

The evolution of signal form: Effects of learned versus inherited recognition*

Masashi Kamo¹, Stefano Ghirlanda², and Magnus Enquist²

¹Department of Biology, Graduate School of Sciences, Kyushu University, Higashi-ku,
Hakozaki 6-10-1, Fukuoka, Japan

²Division of Ethology, Department of Zoology, Stockholm University, Krättriket 7B
S-106 91, Stockholm, Sweden

Reprint of December 3, 2006

Abstract

Organisms can learn by individual experience to recognize relevant stimuli in the environment or they can genetically inherit this ability from their parents. Here, we ask how these two modes of acquisition affect signal evolution, focusing in particular on the exaggeration and cost of signals. We argue first, that faster learning by individual receivers cannot be a driving force for the evolution of exaggerated and costly signals unless signal senders are related or the same receiver and sender meet repeatedly. We argue instead that biases in receivers' recognition mechanisms can promote the evolution of costly exaggeration in signals. We provide support for this hypothesis by simulating coevolution between senders and receivers, using artificial neural networks as a model of receivers' recognition mechanisms. We analyse the joint effects of receiver biases, signal cost and mode of acquisition, investigating the circumstances under which learned recognition gives rise to more exaggerated signals than inherited recognition. We conclude the paper by discussing the relevance of our results to a number of biological scenarios.

*First published in *Proceedings of The Royal Society* **B269**, 1765–1771 (2002). © The Royal Society. Permission to reproduce and distribute in any media is granted provided notice of first publication is retained, no fees are requested and no alterations are made. Correspondance: Masashi Kamo, kamo@biomath-10.biology.kyushu-u.ac.jp.

1 Introduction

The evolution of signal form has recently received a lot of attention (reviewed by Bradbury & Vehrencamp, 1998; Enquist & Arak, 1998; Ryan, 1998). Determinants of form, such as the coding and transmission of information, receiver biases and strategic factors, have been considered. Here, we consider another potential determinant of signal form, namely how recognition of the signal is acquired by receivers. We distinguish between genetically inherited recognition and recognition learned by individual receivers, based on their own experiences with signal senders. In the former case, receivers are born with knowledge obtained during their species' evolutionary history; in the latter, receivers are born naive and will make mistakes while learning the appropriate reactions to the senders' signals.

Both inherited and learned recognition occur in nature (Hogan & Bolhuis, 1994). For example, female frogs recognize the call of conspecific males without any specific learning (Blair, 1964; Salthe & Mecham, 1974). Similarly, snakes recognize food without learning (Arnold, 1981). Examples of where learning is important include mate recognition in birds (ten Cate, 1994) and predators learning to avoid unpalatable prey (Edmunds, 1974). Why this diversity exists is an important issue, but is not the topic of this paper. Rather, we ask whether individually learned and genetically inherited recognition have different consequences for the evolution of signal form. In particular, we are interested in the degree of exaggeration and cost of the signal, noting that exaggeration often entails a bigger cost for senders (either costs in producing the signal or other costs, such as easier detection by predators).

An important factor for the evolution of costly signals is conflict between individuals (Dawkins & Krebs, 1978; Arak & Enquist, 1995). Typically, conflict is due to individuals preferring different courses of action (in gametheoretical analyses, such conflicts can be seen in the payoff matrix of a game). However, in this paper we assume that senders and receivers both benefit from receivers accurately recognizing senders. Nevertheless, conflict can occur due to a lack of knowledge in receivers, causing them to respond in ways that are suboptimal for both themselves and senders. For example, a naive bird may attack a distasteful bug. Once a receiver has learned to recognize senders, the conflict diminishes. Note that, in the case of learned recognition, the conflict reappears with every individual born to the population, whereas in the case of inherited recognition, conflict disappears once recognition has been successfully coded in the genes.

This last remark points out that signal evolution may be affected by whether signal recognition is inherited or individually learned, but gives little clue as to

what the difference might be. To date, most theoretical studies of signalling seem to assume genetically inherited recognition (e.g. Grafen, 1990; Bradbury & Vehrencamp, 1998). The only signalling context, to our knowledge, in which the role of learning has been extensively studied is the evolution of aposematic coloration (Leimar et al., 1986; Guilford, 1990; Mallet & Joron, 1999). In this context, it is sometimes said that unprofitable prey may benefit from using an exaggerated signal because receivers will learn more quickly, thus making fewer mistakes that are costly for both prey and predator.

However, when a sender and a receiver meet, the sender's fitness depends in many cases on the receiver's response only, not on what the receiver learns from the encounter. If a would-be aposematic prey is killed by a predator, for example, it does not matter to the dead prey whether its signal has been effectively learned or not. More generally, learning by receivers has fitness consequences for senders only if the sender and receiver meet repeatedly or the receiver meets relatives to the sender (Luce & Raiffa, 1957; Hamilton, 1964). In this paper, we ignore these possibilities to study the simpler case of unrelated senders and single encounters between senders and receivers. Under such circumstances, exaggerated signals cannot evolve simply because they are easier to learn. Rather, there must be a direct benefit for a mutant bearing an exaggerated signal, compared with the rest of the sender population that employs a less exaggerated one.

Biases in receiver recognition mechanisms can potentially provide such a benefit (Leimar et al., 1986; Guilford, 1990). For example, if receivers learn to respond to a stimulus, an exaggerated form of the stimulus (e.g. bigger, louder or brighter) may elicit a stronger response (Hinde, 1970; Mackintosh, 1974). It will be advantageous for senders to exploit such biases only to the extent that the exaggerated signal is better recognized by receivers and does not carry too high a cost. As receivers will eventually recognize signals without mistakes if recognition is inherited, we may conjecture that senders may benefit from exaggerated, costlier signals mainly in the case of learned recognition. In the following, we examine theoretically the coevolution between senders and receivers in the case of both inherited and learned recognition, considering both cost-free and costly exaggeration. To assess the role of receiver biases, we use artificial neural networks as the receiver recognition mechanism. These models are known to exhibit biases similar to those of animals under many circumstances (Ghirlanda & Enquist, 1998; Kamo et al., 1998; Phelps & Ryan, 1998; Kamo & Iwasa, 2000; Phelps, 2001).

2 Model

2.1 Receivers and senders

As a model of receivers, we use feed-forward artificial neural networks (see Haykin, 1994; Ghirlanda & Enquist, 1998, for a comprehensive presentation of artificial neural networks and a study of their response properties in biologically relevant settings, respectively). Here, we use a network with three interconnected layers. The first layer contains 10 input units, whose activation transmits to five units in the middle layer by means of weights (model synapses) that can amplify or attenuate the transmitted signal. In turn, activation of middle-layer units reaches one output unit by a second layer of connections. The activation of the output unit (a number between zero and unity) is the response of the network to the initial stimulation.

A sender is described by specifying the effect of its signal, s , on the receiver's input units. A signal is thus a set of 10 numbers between zero and unity.

2.2 Fitness

Receiver fitness is determined by the extent to which it discriminates between the sender signal, s , and a second stimulus, b , modelling background stimulation or other stimuli which should be ignored. The optimal receiver responses to s and b are assumed to be 0.9 and 0.1, respectively. Formally, receiver fitness is calculated as follows based on the absolute differences e_b and e_s (errors) between the actual and optimal receiver responses to b and s :

$$f_R = \sqrt{(1 - e_b)(1 - e_s)} \quad (1)$$

Note that f_R reaches its maximum value of 1 when both errors e_b and e_s are zero.

Sender fitness is determined by receiver responses to s and signal cost. We assume that the the same value 0.9 of responding to s is optimal for both senders and receivers. To calculate signal cost, we first define an exaggerated signal as one eliciting a strong response from receiver sense organs. More precisely the exaggeration $x(s)$ of signal s is calculated as

$$x(s) = \frac{1}{n} \sum_{k=1}^n s_k^2 \quad (2)$$

where s_k is the value of signal unit k and $n = 10$ is the number of signal units. Exaggeration so defined varies between 0 (if all signal units are 0) and 1 (if all

signal units are 1), and we further assume that signal cost be proportional to signal exaggeration. The full expression for sender fitness may thus be written as:

$$f_S = 1 - e_s - cx(s) \quad (3)$$

where c is a non-negative number regulating signal cost. The term $1 - e_s$ translates the assumption that accurate recognition by receivers (small e_s) is beneficial to senders as well, while the term $-cx(s)$ is the negative contribution of exaggeration-induced signal cost. In the following we will analyse the evolution of both cost-free signals ($c = 0$) and costly signals with different values of c .

2.3 Learning

To model learning by individual receivers, we use the well-known back-propagation algorithm (LeCun, 1985; Parker, 1985; Rumelhart et al., 1986; Haykin, 1994). This is an iterative procedure whereby, at each iteration, network weights are slightly modified to reduce the difference between the actual output and the optimal response to each stimulus. More iterations of the algorithm and a bigger difference between s and b both result in better approximation of the optimal responses (1). The number of back-propagation iterations can thus be taken to correspond to the amount of experience that receivers have with the relevant stimuli. As a model of biological learning, back-propagation is not fully realistic (e.g. McLaren, 1989; Mitchison, 1989). However, in studies of signal evolution we are not interested in the learning process per se. It is enough that receivers, after experiences with a given sender signal, react realistically to mutant signals: in other words, we need realistic models of generalization. Previous studies have ascertained that feed-forward neural networks, learning by back-propagation, generalize similarly to real animals (Ghirlanda & Enquist, 1998; Phelps & Ryan, 1998; Phelps, 2001).

2.4 Simulations

2.4.1 Inherited recognition

For simplicity, we consider fairly large populations, and we assume that favourable mutations arise at a sufficiently low rate as to allow a successful mutant to invade the population before a fitter one arises. Under such conditions, each population may be assumed to be dominated by one phenotype, called the resident. During each generation, a mutant arises, and its fitness is compared with that of the




Figure 1: Learning curves as a function of the amount of experience (number of back-propagation iterations) and signal exaggeration. The background b is a uniform stimulus of value 0.02, and s is a uniform stimulus set at: solid line, 1.0; short-dashed line, 0.5; solid and dotted line, 0.2; and solid and crossed line, 0.1. Long-dashed lines represent optimal responses to s and b (0.9 and 0.1, respectively). Note that discrimination becomes more accurate when the amount of learning increases and when the difference between the two stimuli increases.

resident phenotype. The phenotype with higher fitness is retained in the population. Furthermore, because mutants are single individuals in a large population, the fitness of both the mutant and the resident phenotypes in one population are calculated based on the interaction with only the resident phenotype in the other population (see Enquist & Arak, 1993, 1994; Kamo et al., 1998, for further details and applications).

The background stimulus b is fixed throughout the course of evolution; all of its units are set at 0.02. Units of s are set at 0.02 (the same as b) at the beginning of each simulation and change by mutation by adding a small random number from a

Gaussian distribution ($m = 0$ and $s = 0.05$) to a few signal units, with the constraint that all unit values stay between zero and unity. On average, one unit is mutated when generating a mutant. Receivers mutate by adding a random number from a Gaussian distribution ($m = 0$ and $s = 0.01$) to several of their network weights. An average of two weights are changed each time a mutant is generated.

2.4.2 Learned recognition

We also examine coevolution where receivers do not inherit any information from the previous generation. In this case, the receiver population does not evolve and it can be modelled simply as a single network whose weights are set at random at the beginning of each generation. Senders evolve in the same way as with inherited recognition. As mutants are rare, we assume that receivers learn for some time based on the resident sender alone, without experience of the mutant. This is modelled by a number of back-propagation iterations where the network learns to discriminate s from b , to an extent that depends on both the characteristics of s and the number of learning iterations (cf. 1). After a number of experiences with the resident sender, receivers will meet the mutant. The fitness of the latter, relative to the resident, can now be assessed by means of 3. That is, we check whether, following initial experiences with the resident sender, the mutant is recognized more accurately than the resident and how much it pays for its signal. The sender with the higher fitness becomes the resident in the next generation (the transitional period where, in an actual population, both mutant and resident types would be abundant can be ignored since the network will prefer the same sender throughout).

3 Results

Generally, simulations reached a steady state after a number of generations, which is a condition in which the observed quantities did not fluctuate appreciably. The evolutionary trajectories are shown in 2. Replications with the same simulation parameters yielded almost the same steady-state values. In the case of learned recognition, the steady state was reached within less than 1000 generations (due to evolution of the sender signal). In the case of inherited recognition, between 10000 and 30,000 generations were needed. Optimal receiver responding was attained in the case of inherited recognition, both with cost-free and costly signals (recall that the same receiver response is optimal for both senders and receivers). In the case of learned recognition and cost-free signals, extreme exag-




Figure 2: Coevolution between senders and receivers. Evolutionary trajectories of exaggeration (solid and dotted line), receiver response to s^- (solid and crossed line), receiver response to s^+ (dashed line) and sender fitness (solid line) are shown for (a) inherited recognition versus non-costly signals, (b) learned recognition versus non-costly signals, (c) inherited recognition versus costly signals, and (d) learned recognition versus costly signals. Long-dashed lines represent optimal responses. We used $c=1$ for costly signals (see 3) and 200 back-propagation iterations for learned recognition. All values are the average of 10 replicates.

generation evolved, as it pays mutant senders to exploit even small biases for exaggerated signals in receivers. That such biases exist is known from previous studies (Ghirlanda & Enquist, 1998) and is also shown below. By contrast, when exaggeration is costly, there is a trade-off between the gain yielded by exaggerating the signal and the loss caused by its greater cost.

The most intriguing result is that senders evolve a more exaggerated, costlier signal when recognition in receivers is learned rather than inherited (cf. 2 c,d). To study this effect in fuller detail, we have run a set of simulations where we varied




Figure 3: Signal exaggeration evolved in simulations with learned recognition as a function of amount of experience (number of back-propagation iterations). Cost parameter, c , is 1. The dashed line corresponds to inherited recognition. Bars on both lines denote s.e.

the amount of experience of receivers with the resident sender, before encountering the mutant. In these simulations, the most exaggerated and costly signals were obtained for intermediate amounts of experience. The figure shows a case where intermediate levels of exaggeration evolve; more extreme values are observed with less costly exaggeration (see 5 below).

To understand this result, we have explored receiver biases during the evolutionary process as follows. At each generation, we start with the signal that has evolved so far and try to improve it by mutation. That is, we generate a new signal and test whether it elicits from receivers a response that is closer to the optimal. We repeat this step 1000 times, each time mutating the signal that has so far elicited the response closest to the optimal. We call the signal thus obtained the “best” signal. Note that evolution is not affected because receivers are not learning about the test signals. In figure 4a, we show the receivers’ output to the best and to the evolved signal as a function of the amount of experience in each generation. We see that senders would not gain much by switching to the best signal when receivers learn either a little or a lot (receivers show little bias). At intermediate amounts of learning, receivers exhibit a larger bias and switching to the best signal would yield bigger fitness gains to mutant senders. In Figure 4b, we also see that the best signal is considerably more exaggerated (costlier) than the evolved one, especially for small amounts of experience. This further inhibits the evolution of exaggeration when receivers learn little.




Figure 4: Receiver biases as a function of amount of experience. Cost parameter, c , is 1. (a) Comparison of receiver output to the signal evolved by sender (open circles) and the “best” signal (filled circles) (whose output is closest to target, see text). Dashed line, s target. (b) Exaggeration of evolved (open circles) and best (filled circles) signals. Each data point is the average of 10 replicates. The standard error in the exaggeration of the best signal varied between 0.064 and 0.003.

Finally, we have explored the joint effect of signal cost and amount of learning (5). The results can be predicted, given the above discussion. When exaggeration is cheap (small c in 3), exaggerated signals evolve easily because it is advantageous to exploit even small receiver biases. When exaggeration is costly, exaggeration evolves only for intermediate amounts of learning.




Figure 5: (a) Signal exaggeration as a function of signal cost and amount of experience by receivers. More exaggerated signals are indicated as lighter areas. For example, black corresponds to $0 < x(s) \leq 0.1$ and white to $0.9 \leq x(s) \leq 1$. (b) Sender fitness as a function of signal cost and length of learning phase. Colour coding as in (a).

4 Discussion

Exaggerated, costly signals may evolve for a number of reasons. For instance, signals may be exaggerated in order to get through to the receiver. Alternatively, the cost may be crucial for maintaining the communication of reliable information (Zahavi, 1975; Grafen, 1990). Exaggerated signals may also evolve in an arms race between senders and receivers, in which the sender tries to manipulate the receiver and the receiver tries to resist being manipulated (Dawkins & Krebs, 1978; Arak & Enquist, 1995; Enquist & Arak, 1998). This requires that there is an evolutionary conflict between senders and receivers (Arak & Enquist, 1995).

In this paper, we have studied another possible explanation for exaggeration in cases where the sender and receiver both benefit from the same receiver response. When signal recognition is learned rather than genetically inherited, our simulations show that costlier exaggeration may evolve. The reason for this is that mutant senders bearing a more exaggerated signal may elicit a response closer to the optimal from receivers. This is a consequence of the generalization gradient emerging during learning. In our simulations, the amount of exaggeration that actually evolves is limited by its cost: when the cost was removed, extreme exaggeration evolved both under learned and inherited recognition. This strategic

scenario agrees with the results from a different type of model that was applied to the evolution of aposematic coloration (Leimar et al., 1986).

Evolutionary processes that involve learning are difficult to study. To our knowledge, no fully satisfactory model of learning exists today. We have chosen artificial neural networks as models of receivers because they do not constrain signal form and are known to generalize realistically when novel signals appear (Ghirlanda & Enquist, 1998). Nevertheless, available learning algorithms for neural networks, including the back-propagation one used here, are difficult to relate to real learning events (learning models such as those proposed by Rescorla & Wagner (1972) and Blough (1975) have similar advantages and disadvantages). Our simulations involved several simplifications of the actual learning sequence, but they capture the main difference between learned and inherited recognition: that receivers are born naive and will thus make mistakes while learning the appropriate behaviour.

To conclude, we discuss reality. For what signalling contexts will our results be relevant? Learned aposematism is an obvious case: predators and unprofitable preys have common interests, but these can not be exploited as long as the predator is inexperienced. Note, however, that our model potentially applies to all situations where senders benefit from being accurately recognized by receivers (examples are interactions between pollinators and flowers, and between fruit eaters and fruits). As stressed in §1, the model is not meant to predict why recognition is learned or inherited in each particular system. Nevertheless, the simulation results provide us with some definite predictions about the co-occurrence of learned versus inherited recognition and signal exaggeration. First, we expect signals that are learned by the receiver to be more exaggerated than signals for which recognition is genetically inherited. Warning signals are often learned by predators (Edmunds, 1974; Guilford, 1990), although genetic factors are also known to play a part. Inherited recognition is predicted to occur in conjunction with less colourful appearances. A potential example of dangerous but inconspicuous organisms is snakes (with a few exceptions). Interestingly, this is associated with innate avoidance responses in frogs (Ewert, 1980) and with strong genetic predispositions to develop fear for snakes in humans and monkeys (Mineka & Cook, 1988). Note that in the latter case we have left the domain of prey-predator interactions.

We may also conjecture that less-striking signals will be found in systems where a specialist predator has been associated with the same unprofitable prey species for a long time. In these cases, inherited recognition may evolve, whereas if predators are generalists (relying more on learning) we expect prey to develop more exaggerated signals. Our simulations also indicate that signals that are either

rarely or very often encountered by receivers might be less exaggerated than signals encountered an intermediate number of times (3). Thus, rare and very common species should be less colourful than those of intermediate density. Because in our simulations we did not consider performance before and during learning as relevant for fitness, the conclusion about rare species is uncertain. That individuals of common species should carry less-exaggerated signals seems to be a more robust prediction, particularly when signalling only occurs at short distances.

It is probably possible to derive further predictions from our results. For instance, an anonymous reviewer has suggested that exaggeration of signals in Múllerian mimicry systems should decrease as more unpalatable species join the system, because the same signal becomes increasingly common. We are, at present, unable to evaluate this and other predictions. As recalled in §1 and at the beginning of this section, many factors can affect the evolution of exaggeration. Not all of these are mutually exclusive (for instance, receiver biases and kinship in senders might contribute independently). We therefore expect that extensive studies, including both comparative analysis and tracing of the ontogeny of recognition, will be necessary to assess the relative importance of each factor in concrete biological situations.

Acknowledgements

We thank two anonymous referees for stimulating and insightful comments. This research has been supported by the Swedish National Research Council and Marianne och Marcus Wallenberg Stiftelse. M.K. has been supported by the Japan Society for the Promotion of Science. S.G. thanks the Group for Theoretical Physical Chemistry of the Department of Chemistry at Università di Roma La Sapienza.

References

- Arak, A. & Enquist, M. 1995 Conflict, receiver bias and the evolution of signal form. *Phil. T. Roy. Soc. Lond. B* 349, 337–344.
- Arnold, S. 1981 Behavioral variation in natural populations. I. Phenotypic, genetic, and environmental correlations between chemoreceptive responses to prey in the garter snake, *Thamnophis elegans*. *Evolution* 35, 489–509.
- Blair, W. F. 1964 Isolating mechanisms and interspecies interactions. *Q. Rev. Biol.* 39, 334–344.

- Blough, D. S. 1975 Steady state data and a quantitative model of operant generalization and discrimination. *J. Exp. Psychol. Anim. B* 104(1), 3–21.
- Bradbury, J. & Vehrencamp, S. 1998 *Principles of animal communication*. Sunderland Massachusetts: Sinauer Associates Inc.
- Dawkins, R. & Krebs, J. R. 1978 Animal signals: information or manipulation? In *Behavioural ecology. An evolutionary approach* (eds. J. R. Krebs & N. B. Davies), pp. 282–309. Oxford: Blackwell Scientific Publications.
- Edmunds, M. 1974 *Defence in animals. A survey of anti-predator defences*. Harlow: Longman Group Limited.
- Enquist, M. & Arak, A. 1993 Selection of exaggerated male traits by female aesthetic senses. *Nature, Lond.* 361, 446–448.
- Enquist, M. & Arak, A. 1994 Symmetry, beauty and evolution. *Nature, Lond.* 372, 169–172.
- Enquist, M. & Arak, A. 1998 Neural representation and the evolution of signal form. In *Cognitive ethology* (ed. R. Dukas), pp. 1–420. Chicago: Chicago University Press.
- Ewert, J.-P. 1980 *Neuroethology*. Berlin Heidelberg New York: Springer-Verlag.
- Ghirlanda, S. & Enquist, M. 1998 Artificial neural networks as models of stimulus control. *Anim. Behav.* 56, 1383–1389.
- Grafen, A. 1990 Sexual selection unhandicapped by the Fisher process. *J. Theor. Biol.* 144, 473–516.
- Guilford, T. 1990 The evolution of aposematism. In *Insect defences* (eds. D. L. Evans & J. O. Schmidt). Albany: State University of New York Press.
- Hamilton, W. 1964 The genetic evolution of social behaviour. I-II. *J. Theor. Biol.* 7, 1–16.
- Haykin, S. 1994 *Neural Networks: A Comprehensive Foundation*. New York: Macmillan, 2nd edn.
- Hinde, R. A. 1970 *Animal Behaviour*. Tokyo: McGraw-Hill Kogakusha, Ltd, 2nd edn.

- Hogan, J. & Bolhuis, J., eds. 1994 *Causal mechanisms of behavioural development*. Cambridge: Cambridge University Press.
- Kamo, M. & Iwasa, Y. 2000 Evolution of preference for consonances as a by-product. *Evol. Ecol. Res.* 2, 375–383.
- Kamo, M., Kubo, T. & Iwasa, Y. 1998 Neural network for female mate preference, trained by a genetic algorithm. *Phil. T. Roy. Soc. Lond. B* 353, 399–406.
- LeCun, Y. 1985 Une procédure d'apprentissage pour reseau a seuil asymmetrique. *Cognitiva* 85, 599–604.
- Leimar, O., Enquist, M. & Sillén-Tullberg, B. 1986 Evolutionary stability of aposematic coloration and prey unprofitability: a theoretical analysis. *Am. Nat.* 128(4), 469–490.
- Luce, R. D. & Raiffa, H. 1957 *Games and decisions*. New York: John Wiley.
- Mackintosh, N. 1974 *The psychology of animal learning*. London: Academic Press.
- Mallet, J. & Joron, M. 1999 Evolution of diversity in warning color and mimicry: polymorphism, shifting balance, and speciation. *Annu. Rev. Ecol. Syst.* 30, 201–233.
- McLaren, I. 1989 The computational unit as an assembly of neurons: an implementation of an error correcting learning algorithm. In Miall et al. (1989).
- Miall, C., Durbin, R. M. & Mitchison, G. J., eds. 1989 *The computing neuron*. Wokingham, England: Addison-Wesley.
- Mineka, S. & Cook, M. 1988 Social learning and the acquisition of snake fear in monkeys. In *Social Learning: Psychological and biological perspectives* (eds. T. R. Zentall & B. G. J. Galef). Hillsdale, NJ: Lawrence Erlbaum.
- Mitchison, G. 1989 Learning algorithms and networks of neurons. In Miall et al. (1989).
- Parker, G. 1985 Models of parent-offspring conflict. V. Effects of the behaviour of the two parents. *Anim. Behav.* 33, 519–533.

- Phelps, S. 2001 History's lessons. In *Anuran Communication* (ed. M. Ryan). London: Smithsonian Institution Press.
- Phelps, S. & Ryan, M. 1998 Neural networks predict response biases of female túngara frogs. *Proc. R. Soc. Lond. B* 265, 279–285.
- Rescorla, R. A. & Wagner, A. R. 1972 A theory of Pavlovian conditioning: Variations in the effectiveness of reinforcement and nonreinforcement. In *Classical conditioning: current research and theory* (eds. A. H. Black & W. F. Prokasy). New York: Appleton-Century-Crofts.
- Rumelhart, D. E., Hinton, G. E. & Williams, R. J. 1986 Learning internal representation by back-propagation of errors. *Nature, Lond.* 323, 533–536.
- Ryan, M. 1998 Sexual selection, receiver bias, and the evolution of sex differences. *Science* 281, 1999–2003.
- Salthe, S. N. & Mecham, J. S. 1974 Reproductive and courtship patterns. In *Physiology of the amphibia* (ed. B. Lofts), vol. 2, pp. 310–473. New York: Academic Press.
- ten Cate, C. 1994 Perceptual mechanisms in imprinting and song learning. In Hogan & Bolhuis (1994).
- Zahavi, A. 1975 Mate selection - a selection for a handicap. *J. Theor. Biol.* 53, 205–214.