently plausible theories only one is internally consistent. An excellent example of the value of a good model in theorizing about communication is the story of the handicap principle. This idea was introduced by Zahavi (1975), who proposed that signallers sacrifice some of their fitness (i.e., impose a handicap on themselves) in order to produce signals that will be believed by receivers. When the handicap principle was first introduced, it was generally not accepted by theoretical biologists. Simple population-genetic models seemed to show that it could not be evolutionarily stable. However, an elaborate mathematical model developed by Grafen (1990) appears to have vindicated Zahavi's idea and has made the handicap principle a respectable explanatory construct.

Simulation models of the evolution of communication have been put forward before, but have rarely considered the general case that is implied by Krebs and Dawkins's theory: the possibility that different kinds of communication may evolve under conditions of conflict and of co-operation. Earlier models have often been constructed such that honest signalling

was always in the interests of both signallers and receivers. Thus, only co-operative communication systems could possibly emerge. For example, Werner and Dyer (1991) postulated blind, mobile males and sighted, immobile females: the evolution of a signalling system was in the interests of both parties as it allowed mating to take place at better-than-chance frequencies. In MacLennan and Burghardt's (1994) model, signallers and receivers were rewarded if and only if they engaged in successful communicative interactions.

Other models (Ackley & Littman, 1994; Oliphant, 1996) have looked at the special case where communication would benefit receivers, but the potential signallers are indifferent. Oliphant argues that this is a good way to model the evolution of alarm calls, for example: if one bird in a flock spots an approaching hawk, it is clear that its conspecifics would benefit from an alarm call. However, why should the bird in question, considered as a product of its selfish genes, give the call?

Finally, some simulation models have considered the evolution of communication in situations where the two parties appear to have conflicting interests. Wheeler and de Bourcier (1995) modelled aggressive territorial signalling. Bullock (1997) constructed a general model in which signallers of varying degrees of quality solicited receivers for a favourable response; receivers were rewarded for responding positively only to high-quality signallers. A secondary goal of the current chapter is to try to position this earlier simulation work in an overarching theoretical context.

3 Expensive hype and conspiratorial whispers

Krebs and Dawkins (1984) view signalling as a competitive affair involving mind-reading and manipulation. Mind-reading consists of one animal exploiting tell-tale predictors about the future behaviour of another, e.g., a dog noticing the bared teeth of an opponent, concluding that it is about to attack, and fleeing in order to avoid injury. Manipulation is what happens when those being mind-read fight back, influencing the behaviour of the mind-readers to their own advantage. For example, a dog could bare its teeth despite not having the strength or inclination to attack, and thus scare off its mind-reading opponent. The authors predict evolutionary arms races between manipulative signallers and sceptical receivers: "selection will act simultaneously to increase the power of manipulators *and* to increase resistance to it" (p. 390). The result will be increasingly costly, exaggerated signals; examples from nature include the roars of red deer stags and the elaborate tails of peacocks.

Krebs and Dawkins admit, however, that not all interactions are competitive in nature. There are some situations in which it is to the receiver's advantage to be manipulated by the signaller. For example, a pack-hunting predator may attempt to recruit a conspecific in order to bring down prey too large for either to tackle alone. Foraging bees, on returning to the hive, may indicate to their closely related hive-mates the direction and distance to a source of nectar. In these cases the receiver's compliance is to the benefit of both parties, i.e., there exists the possibility of co-operation. Krebs and Dawkins argue that when the two parties share a common interest in this way, then a different kind of signal co-evolution will result. Specifically, there will be selection for signals that are as energetically cheap as possible while still being detectable; Krebs and Dawkins suggest the phrase "conspiratorial whispers" to describe these signals. Rather than signallers needing to be more and more extravagant in their attempts to persuade receivers, the opposite process occurs: receivers are eager to be persuaded, and selection will favour subtle signalling and low response thresholds.

Krebs and Dawkins's argument has been influential but no formal justification of it exists. The models presented here will test their prediction that evolved signals will necessarily be more costly when there is a conflict of interests than when the two parties have common




Figure 1: Possible communication scenarios classified by their effects on the fitness of each participant.

interests. In order to do so, it will first be necessary to determine whether communication should be expected *at all* when signallers and receivers have a genuine conflict of interests.

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 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 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 incurs 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), where 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 animal 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 the two parties will each be 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. Reliable communication requires, on average, honest signallers and trusting receivers, and thus will only develop when $P_S > 0$ and $P_R > 0$, i.e., when both participants are selected to participate. However, real animals sometimes communicate despite apparent conflicts of interest (Hinde, 1981). Recent models (Grafen, 1990; Bullock, 1997) have established




Figure 2: Extended form of the simple signalling game. The shaded cell in each chart icon indexes the relevant payoff value in Table 1.

that, in certain situations where communication would otherwise be unstable, increasing the production costs of the signal can lead to a prediction of evolutionarily stable signalling. Therefore, in the current model, 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.

5 A simple signalling game

If the signalling interaction is to involve information transmission, and allow for the possibilities of 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 2 shows the extended form of a simple action-response game that fulfils this requirement. 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 some feature of the environment that only the signaller has detected, e.g., noticing an approaching predator. 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 1 specifies the payoff matrix. Hurd (1995), Oliphant (1996), and Bullock (1997) used similar games with different payoff structures.

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

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

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 in the signalling game 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 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. This will be a first step towards assessing Krebs and Dawkins's conspiratorial whispers theory.

6 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 strategy pair is the conjunction of a signalling and a response strategy; e.g., (NS/NS, Pos/Pos) is the strategy pair that specifies never signalling and always responding positively.

The strategy pair (NS/Sig, Neg/Pos) specifies signalling only in the high state, and responding positively only to signals—call this the honest strategy. It can be shown that honesty will be an evolutionarily stable strategy (ESS; Maynard Smith 1982) 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 in co-operative contexts.

Of the 16 possible strategy pairs, 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 strategy pairs 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 any

	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 2: Genetic specification of strategies.

mixed strategy involving (NS/NS, Pos/Pos) and (NS/NS, Pos/Neg), both non-signalling strategies where the receiver always responds positively, can be an ESS 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 An evolutionary simulation model

Game theory is limited to describing equilibria; 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; see Mitchell (1996) for an introduction to this technique. Each individual could play both signalling and receiving roles, and a strategy pair was specified by a four-bit genotype as shown in table 2. The population size was 100, the mutation rate was 0.01 per locus, and 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 individual's fitness score was the total payoff from these games. For breeding purposes, the fitness scores were normalized by subtracting the minimum score from each, and 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.

In order to see how communication might arise from a non-communicative context, the initial population for the genetic algorithm was not randomly generated (as is usually done), but was constructed in such a way that no communication occurred. 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. 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




Figure 3: Mean communication index by P_S and P_R . Each point is a mean calculated over 25 runs.

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 $\chi^2_{max} = 500$, due to the 500 games played in the final, snapshot generation) indicate near-perfect communication.

Figure 3 shows the average values of the communication index over the repeated simulations; the co-operative quadrant is at the left rear of the graph. Clearly the conditions for the *stability* of the honest strategy, as established by the game-theoretic analysis in section 6, are not the same as the conditions established here for the evolution of honest signalling from non-communicative beginnings. 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 ESS described earlier: although communication would result in a higher average fitness, the high value of P_R keeps the receivers responding positively at all times, removing any incentive for the signallers to bother signalling. Note that under no circumstances does stable communication occur when there is any conflict of interest between the two parties, i.e., outside the co-operative quadrant.

8 A game with variable signal 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 2) was therefore extended to incorporate signals of differing costs.

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 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 can be shown that the set of all mixed strategies involving these two is an ESS under the familiar conditions:

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

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). Further extensions of the game, by adding ever more costly signalling options, do 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 the discrete signals used in the games presented so far have had an unwarranted effect on the results.

9 Simulation model with continuous signal costs and reception threshold

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 values 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. 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.

Figure 4 shows the average values of the communication index. The results are qualitatively similar to those for the discrete simulation model: communication occurs in a limited region of the co-operative quadrant, and never outside it.

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




Figure 4: Mean communication index by P_S and P_R in the continuous simulation. Each point is a mean calculated over 25 runs.




Figure 5: Mean cost of high-state signals by P_S and P_R . Each point is a mean calculated over 25 runs.

outside the region where communication was established: Figure 5 shows the mean values of C_{high} . 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 5. Figure 6 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 4 that $P_R = 1$ is approximately the point where significant communication is established.

The threshold values showed corresponding variation. Figure 7 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 8 shows the cross-sectional results for $P_S = 5$.

Figure 9 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 $P_R = 1$, 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




Figure 6: Cross-sectional means (± 1 s.e.) for high-state signal costs with $P_S = 5$. Each point is a mean calculated over 25 runs.




Figure 7: Mean threshold value by P_S and P_R . Each point is a mean calculated over 25 runs.




Figure 8: Cross-sectional mean threshold values (± 1 s.e.) with $P_S = 5$. Each point is a mean calculated over 25 runs.




Figure 9: Cross-sectional means: cost of high and low signals, and reception threshold. $P_S = 5$. Each point is a mean calculated over 25 runs.




Figure 10: Approximate predicted results for Figure 9 according to discrete-cost game-theoretic model.

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 10 would be expected if the soft-loud signalling game accurately modelled the continuous case.

10 Discussion of results

In all of the models presented, 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. 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. However, although the results from the second simulation model do not confirm Krebs and Dawkins’s conspiratorial whispers theory, they definitely suggest a modification of it. As Figure 9 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 costly manipulative signals 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. Costly signals evolve when honest signalling is highly profitable to the signaller, but only marginally so to the receiver. For example, if a juvenile benefits by honestly signalling extreme hunger to its parent (because the parent responds by feeding it), but the net inclusive-fitness payoff to the parent is only slight, then costly signals by the juvenile are expected.

There are two qualifications that must be made concerning the results. Firstly, 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 potentially used to transmit information about a binary hidden state, the signalling game is different from those employed by Hurd (1995) and Oliphant (1996). Hurd’s game 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 both Hurd’s and Oliphant’s 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.

Secondly, it must be stressed that 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 were discussing the likely evolution of signals in complex real-world cases, and could therefore appeal to the effects of differing mutation rates in signallers and receivers, and the exploitation of behaviours that had originally been selected for other purposes, etc. Communication in the predicted costly signalling arms races was not 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 simple signalling model all this complexity is abstracted into the base fitness payoffs for signallers and receivers, and there is no guarantee that any transient, unstable evolved communication systems will be detected.

11 Implications for theories of language evolution

Where does all this leave Knight (1998) and others who wish to use Krebs and Dawkins's ideas as part of the foundations of a theory of language evolution? The simulations seem to show that the costly-arms-race/conspiratorial-whispers theory is simply not correct, at least not without modification. However, that is only to say that communication is not expected to evolve under conditions of conflicting interest in a simple action-response game. The suspicion that it *would* have evolved (of which the author is manifestly guilty) can now be put down to careless interpretation of Krebs and Dawkins's talk about the possibilities of "manipulative signals". Thus we are reminded of the value of formal modelling: when considering Krebs and Dawkins's verbal argument, it is easy to come away with the impression that communication will readily occur given a conflict of interests (and will involve high signal costs). The simulation models demonstrate the falsity of that impression.

Nothing has been established as to the success of Krebs and Dawkins's theory in more complex scenarios, however. For example, a single communication system may be subject to contemporaneous co-operative and competitive usage, e.g., when social animals use the same signal repertoire to communicate with both in-group and out-group conspecifics, as Knight (1998) suggested. Dessalles (1998, this volume) describes a scenario in which competition for social status provides a new currency which can offset the costs of signal production. In such cases, costly signalling between agents with conflicting interests might well be evolutionarily stable—the issue is appropriate for further simulation modelling.

The immediate implications of the results presented should therefore not be over-stated. Nevertheless, even such simple models start to put constraints on theories of proto-language and language evolution. Given the results from the second simulation model, and supposing one suspects that language originated in the food and alarm calls of early hominids, then one has to establish that the balance of co-operative payoffs would have allowed communication to evolve. If one's theory of language evolution requires low-cost signalling, then the "payoff window" will be even narrower. It is through exploring these sorts of constraints that our theories about the evolution of simple signalling systems will eventually connect up with our theories about language.

Acknowledgements

This work was conducted while at the University of Sussex, and I am grateful to the Association of Commonwealth Universities and the British Council for financial support. I would also like to thank Seth Bullock and Ezequiel Di Paolo for valuable discussions.

References

- Ackley, D. H., & Littman, M. L. (1994). Altruism in the evolution of communication. In Brooks, R. A., & Maes, P. (Eds.), *Artificial Life IV*, pp. 40–48. MIT Press, Cambridge, MA.
- Bickerton, D. (1998). Catastrophic evolution: The case for a single step from protolanguage to full human language. In Hurford, J. R., Studdert-Kennedy, M., & Knight, C. (Eds.), *Approaches to the Evolution of Language: Social and Cognitive Bases*, pp. 341–358. Cambridge University Press, Cambridge.
- Bullock, S. (1997). An exploration of signalling behaviour by both analytic and simulation means for both discrete and continuous models. In Husbands, P., & Harvey, I. (Eds.),

- Proceedings of the Fourth European Conference on Artificial Life (ECAL'97)*, pp. 454–463. MIT Press / Bradford Books, Cambridge, MA.
- Calvin, W. H. (1983). A stone's throw and its launch window: Timing precision and its implications for language and hominid brains. *Journal of Theoretical Biology*, 104, 121–135.
- Dawkins, R. (1976). *The Selfish Gene*. Oxford University Press, Oxford.
- Dessalles, J.-L. (1998). Altruism, status, and the origin of relevance. In Hurford, J. R., Studdert-Kennedy, M., & Knight, C. (Eds.), *Approaches to the Evolution of Language: Social and Cognitive Bases*, pp. 130–147. Cambridge University Press, Cambridge.
- Gould, S. J. (1987). The limits of adaptation: Is language a spandrel of the human brain?. Unpublished paper delivered to the Center for Cognitive Science, MIT.
- Grafen, A. (1990). Biological signals as handicaps. *Journal of Theoretical Biology*, 144, 517–546.
- Hamilton, W. D. (1964). The genetical evolution of social behaviour. *Journal of Theoretical Biology*, 7, 1–52.
- Hinde, R. A. (1981). Animal signals: Ethological and games-theory approaches are not incompatible. *Animal Behaviour*, 29, 535–542.
- Hurd, P. L. (1995). Communication in discrete action-response games. *Journal of Theoretical Biology*, 174, 217–222.
- Knight, C. (1998). Ritual/speech coevolution: A solution to the problem of deception. In Hurford, J. R., Studdert-Kennedy, M., & Knight, C. (Eds.), *Approaches to the Evolution of Language: Social and Cognitive Bases*, pp. 68–91. Cambridge University Press, Cambridge.
- Krebs, J. R., & Dawkins, R. (1984). Animal signals: Mind reading and manipulation. In Krebs, J. R., & Davies, N. B. (Eds.), *Behavioural Ecology: An Evolutionary Approach* (Second edition), pp. 380–402. Blackwell, Oxford.
- MacLennan, B. J., & Burghardt, G. M. (1994). Synthetic ethology and the evolution of cooperative communication. *Adaptive Behavior*, 2(2), 161–188.
- Maynard Smith, J. (1982). *Evolution and the Theory of Games*. Cambridge University Press, Cambridge.
- Maynard Smith, J., & Harper, D. G. C. (1995). Animal signals: Models and terminology. *Journal of Theoretical Biology*, 177, 305–311.
- Mitchell, M. (1996). *An Introduction to Genetic Algorithms*. MIT Press / Bradford Books, Cambridge, MA.
- Oliphant, M. (1996). The dilemma of Saussurean communication. *BioSystems*, 37, 31–38.
- Trivers, R. L. (1971). The evolution of reciprocal altruism. *Quarterly Review of Biology*, 46, 35–57.
- Trivers, R. L. (1974). Parent-offspring conflict. *American Zoologist*, 14, 249–264.

- Werner, G. M., & Dyer, M. G. (1991). Evolution of communication in artificial organisms. In Langton, C. G., Taylor, C., Farmer, J. D., & Rasmussen, S. (Eds.), *Artificial Life II*. Addison-Wesley, Redwood City, CA.
- Wheeler, M., & de Bourcier, P. (1995). How not to murder your neighbor: Using synthetic behavioral ecology to study aggressive signaling. *Adaptive Behavior*, 3(3), 273–309.
- Zahavi, A. (1975). Mate selection—A selection for a handicap. *Journal of Theoretical Biology*, 53, 205–214.