How can automimicry persist when predators can preferentially consume undefended mimics?

Graeme D. Ruxton¹ and Michael P. Speed^{2,*}

¹*Division of Environmental & Evolutionary Biology, Institute of Biomedical and Life Sciences, Graham Kerr Building, University of Glasgow, Glasgow G12 8QQ, UK*

²*School of Biological Sciences, University of Liverpool, Liverpool L69 7ZB, UK*

It is common for species that possess toxins or other defences to advertise these defences to potential predators using aposematic ('warning') signals. There is increasing evidence that within such species, there are individuals that have reduced or non-existent levels of defence but still signal. This phenomenon (generally called automimicry) has been a challenge to evolutionary biologists because of the need to explain why undefended automimics do not gain such as a fitness advantage by saving the physiological costs of defence that they increase in prevalence within the population, hence making the aposematic signal unreliable. The leading theory is that aposematic signals do not stop all predatory attacks but rather encourage predators to attack cautiously until they have identified the defence level of a specific individual. They can then reject defended individuals and consume the undefended. This theory has recently received strong empirical support, demonstrating that high-accuracy discrimination appears possible. However, this raises a new evolutionary problem: if predators can perfectly discriminate the defended from the undefended and preferentially consume the latter, then how can automimicry persist? Here, we present four different mechanisms that can allow non-trivial levels of automimics to be retained within a population, even in the extreme case where predators can differentiate defended from undefended individuals with 100% accuracy. These involve opportunity costs to the predator of sampling carefully, temporal fluctuation in predation pressure, predation pressure being correlated with the prevalence of automimicry, or developmental or evolutionary constraints on the availability of defence. These mechanisms generate predictions as to the conditions where we would expect aposematically signalling populations to feature automimicry and those where we would not.

Keywords: automimicry; aposematism; secondary defence; toxins; predator; prey

1. INTRODUCTION

The co-evolution of anti-predator defences and advertisements of those defences in aposematic signalling systems continue to be a productive area of study among evolutionary ecologists (Lindström 1999; Maynard Smith & Harper 2003; Ruxton *et al.* 2004). However, there is mounting evidence from a variety of species of heterogeneity in levels of defence between identical looking individuals in the same local population (reviewed in Ruxton *et al.* 2004, pp. 176–179), a phenomenon often called automimicry (first coined in Brower *et al.* 1967). Explaining the evolution of automimicry has been a challenge to theoreticians (Brower *et al.* 1970; Pough *et al.* 1973; Gibson 1984; Gullan & Cranston 1994). If undefended individuals (automimics) can avoid predation through aposematic coloration, but avoid the energetic costs of investment in defences, then they would be expected to have a higher relative fitness than members of the population that do invest in defences. Eventually, the frequency of undefended individuals would be sufficiently high that it would no longer pay predators to heed aposematic signals, predators would ignore the signal and the evolutionary pressure for its maintenance would be lost. However, in the natural world we seem to see automimicry being maintained in stable signalling systems

(Cohen 1985; Alonso-Mejia & Brower 1994; Ritland 1994). Hence, the challenge has been to explain what stops the frequency of automimics relentlessly growing. The leading explanation has been the suggestion of Guilford (1994) that aposematic signals do not act to stop predators from attacking individuals with that signal, but rather act as a 'go-slow signal'. That is, at least some predators handle signalling individuals cautiously, and are able to determine the defence level of a specific individual before committing themselves to a full-blown attack; they then release defended individuals relatively unharmed but ingest undefended individuals.

This theory has received strong support from recent work with chicks eating artificial foods (Gamberale-Stille & Guilford 2004). In this experiment, chicks were able to 'taste-reject' 98% of unpalatable prey, but consume 98% of palatable but identical looking automimics. This lower 'survivorship' of automimics was true in a number of trials differing in the frequencies of automimics in the population (also Fink & Brower 1981). These results suggest a plausible explanation for why automimics do not increase in frequency, but they raise another question: how do automimics persist at all when predators can selectively consume them? That is, these results explain how an aposematic signalling system need not inevitably become destabilized by automimics but they do not explain how significant levels of automimicry can persist. If predators

* Author for correspondence (speedm@liverpool.ac.uk).

can distinguish automimics from protected individuals as effectively as suggested in the experiments by [Gamberale-Stille & Guilford \(2004\)](#), it seems that automimicry is untenable. However, in this paper, we will suggest a number of scenarios that allow substantial levels of automimicry to persist, even in a situation of extreme hostility to automimics, in which predators can sample prey and then distinguish defended from undefended individuals with 100% fidelity, such that defended individuals always survive an attack and undefended ones never do.

2. CAREFUL SAMPLING AND DISCRIMINATION OF PREY INDIVIDUALS MAY BE EXPENSIVE FOR PREDATORS

Let us assume that a given predator encounters individuals of species A and B at rates λ_A and λ_B , respectively. All individuals of species B are undefended and yield a single unit of energetic value to the predator if ingested; handling an encountered B individual takes time t_B . For species A, only a fraction p of encountered individuals are undefended (and again yield a single unit of energetic value if ingested), the remaining $1-p$ of species A have toxic defences that make them unattractive as prey. If the predator encounters a B individual, then it ingests it; if it encounters an A individual, then it could either ignore it or it can invest a time t_E in evaluating whether an individual is defended or not. If the predator then finds the individual to be undefended, an extra handling time t_A is required to consume it. If the predator ignores encountered A individuals and restricts its diet to B individuals, then its energy gain rate is given by standard time-budget methodology (e.g. [Holling 1965](#); [Lendrem 1986](#); [Stephens & Krebs 1986](#)) as:

$$R_B = \frac{\lambda_B}{1 + \lambda_B t_B}. \quad (2.1)$$

Similarly, if the predator adds undefended A individuals to its diet, then the energy gain rate when expanding the diet to both species becomes:

$$R_{AB} = \frac{\lambda_B + p\lambda_A}{1 + \lambda_B t_B + \lambda_A t_E + p\lambda_A t_A}. \quad (2.2)$$

Hence, undefended A individuals should only be incorporated into the diet if $R_{AB} > R_B$; using (2.1) and (2.2), this criterion for including undefended members of A in the diet becomes a condition on the fraction of automimics p :

$$p > \frac{t_E}{\frac{1}{\lambda_B} + t_B - t_A}. \quad (2.3)$$

If the right side of (2.3) is negative (i.e. if $t_A > t_B + \lambda_B^{-1}$) then including A individuals in the diet is unattractive even when all individuals in the population are undefended, and so automimics can rise to fixation. Essentially, in this case the long handling time of otherwise undefended individuals is sufficient to make them unattractive as prey, even without further defences. Similarly, if the right side of (2.3) is greater than one (i.e. $t_E + t_A > t_B + \lambda_B^{-1}$), then again the time spent consuming an A individual could be more profitability spent searching for B individuals, and so again we would expect automimics to face no predation pressure, even if they rise to fixation in the A population.

The interesting case occurs when $t_E + t_A < t_B + \lambda_B^{-1}$, and so the calculated value of p in (2.3) lies between zero and one. In this case, if the fraction of automimics remains below the value obtained from (2.3), then they will not be subject to predator pressure, and so (if there is a cost to defence), the frequency of automimics in the population will increase. However, if their prevalence rises above this critical value, then the predator will expand their diet to include undefended members of species A until such times as the prevalence of automimics falls below the critical value again. Hence, we would expect the fraction of the population made up of undefended individuals to remain relatively close to p through time.

From inspection of (2.3), we can see that the fraction of automimics that can be tolerated without incurring predation costs increases with the cost of sampling to differentiate defended individuals from undefended automimics (increasing t_E), increasing profitability of alternative prey (either increasing λ_B or decreasing t_B) or decreasing profitability of automimics (increasing t_A). Hence, we can see that optimal diet selection by predators may be one mechanism by which automimicry can persist even if predators are theoretically able to sample the prey population and perfectly distinguish defended individuals from automimics. This occurs because there are circumstances where it not profitable for predators to invest time in such sampling, and they do best by ignoring all encountered individuals of the population exhibiting automimicry.

3. PREDATION PRESSURE IS VARIABLE

Imagine, we have a prey species with a very simple life history strategy: individuals must survive for a time T before a single round of reproduction. We assume two types within that species: defended (D) and undefended (U). Defence against predators is costly, and this cost is felt in the potential number of offspring produced ($N_D < N_U$). In cases such as the monarch and queen butterflies in which defence is gained by larval use of particular toxin-conferring host plants ([Brower *et al.* 1967](#)), energetic costs may be incurred by: (i) ovipositing females searching for the correct host plants and (ii) by their larval forms metabolizing and/or storing the toxin such that autotoxicity is prevented (cf. [Zalucki *et al.* 2001](#)). Furthermore, when toxins are self-generated by prey there may still be costs from searching for the right nutritional precursors, from biosynthesis itself and again from prevention of autotoxicity ([Bowers 1992](#); [Ruxton *et al.* 2004](#)).

The benefit of defence lies in the likelihood of surviving long enough to reproduce. For a defended individual, we assume that survival is certain no matter how many times it is attacked by predators, whereas any attack on an undefended individual is always fatal. We assume that attacks on any individual happen as a Poisson process at rate λ . Thus, in a small interval of time dt , the probability of an attack occurring is simply λdt , and so the probability of an undefended organism surviving to reproduce is given by $\exp(-\lambda T)$. If we use the probability of surviving to reproduce multiplied by the number of offspring that can be produced as a measure of fitness, then the fitness of a defended individual is simply N_D , whereas that of an undefended individual is $N_U \exp(-\lambda T)$. It is easy to show that there is a critical value of predation rate λ_c , such that

the fitness of the two types is identical:

$$\lambda_c = \left(\frac{-1}{T}\right) \ln\left(\frac{N_D}{N_U}\right). \quad (3.1)$$

If $\lambda < \lambda_c$, then we would expect the fitness of the undefended type to be greater and so this type to increase within the population; whereas the defended type should increase in frequency if $\lambda > \lambda_c$. If, however, the predation rate fluctuates either side of the critical value then the changing direction of selection would allow both defended and undefended organisms to be retained at non-trivial levels within the population. This occurs despite defence being costly, and despite the ability of predators to differentiate and preferentially consume undefended individuals.

4. PREDATION PRESSURE INCREASES WITH THE PREVALENCE OF AUTOMIMICRY

Here, we assume an essentially identical ecological situation to the last section except that rather than being dependent of time, the rate at which individuals are attacked (λ) is an increasing function of the frequency of automimics in the population (p). For example, we could describe λ by

$$\lambda(p) = \frac{\lambda_{\max} p}{p + q}, \quad (4.1)$$

where $\lambda_{\max}$ and q are constants ($\lambda_{\max}$ is the maximum attainable attack rate and q is the value of p that leads to attack rate being half of its maximal value). If this is the case, it is easy to show that there is a critical value of p (p_c) such that the fitness of the defended and undefended types is equal:

$$p_c = \frac{q}{\frac{1}{\ln\left(\frac{N_U}{N_D}\right)} - \lambda_{\max} T}. \quad (4.2)$$

If the right side of (4.2) is greater than one, then no matter how frequent undefended individuals become, predation rate never becomes sufficiently high to warrant investment in defence, and we would expect the undefended type to rise to fixation. Similarly, if the right side is negative, this implies that even when the predation rate is maximal (i.e. $\lambda = \lambda_c$), the cost of defence is still too great to give defended individuals a fitness advantage, and again we would expect a monomorphic population of undefended individuals to evolve. However, in the case where p_c is between zero and one, we have a situation where if the prevalence of automimics is lower than this critical value, then the automimics will enjoy a fitness advantage over defended individuals. Whereas if the prevalence of automimics rises above this level, then the opposite is true, and the defended morph should grow in prevalence. Hence, in the case where there is a non-trivial value of p_c , we would expect the population to evolve a stable polymorphism with the fraction of the population being undefended being close to p_c . Again, such polymorphism in defence can occur despite defences being costly and despite predators being able to differentiate and selectively consume undefended individuals (see also Till-Bottraud & Gouyon 1992; Speed *et al.* in press).

5. THERE IS A DEVELOPMENTAL OR ENVIRONMENTAL CONSTRAINT ON THE USE OF DEFENCE

In the three scenarios described above, we describe how a polymorphism between similar-looking defended and undefended members of a population can persist in time. In this case, persistence of the polymorphism is achieved by natural selection, and there are some circumstances within these scenarios where lack of defence is adaptive. However, it is theoretically possible to find variation in defence occurring in a natural population, the cause of which is not adaptive, since undefended individuals would benefit from being defended, but are prevented from achieving this by some constraint. An example of this might occur in caterpillars where toxins have to be sequestered from host plants. Environmental scarcity of preferred host plants may cause some females to lay their young on plants that do not have the appropriate toxins for their offspring. More generally, we should be cautious about taking within-population variation in defence as sufficient proof that automimicry is adaptive. Rather, while variation is a necessary condition, this should be backed up by demonstrating that ecological circumstances can vary (either naturally or by experimental manipulation) such that either high or low defence is selectively favoured in the manner predicted by the last three sections.

6. DISCUSSION

If the mechanism allowing automimics to persist is optimal diet selection by predators, then we can make certain predictions about the ecological circumstances in which automimics can flourish. They can exist at ecologically relevant frequencies if their potential predators have alternative prey that is abundant, if the time to evaluate whether an individual is defended or not is long, or if automimics take a long time to assess or otherwise handle and/or are of low energetic value, relative to alternative prey types. Clearly, all species will be under selection pressure to make themselves uneconomical to potential predators, hence the interesting variable here is the time taken to evaluate whether an individual is defended or not. This produces selection on automimics to resemble defended individuals as accurately as possible. Conversely, since there is likely to be a cost to being sampled even if this leads to rejection, defended individuals will be under selection pressure to allow differentiation to occur as quickly as possible in a predation sequence. This can lead to evolutionary pressure for reliable, self-advertising physical defences such as spines that can be seen at a distance and hence are very difficult or impossible for undefended individuals to successfully mimic. Even for toxic defences, automimicry should lead to evolutionary pressure for the predator to detect these defences as early in the predation process as possible. Toxins should, therefore, be very close to (if not on) the surface of the animal itself, concentrated in areas that predators first come into contact with, and of such a chemical nature that they can be detected as easily as possible (cf. Poulton 1890).

In our simple model of predator selection, we assumed that the predator must avoid defended prey at all costs. There may be situations where the level of defence of defended individuals might be more moderate, such that

predators would prefer to avoid them if possible, but occasional ingestion can be tolerated (Kokko *et al.* 2003; Sherratt *et al.* 2004). This could lead to an increased behavioural repertoire being open to the predator. Instead of ignoring all individuals, or sampling carefully to weed out defended individuals, it could simply ingest all individuals of the automimicked population that it encounters. This would not be an attractive strategy if automimics are rare, but if automimics are plentiful then this may become an attractive strategy (if unselective processing allows a time and/or energy saving that compensates for occasional consumption of defended individuals). Under such circumstances, both defended and undefended individuals would suffer similar predation costs, and so—assuming defences are costly—automimics would rise to fixation. Hence, in order to avoid this outcome, defended prey must invest in defences of sufficient potency that unselective ingestion is not an attractive option to predators. We might, therefore, expect to see adaptive automimicry occur only in species where the defences adopted by some individuals in the population are of high potency to potential predators.

A key assumption of scenarios 2 and 3 is that defence is costly. If there is no cost to defence in these two scenarios then the defended morph always has higher fitness and we would not expect to find automimicry persisting in such systems. However, there is now growing empirical evidence that defences are costly in a wide variety of, though not necessarily all, circumstances (Ruxton *et al.* 2004, pp. 59–63).

We also demonstrate that a variable predation pressure can provide a theoretical mechanism by which both defended and undefended types can co-exist in the long term, even if defence is costly and predators can differentiate and preferentially consume undefended individuals. Temporal fluctuation in per-capita rate of attack can be driven by variation in the numbers of either the focal species, the predator population or the availability of alternative prey types. These fluctuations could be driven by either intrinsically unstable population dynamics, by external factors such as weather patterns and by migration. In the earliest theoretical paper on automimicry, Brower *et al.* (1970) argued that for the monarch, *Danaus plexippus*, predation pressure is likely to be highly variable because of migratory change, writing:

in most of North America [the number of prey per predator] is small during the northward migration in the spring. By late summer successive generations have increased the size of the butterfly populations all over North America and [the number of prey per predator] is large.

(Brower *et al.* 1970, p. 1064)

Furthermore, both defended and undefended phenotypes could be maintained if there is no temporal variation in predator number, but spatial variation combined with migration such that there is mixing of prey individuals between areas where defended individuals have an advantage and areas where undefended individuals have the advantage.

It is important to note that fluctuation itself is not a sufficient condition for retention of both phenotypes in the population: the fluctuation must be such that predation pressure repeatedly moves from one side of the critical

value of attack rate to the other; switches must also be sufficiently frequent that neither type is able to rise to fixation before selection changes to favour the other type. Further, if reproduction is synchronized, then switches must not be so frequent that although switches occur the average frequency of attacks over the course of a generation is always on one side of the critical value or the other.

(a) What range of chemical defences might show automimicry?

In the literature, automimicry is usually associated with constitutive defences that are not controlled responsively by the individual animal (e.g. a toxin stored in a butterfly's cuticle that is continually present within an individual). However, the scope of automimicry may be wider than this. In many arthropod species, constitutive chemical defences can be used responsively by defensive secretion (e.g. when a ladybird carries out costly reflex bleeding, Whitman *et al.* 1990; Grill & Moore 1998). This leads to the potential for automimicry in the sense that animals may be able to choose whether or not to 'fire off' for a given level of predatory threat, in which case we hypothesize that there may be automimics that tend to reserve their toxins and automodels that more readily deplete them and pay the costs of restoring them subsequently. Although, the data are to our knowledge currently lacking, it is an interesting possibility that prey that utilize their toxins responsively may contain a proportion of such automimics. Should this be the case, then models like those described here may apply (albeit in a modified form).

Furthermore, even if all capable prey always use their defensive secretion when attacked, there will still be some level of automimicry if the defensive act quickly depletes the defensive store, leaving recently disturbed individuals as effectively undefended automimics for a refractory period (e.g. Holloway *et al.* 1991). In this case, automimicry is a constraint of a particular defensive system and it does not need the treatments described here to explain it. However, if attacks were frequent (relative to the time required to recover defences) and predators could costlessly discriminate defended from undefended-depleted individuals, the entire system of defensive secretion would be unstable and some alternative defence would be necessary. Hence, prohibitive handling costs for predators, like those described in the first model of this paper, reduce the predatory burden on such transient obligate automimics, thereby rendering the overall strategy of defensive secretion (and a subsequent refractory period of vulnerability) profitable and stable. In other cases where the defence is truly inducible (i.e. a toxin is not present in an animal until it is induced by a predatory cue), other modelling schemes apply (Clark & Harvell 1992; Tollrian & Harvell 1999).

(b) How best might the assumptions and predictions of this theory be tested?

This paper was prompted by Gamberale-Stille & Guilford's (2004) demonstration that chicks can differentially consume undefended artificial prey, while rejecting very similar looking defended individuals. Since, all our theory rests on this behaviour, the most important empirical task is to explore its generality and fine detail.

Thus, it would be useful to explore whether similar selective behaviour is found in experiments that are closer to ecologically relevant conditions. Particularly, interesting would be experiments that utilized natural prey. Specifically, if it were possible to assay individuals of a prey species for strength of defence prior to offering them to predators, then one could test whether less defended individuals were preferentially consumed. This may be possible with prey species such as ladybirds that defend themselves by releasing chemically loaded reflex blood, a process that can be induced by experimenters, allowing the blood to be analysed for level of defence without destroying the prey (Holloway *et al.* 1991). A further benefit to such an experiment is that one could explore whether predators can differentiate between two non-zero levels of defence: Gamberale-Stille & Guilford's birds differentiated between defended and entirely undefended prey.

The key assumption of the first mechanism described in this paper is that careful sampling and discrimination of prey individuals is expensive (in time spent) for predators. It should be possible to test this assumption with, for example, similar chick experiments to those of Gamberale-Stille & Guilford (2004) by adding a control manipulation where only undefended prey are offered to the birds and no discrimination is required. One theoretical prediction from our first model that should be particularly amenable to testing is the prediction in equation (2.3) for when a prey type featuring automimicry should or should not be included in a predator's diet. All of the parameters of this equation could be controlled in laboratory experiments allowing predictions as to whether predators would exploit a given prey type that features automimicry: these predictions then being tested by observation of real predators. More challenging to organize, but still within the bounds of possibility would be to use equation (2.3) combined with empirically derived estimates of the necessary parameter values for a natural system, then compare the actual level of automimicry found in real prey populations with that predicted by equation (2.3).

Testing the mechanism that relates temporal variation in predation pressure to the existence of automimicry may be more challenging because obtaining quantitative predictions from the models would involve obtaining an estimate of the cost of defence. Measurements of the cost of defence are often technically challenging, although there have been notable successes (reviewed in Ruxton *et al.* 2004). However, a useful first step would be to measure temporal variation in predation pressure on a natural prey species that features automimicry, then equation (3.1) can be used to estimate how large the cost of defence would have to be for temporal variation in predation pressure to explain the maintenance of automimicry in this prey species. The plausibility of this estimated cost could then be evaluated in terms of our understanding of the physiology and life history of the species, and further could inform the design of experiments aimed at quantifying the cost empirically. Similarly, an important first step in exploring the importance of our third mechanism would be to use either natural variation in subpopulations in level of automimicry, or controlled variation in artificial-prey experiments to explore whether there is any evidence that predators do increase their

attack rate on populations that have higher levels of automimicry.

Particularly with reference to our final non-adaptive explanation for within-population variance in level of defence, it would be valuable to empirically explore whether an example of naturally occurring automimicry appears to be adaptive. Specifically, this requires exploring whether ecological circumstances can vary (either naturally or by experimental manipulation) such that either high or low defence is selectively favoured in the manner predicted by the first three explanations of this paper. This could be done by, for example, experimentally reducing predation pressure on a prey population that features automimics. If automimicry is driven by an adaptive explanation such as the first three given in this paper, then we would expect that this should lead to an increase in the frequency of automimics within the population. Such experiments could be carried out in laboratory colonies or in the wild using predator-excluding enclosures.

7. CONCLUSIONS

Despite its obvious evolutionary and ecological importance, automimicry has been subjected to very little theoretical scrutiny—in fact a tiny fraction of theoretical mimicry papers focus on automimicry (of 46 models listed in Ruxton *et al.* 2004, two involve automimicry). Early theoretical models simply estimated the effects of automimics on predation levels within populations (Brower *et al.* 1970), and other nonmathematical treatments considered the role of parasitoids in the maintenance of the phenomenon (Gibson 1984). Here, we have shown that the persistence of automimicry can be mathematically explained from an ecological viewpoint when the economics, variability and frequency-dependent properties of predation are considered in model formulation. We also argue that defended individuals will be under selection pressure to allow their differentiation from undefended individuals to occur as quickly as possible in a predation sequence. We hope that such models help to stimulate renewed empirical interest in this widespread and important phenomenon.

REFERENCES

- Alonso-Mejia, L. & Brower, L. P. 1994 From model to mimic: age-dependent unpalatability in monarch butterflies. *Experientia* **50**, 176–181. (doi:10.1007/BF01984960)
- Bowers, M. D. 1992 The evolution of unpalatability and the cost of chemical defense in insects. In *Insect chemical ecology: an evolutionary approach* (ed. B. D. Roitberg & M. B. Isman), pp. 216–244. New York: Chapman & Hall.
- Brower, L. P., Brower, J. V. Z. & Corvino, J. M. 1967 Plant poisons in a terrestrial food chain. *Proc. Natl Acad. Sci. USA* **57**, 893–898.
- Brower, L. P., Pough, F. H. & Meck, H. R. 1970 Theoretical investigations of automimicry. I. Single trial learning. *Proc. Natl Acad. Sci. USA* **66**, 1059–1064.
- Clark, C. W. & Harvell, C. D. 1992 Inducible defenses and the allocation of resources—a minimal model. *Am. Nat.* **139**, 521–539. (doi:10.1086/285342)
- Cohen, J. A. 1985 Differences and similarities in cardenolide contents of queen and monarch butterflies in Florida and their ecological and evolutionary implications. *J. Chem. Ecol.* **11**, 85–103. (doi:10.1007/BF00987608)

- Fink, L. S. & Brower, L. P. 1981 Birds can overcome the cardenolide defence of monarch butterflies in Mexico. *Nature* **291**, 67–70. (doi:10.1038/291067a0)
- Gamberale-Stille, G. & Guilford, T. 2004 Automimicry destabilizes aposematism: predator sample-and-reject behaviour may provide a solution. *Proc. R. Soc. B* **271**, 2621–2625. (doi:10.1098/rspb.2004.2893)
- Gibson, D. O. 1984 How is automimicry maintained. In *The biology of butterflies* (ed. R. I. Vane-Wright & P. R. Ackery), pp. 163–165. London: Academic Press.
- Grill, C. P. & Moore, A. J. 1998 Effects of a larval antipredator response and larval diet on adult phenotype in an aposematic ladybird beetle. *Oecologia* **114**, 274–282. (doi:10.1007/s004420050446)
- Guilford, T. 1994 Go-slow signaling and the problem of automimicry. *J. Theor. Biol.* **170**, 311–316. (doi:10.1006/jtbi.1994.1192)
- Gullan, P. J. & Cranston, P. S. 1994 *The insects: an outline of entomology*. London: Chapman & Hall.
- Holling, C. S. 1965 The functional response of predators to prey density, and its role in mimicry and population regulation. *Mem. Entomol. Soc. Canada* **45**, 1–60.
- Holloway, G. J., de Jong, P. W., Brakefield, P. M. & de Vos, H. 1991 Chemical defence in ladybird beetles (Coccinellidae). I. Distribution of coccinelline and individual variation in defence in 7-spot ladybirds (*Coccinella septempunctata*). *Chemoecology* **2**, 7–14. (doi:10.1007/BF01240660)
- Kokko, H., Mappes, J. & Lindström, L. 2003 Alternative prey can change model-mimic dynamics between parasitism and mutualism. *Ecol. Lett.* **6**, 1068–1076. (doi:10.1046/j.1461-0248.2003.00532.x)
- Lendrem, D. 1986 *Modelling in behavioural ecology*. Beckenham: Croom Helm.
- Lindström, L. 1999 Experimental approaches to studying the initial evolution of conspicuous aposematic signalling. *Evol. Ecol.* **13**, 605–618.
- Maynard Smith, J. & Harper, D. G. C. 2003 *Animal signals*. Oxford: Oxford University Press.
- Pough, F. H., Brower, L. P., Meck, H. R. & Kessell, S. R. 1973 Theoretical investigations of automimicry: multiple trial learning and the palatability spectrum. *Proc. Natl Acad. Sci. USA* **70**, 2261–2265.
- Poulton, E. B. 1890 *The colours of animals: their meaning and use especially considered in the case of insects*. London: Kegan Paul, Trench, Trubner & co. Ltd.
- Ritland, D. B. 1994 Variation in palatability of queen butterflies (*Danaus gilippus*) and implications regarding mimicry. *Ecology* **75**, 732–746.
- Ruxton, G. D., Sherratt, T. N. & Speed, M. P. 2004 *Avoiding attack: the evolutionary ecology of crypsis, aposematism and mimicry*. Oxford: Oxford University press.
- Sherratt, T. N., Speed, M. P. & Ruxton, G. D. 2004 Natural selection on unpalatable species imposed by state-dependent foraging behaviour. *J. Theor. Biol.* **228**, 217–226. (doi:10.1016/j.jtbi.2003.12.009)
- Speed, M. P., Broom, M. & Ruxton, G. D. In press. Automimicry and the evolution of discrete defences. *Biol. J. Linn. Soc.*
- Stephens, D. & Krebs, J. R. 1986 *Foraging theory*. Princeton: Princeton University Press.
- Till-Bottraud, I. & Gouyon, P.-H. 1992 Intra- versus interplant Batesian mimicry? A model on cyanogenesis and herbivory in clonal plants. *Am. Nat.* **139**, 509–520. (doi:10.1086/285341)
- Tollrian, R. & Harvell, C. D. 1999 The evolution of inducible defenses: current ideas. In *The ecology and evolution of inducible defenses* (ed. R. Tollrian & C. D. Harvell), pp. 306–322. Princeton University Press.
- Whitman, D. W., Blum, M. S. & Alsop, D. W. 1990 Allomones: chemicals for defense. In *Insect defenses: adaptive mechanisms and strategies of prey and predators* (ed. D. L. Evans & J. O. Schmidt), pp. 229–251. New York: State University of New York Press.
- Zalucki, M. P., Brower, L. P. & Alonso, A. 2001 Detrimental effects of latex and cardiac glycosides on survival and growth of first-instar monarch butterfly larvae *Danaus plexippus* feeding on the sandhill milkweed *Asclepias humistrata*. *Ecol. Entomol.* **26**, 212–224. (doi:10.1046/j.1365-2311.2001.00313.x)