

# Context-Dependent Discrimination and the Evolution of Mimicry

Øistein Haugsten Holen<sup>1,\*</sup> and Rufus A. Johnstone<sup>2,†</sup>

1. Centre for Ecological and Evolutionary Synthesis, Department of Biology, University of Oslo, P.O. Box 1050 Blindern, N-0316 Oslo, Norway;

2. Department of Zoology, University of Cambridge, Downing Street, Cambridge CB2 3EJ, United Kingdom

Submitted July 11, 2005; Accepted November 7, 2005;  
Electronically published January 9, 2006

---

**ABSTRACT:** Many mimetic organisms have evolved a close resemblance to their models, making it difficult to discriminate between them on the basis of appearance alone. However, if mimics and models differ slightly in their activity patterns, behavior, or use of microhabitats, the exact circumstances under which a signaler is encountered may provide additional clues to its identity. We employ an optimality model of mimetic discrimination in which signal receivers obtain information about the relative risk of encountering mimics and models by observing an external background cue and flexibly adjust their response thresholds. Although such flexibility on the part of signal receivers has been predicted by theory and is supported by empirical evidence in a range of biological settings, little is known about the effects it has on signalers. We show that the presence of external cues that partly reveal signaler identity may benefit models and harm mimics, harm both, or even benefit both, depending on ecological circumstances. Moreover, if mimetic traits are costly to express, or mimics are related to their neighbors, context-dependent discrimination can dramatically alter the outcome of mimetic evolution. We discuss context-dependent discrimination among signal receivers in relation to small-scale synchrony in model and mimic activity patterns.

*Keywords:* Batesian mimicry, aggressive mimicry, conditional strategies, signal detection theory, context-dependent discrimination, predation.

---

Animals discriminate between signals in a wide range of settings, including mate choice, foraging, and parental in-

vestment. The optimal response to a certain type of signaler may depend on context, and consequently, animals often show great flexibility in how they respond to signalers in various contexts (e.g., Jennions and Petrie 1997; Sherman et al. 1997; Liebert and Starks 2004). Phenotypic variation in signaler traits, environmental degradation of signals, and cognitive constraints may render it impossible for a signal receiver to discriminate perfectly. Instead, the best option may be to strike an optimal balance between two types of discrimination error: erroneous rejection of desirable signalers and erroneous acceptance of undesirable signalers. The optimal discrimination strategy will depend on the costs associated with each of the two types of error, on the probabilities of encountering desirable and undesirable signalers, and on the typical similarity and overlap in perceived signals (e.g., Getty 1985; Reeve 1989; Wiley 1994). If the proportion of desirable and undesirable signalers or the costs and benefits of responding to them change between contexts, signal receivers should be flexible and adopt different acceptance thresholds to optimize the balance of discrimination errors in each context (Reeve 1989; Sherman et al. 1997).

Signal discrimination plays a central role in Batesian mimicry and aggressive mimicry. Aggressive mimics are predators or parasites that use mimicry to deceive victims, such as avian brood parasites that lay mimetic eggs in the nests of other bird species (Rothstein and Robinson 1998; Davies 2000) and rewardless orchids that mimic female wasps and bees to attract males as pollinators (Nilsson 1992). Batesian mimics are palatable prey that deter predators from attacking by resembling unpalatable prey; a familiar example is hoverflies that mimic wasps, honeybees, or bumblebees (Gilbert 2005). Mimics often resemble their models very closely, which makes discrimination errors more likely. Many mathematical models based on signal detection theory have been constructed to investigate signal discrimination in Batesian mimicry and aggressive mimicry (e.g., Oaten et al. 1975; Getty 1985; Rodríguez-Gironés and Lotem 1999; Johnstone 2002; Sherratt 2002; Holen and Johnstone 2004). Such models typically predict that signal receivers should be easier to

\* Corresponding author; e-mail: o.h.holen@bio.uio.no.

† E-mail: raj1003@hermes.cam.ac.uk.

deceive (i.e., should more often misclassify mimics as models) when the ratio of mimics to models is low and when the cost of mistaking a model for a mimic is high relative to the cost of mistaking a mimic for a model; both predictions have empirical support (e.g., Brower 1960; Alcock 1970; Goodale and Sneddon 1977; Nonacs 1985; Alvarez 1996; Lindström et al. 1997; Brooke et al. 1998; Robert and Sorci 1999) and can be derived from other types of mimicry models (see Ruxton et al. 2004 for a review of models of Batesian mimicry).

Several experimental studies show that hosts of the common cuckoo *Cuculus canorus* may use the presence of an adult cuckoo near the nest as an indicator of a high probability of parasitism and may subsequently be more likely to reject eggs when the adult cuckoo has been observed (e.g., Davies and Brooke 1988; Moksnes et al. 1993; Bártol et al. 2002; Davies et al. 2003). Hosts have also been found to adjust their level of discrimination as the probability of parasitism declines through the season (Alvarez 1996; Brooke et al. 1998); the proximate cue may be encounters with adult cuckoos. Since host birds that routinely eject mimetic eggs may sometimes inappropriately eject some of their own eggs (Davies and Brooke 1988; Marchetti 1992; Welbergen et al. 2001), flexibility in discrimination could serve to reduce costly recognition errors when the risk of being parasitized is low (Davies and Brooke 1988; Davies 2000).

On a more general note, if mimics and models differ in their daily activity patterns, in behavior, or in their use of microhabitats, the exact circumstances under which a signaler is encountered could provide signal receivers with some information about its identity. For example, a predator may come to associate the noxious taste of a defended model not only with its appearance but also with some external background cue that tends to be present when the model is encountered (e.g., Uésugi 1996). When the mimetic resemblance is good, such background cues could potentially be more informative about signaler identity than minuscule differences in appearance. For instance, the thynnine wasp *Neozeleboria cryptoides* fails to discriminate between the chemical attractants of female conspecifics and sexually deceptive *Chiloglottis trapeziformis* orchids but learns to avoid the location of the orchid patch after having been deceived (Wong and Schiestl 2002; Schiestl et al. 2003; Wong et al. 2004). Data from experimental psychology show that learned discrimination and the extinction of conditioned responses can be context-dependent (Bouton 1993; Bouton 1994), which suggests that context could be important in mimicry systems that involve learning, such as Batesian mimicry. A recent study showed that predatory bats *Trachops cirrhosus* could quickly be trained to reverse their preferences for cues of

palatable and poisonous prey that had been acquired in the wild (Page and Ryan 2005).

Although factors that favor context-dependent discrimination by signal receivers are relatively well understood (e.g., Sherman et al. 1997), little is known about the effects it may have on signalers. In this article, we investigate such effects using mimicry as an example. For this purpose, we employ a signal detection model in which receivers flexibly adjust their response thresholds depending on whether the relative risk of encountering mimics is high or low.

### The Model

Suppose that a signal receiver encounters models and mimics sequentially in a random fashion. In each encounter, the probability that the signaler is a mimic is  $p$  and the probability that it is a model is  $1-p$ . The signal receiver attempts to classify signalers as models or mimics (and respond accordingly) on the basis of a continuous trait such as size or color. Both mimics and models benefit from being classified as models. The model population is monomorphic with trait value standardized to 0. The mimic population is monomorphic with trait value  $m$ , by convention positive, with trait values closer to 0 signifying better mimicry. Signal receivers do not perceive the trait perfectly. Instead, the perceived trait value of a signaler follows a standard normal distribution centered on the true trait value. At each encounter with a signaler, there are four possible outcomes: the signal receiver may correctly identify a model, erroneously classify a model as a mimic, erroneously classify a mimic as a model, or correctly identify a mimic. The payoffs to the signal receiver for these four outcomes are  $v_{11}$ ,  $v_{12}$ ,  $v_{21}$ , and  $v_{22}$ , respectively. Assuming that correct classifications yield higher payoff, we must have  $v_{11} > v_{12}$  and  $v_{22} > v_{21}$ . Because the perceived trait values of models and mimics overlap, the signal receiver cannot simultaneously minimize the risk of erroneously classifying models as mimics and the risk of erroneously classifying mimics as models. The costs of these two types of mistakes are given by  $v_{11} - v_{12}$  and  $v_{22} - v_{21}$ , respectively. As is common in signal detection models of mimicry (but see Getty 1985), we assume that the signal receiver maximizes payoff per encounter. Studies in experimental psychology indicate that animals tend to use shortsighted decision rules (Stephens and Anderson 2001; Stephens 2002). Although it is not predicted by classical rate-maximizing foraging models, shortsighted foraging strategies may be expected if future food benefits are strongly discounted (e.g., Kagel et al. 1986) or when the signal receiver suffers information-processing constraints (Stephens 2002). Shortsightedness may also be a reasonable assumption when handling times are very short relative to search time.

Our assumptions are equivalent to an equal-variance, normal-normal receiver operating characteristic (ROC; see Egan 1975). The optimal strategy for a signal receiver that maximizes expected benefit per encounter is to classify as models only those signalers whose perceived trait value falls below a certain threshold. The probabilities that a model and a mimic are perceived to have trait values below the threshold  $t$  are equal to  $Z(t)$  and  $Z(t - m)$ , respectively, where  $Z$  is the cumulative standard normal distribution. The expected payoff per encounter is

$$(1 - p)Z(t)v_{11} + (1 - p)[1 - Z(t)]v_{12} + pZ(t - m)v_{21} + p[1 - Z(t - m)]v_{22}, \quad (1a)$$

and it is maximized by adopting the threshold  $t^*$  as given by

$$t^*(m, K) = \frac{m}{2} - \frac{\ln(K)}{m},$$

where  $K = \frac{p(v_{22} - v_{21})}{(1 - p)(v_{11} - v_{12})}$ . (1b)

From the perspective of the models, the “mimetic load”  $K$  may be viewed as a partial measure of the negative ecological impact that mimics have on models. From the perspective of the signal receivers, the mimetic load may be seen as the incentive to identify mimics. The mimetic load is high when mimics are common and/or when the cost of misclassifying a mimic is high relative to the cost of misclassifying a model. The greater the mimetic load, the lower the threshold the signal receivers should set (leading to fewer misclassifications of mimics but more misclassifications of models).

In Batesian mimicry, the predator (signal receiver) benefits from attacking palatable prey (mimics) but should avoid attacking defended prey (models), and the payoffs satisfy  $v_{22} > v_{21} = v_{11} > v_{12}$ . In this case, the mimetic load is high when the proportion of mimics is high ( $p$  is high); when mimics are profitable, so that the cost of failing to attack a mimic is high ( $v_{22} - v_{21}$  is high); and when models are weakly defended, so that the cost of attacking a model is low ( $v_{11} - v_{12}$  is low). (We assume that the payoffs from attacking mimics and models are fixed, but see Sherratt [2003] for an approach that takes the predator’s hunger state into account.) In aggressive mimicry, the signal receiver benefits from interacting with models but should ignore or reject the harmful mimics, and the payoffs satisfy  $v_{11} > v_{12} = v_{22} > v_{21}$ . In this case, the mimetic load is high when the proportion of mimics is high ( $p$  is high), the cost of interacting with a mimic is high ( $v_{22} - v_{21}$  is high),

and the cost of failing to interact with a model is low ( $v_{11} - v_{12}$  is low).

#### *External Cues Provide Information about Signaler Identity*

Suppose now that signal receivers encounter signalers in different “recognition contexts” (Reeve 1989) and that the mimic-to-model ratio differs between these contexts. For instance, the ratio of Batesian mimics to models may be higher in certain patch types (e.g., on certain plant species) than in others, so patch type may be informative regarding the identity of signalers. Likewise, a worker of a social insect that discriminates between nest mates and mimetic social parasites may experience different mimic-to-model ratios inside and outside the colony. In avian brood parasitism, the sight of an adult brood parasite near the nest could indicate a higher-than-average risk of receiving a mimetic egg in the current clutch. In each case, the signaler can obtain information about whether an encounter is more or less likely than average to involve a mimic by assessing an external background cue (a “non-phenotypic” cue sensu Sherman et al. 1997).

The standard approach to signal discrimination is to assume that perceived signal values are unimodally (and most often normally) distributed (e.g., Oaten et al. 1975; Reeve 1989; Wiley 1994; Rodríguez-Gironés and Lotem 1999). If we consider an environment that is sufficiently fine grained relative to encounter rate, we may think of the perceived external cue as a random variable that in each signal encounter is drawn from a distribution characteristic of the given signaler type. If the respective distributions are unimodal, the standard approach may give us insight. However, this will not often be the case. Consider, for instance, a patchy environment in which mimics and models occupy the same patch types but divide their time differently between them. Although an individual signaler may spend any fraction of its time in a given patch type, this fraction cannot be observed directly by the signal receiver; in any given encounter, the signaler is observed to be in only one patch. If the patch types are dissimilar, the distribution of perceived patch characteristics for a signaler type may thus have several widely separated peaks.

We consider a discrete cue environment, and for simplicity, we assume it to consist of a single cue that with perfect accuracy is perceived as present or absent. Models and mimics differ in their propensity to be associated with the cue, and the cue is present with probability  $a$  in encounters that involve mimics and with probability  $b$  in encounters that involve models. We may by convention and with no loss of generality assume that  $a \geq b$  (since this is just a question of how we define the cue). We focus on the case where the background cue by itself is never sufficient to guarantee correct classification of a signaler; that is, where

$1 > a \geq b > 0$ . Letting “cue” denote the event that the cue is present and “mim” denote the event that a signaler is a mimic, then  $a = P(\text{cue}|\text{mim})$ , and  $b = P(\text{cue}|\text{mim}^c)$ , where the superscript C denotes the complement of the event. A straightforward calculation shows that  $P(\text{cue}) = P(\text{cue}|\text{mim})P(\text{mim}) + P(\text{cue}|\text{mim}^c)P(\text{mim}^c) = ap + b(1 - p)$ . According to Bayes’s rule, the conditional probability that a signaler is a mimic given that the cue is present is

$$\begin{aligned} P(\text{mim}|\text{cue}) &= \frac{P(\text{cue}|\text{mim})P(\text{mim})}{P(\text{cue})} \\ &= \frac{pa}{pa + (1 - p)b}, \end{aligned} \tag{2a}$$

while the conditional probability that a signaler is a mimic given that the cue is not present is

$$\begin{aligned} P(\text{mim}|\text{cue}^c) &= \frac{P(\text{cue}^c|\text{mim})P(\text{mim})}{P(\text{cue}^c)} \\ &= \frac{p(1 - a)}{p(1 - a) + (1 - p)(1 - b)}. \end{aligned} \tag{2b}$$

After substituting  $P(\text{mim}|\text{cue})$  for  $p$  in equation (1a), it turns out that a signal receiver alerted by the presence of the cue should adopt the optimal threshold

$$t_c^*(m) = \frac{m}{2} - \frac{\ln(K_C)}{m}, \text{ where } K_C = K\frac{a}{b}, \tag{3a}$$

while substituting  $P(\text{mim}|\text{cue}^c)$  for  $p$  in equation (1a) shows that a signal receiver unalerted by the cue should adopt the threshold

$$\begin{aligned} t_N^*(m) &= \frac{m}{2} - \frac{\ln(K_N)}{m}, \\ \text{where } K_N &= K\frac{(1 - a)}{(1 - b)}. \end{aligned} \tag{3b}$$

The aggregated parameters  $K_C$  and  $K_N$  give the effective mimetic loads, or incentives to identify mimics, in presence and absence of the cue. If  $a = b$ , signal receivers do not benefit from paying attention to the cue. If  $a > b$ , the presence of the cue indicates a higher-than-average probability that encountered signalers are mimics. It is clear from equations (3a) and (3b) that  $t_N^* > t_c^*$  when  $a > b$ , meaning that the signal receiver should adopt a stricter (lower) threshold and classify a higher proportion of signalers as mimics when alerted by the cue. If  $a$  increases and/or  $b$  decreases, the cue becomes a more reliable indicator of the identity of the signaler, and the difference

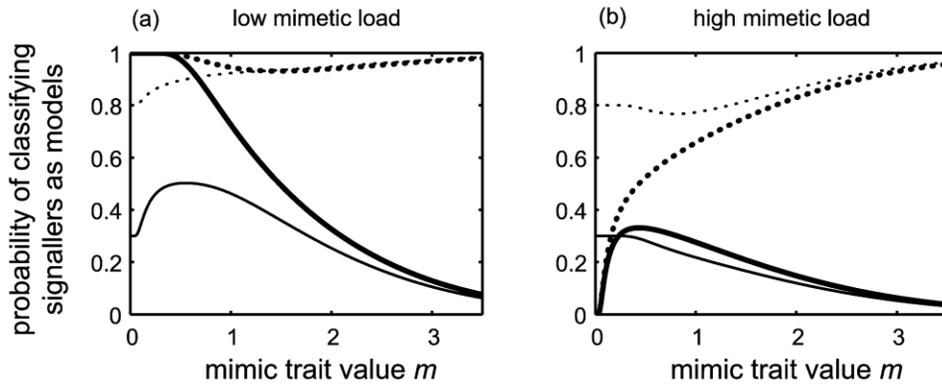
between the response thresholds adopted in presence and absence of the cue will increase. Note that for a given cue reliability, both  $t_c^*$  and  $t_N^*$  decreases with increasing mimetic load  $K$ .

### Optimal Classification

We assume that the distribution of background cues in the environment is sufficiently fine grained so that in any given encounter, an individual mimic will be associated with the cue with probability  $a$ , and an individual model will be associated with the cue with probability  $b$ . Signalers will experience different response thresholds in different encounters, and the overall probabilities that models and mimics are classified as models by signal receivers equal  $bZ(t_c^*(m)) + (1 - b)Z(t_N^*(m))$  and  $aZ[t_c^*(m) - m] + (1 - a)Z[t_N^*(m) - m]$ , respectively. First, consider the case where the external cue is completely uninformative ( $a = b$ ), so that  $t_c^*(m) = t_N^*(m) = t^*(m)$ . By differentiation, it is easy to show that if the mimetic load is low ( $K < 1$ ), a population of more accurate mimics will be classified as models with a higher probability than a population of less accurate mimics (fig. 1a, *thick solid line*). However, if the mimetic load is high ( $K > 1$ ), populations of sufficiently accurate mimics (i.e., with  $m < [2 \ln K]^{1/2}$ ) will be classified as models with a lower probability than slightly less accurate ones (fig. 1b, *thick solid line*). When mimicry is perfect, no discrimination is possible, and because signal receivers benefit from minimizing the occurrence of the classification mistake that has the highest associated cost (which is reflected in  $K$ ), they should either classify all signalers as models (when  $K < 1$ ) or classify all signalers as mimics (when  $K > 1$ ); see figure 1. These results are in accordance with predictions from other signal detection models (e.g., Oaten et al. 1975; Getty 1985; Rodríguez-Gironés and Lotem 1999).

The picture becomes more complex when the external cue is informative (i.e., when  $a > b$ ), since this makes the effective mimetic load higher in presence of the cue than in absence of the cue (i.e.,  $K_C > K_N$ ). For instance, the probability that mimics are classified as models may now decrease as the mimetic resemblance increases even if the overall mimetic load is  $< 1$  (fig. 1a, *thin solid line*). This is possible because the effective mimetic load in presence of the cue ( $K_C$ ) is  $> 1$ . When  $K_C > 1 > K_N$  and mimics are perfectly accurate, signalers are classified as models if and only if the cue is not perceived; models are thus correctly identified with probability  $1 - b$ , while mimics are classified as models with probability  $1 - a$  (fig. 1a, 1b, *thin lines*).

When do we expect signal receivers to pay attention to informative cues? Suppose that the value of the information gained from using cues can be measured as the



**Figure 1:** Probabilities that signal receivers correctly identify models (*dotted lines*) and erroneously classify mimics as models (*solid lines*) are shown as a function of mimic trait value. The model trait value is standardized to 0; a mimic trait value closer to 0 signifies better mimicry. Thin lines indicate that an external cue provides signal receivers with information about the relative risk of encountering mimics ( $a = 0.7$ ,  $b = 0.2$ ). Thick lines indicate that no such cue is available, or the cue is completely uninformative ( $a = b$ ). Parameters: (a),  $K = 1/3$ ; (b),  $K = 1.1$ .

increase in expected payoff per encounter to the signal receiver. The change in payoff equals

$$\begin{aligned} & (1 - p)(v_{11} - v_{12})[bZ(t_C^*(m)) + (1 - b)Z(t_N^*(m)) \\ & - Z(t^*(m))] - p(v_{22} - v_{21})[aZ(t_C^*(m) - m) \\ & + (1 - a)Z(t_N^*(m) - m) - Z(t^*(m) - m)]. \end{aligned} \quad (4)$$

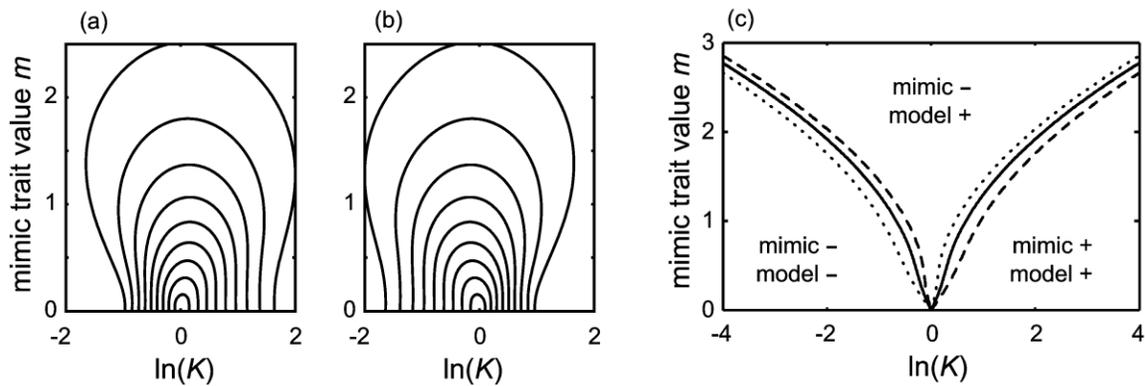
Numerical explorations show that signal receivers benefit most from evaluating cues when  $K$  is close to 1 and mimetic accuracy is high (illustrative cases are shown in fig. 2a, 2b). Signalers may be affected by the use of cues in various ways. The expressions in the first and last pair of square brackets of equation (4) give the resulting change in probability of being classified as a model for models and mimics, respectively. The use of cues may be favorable for models but not mimics, detrimental for both, or favorable for both (fig. 2c; see also fig.1). For a given level of mimetic accuracy and cue reliability, the following pattern is found: when  $K$  is sufficiently low, use of cues is detrimental to both mimics and models. By attending to the cue, the signal receiver increases the net payoff by reducing the expected cost of misclassifying mimics at the price of a smaller increase in the expected cost of misclassifying models. When  $K$  is intermediate, use of cues is detrimental to mimics but beneficial to models. The signal receiver simultaneously reduces its expected costs of misclassifying mimics and misclassifying models. When  $K$  is sufficiently high, use of cues is beneficial to both mimics and models. The signal receiver reduces the expected cost of misclassifying models at the price of a smaller increase in the expected cost of misclassifying mimics. However, using cues to optimally adjust response thresholds cannot

favor mimics and not models, since this would imply that the adjustments increased the probability of misclassifying mimics without decreasing the probability of misclassifying models, which would reduce signal receiver payoff.

### Mimetic Evolution

We have shown that paying attention to an informative background cue can change the overall probabilities of signal receivers correctly identifying models and erroneously misclassifying mimics. Moreover, it causes the level of discrimination that signalers experience to vary between encounters. This has implications for the evolution of mimicry. All other things being equal, an individual mimic always benefits from a closer resemblance to the model. However, the magnitude of this benefit, and thus the strength of selection favoring more accurate mimicry, may change when receivers attend to an informative cue. If there are opposing selection pressures that favor less accurate mimicry, then the balance between the two selective pressures, and the resulting outcome, may also change.

Below, we consider two scenarios in which there are selective pressures opposing the evolution of accurate mimicry. First, in a population that exhibits local relatedness, kin selection may favor less accurate mimics (Johnstone 2002). Second, if mimetic traits may interfere negatively with other functions of a mimetic organism, then this cost of accurate mimicry may outweigh the benefits to be gained (Holen and Johnstone 2004). In both cases, we assume that the response threshold adopted by signal receivers at any time is determined by the current typical level of mimetic accuracy, thus restricting ourselves to mimicry systems where mimetic discrimination has a



**Figure 2:** *a, b*, A signal receiver that attends to informative background cues when discriminating between mimics and models will be able to increase its payoff per encounter relative to a signal receiver that ignores background cues. A contour plot of the magnitude of this increase is shown. The contours are spaced at 10% intervals of the increase that may be obtained when  $m = 0$  (perfect mimicry) and  $\ln(K) = 0$ ; more than 90% of this increase may be obtained inside the innermost contour but less than 10% outside the outermost contour. Parameters:  $(v_{22} - v_{21}) = (v_{11} - v_{12})$ . (*a*),  $a = 0.9$ ,  $b = 0.3$ . (*b*),  $a = 0.7$ ,  $b = 0.1$ . *c*, The use of informative cues by signal receivers to optimally adjust their discrimination strategies may have a positive effect on models and a negative effect on mimics, a negative effect on both, or a positive effect on both. Parameters: *solid line*,  $a = 0.8$ ,  $b = 0.2$ ; *dashed line*,  $a = 0.9$ ,  $b = 0.3$ ; *dotted line*,  $a = 0.7$ ,  $b = 0.1$ .

learning component. We show that the use of an informative cue by signal receivers can dramatically alter the evolutionary outcome.

#### Kin Selection

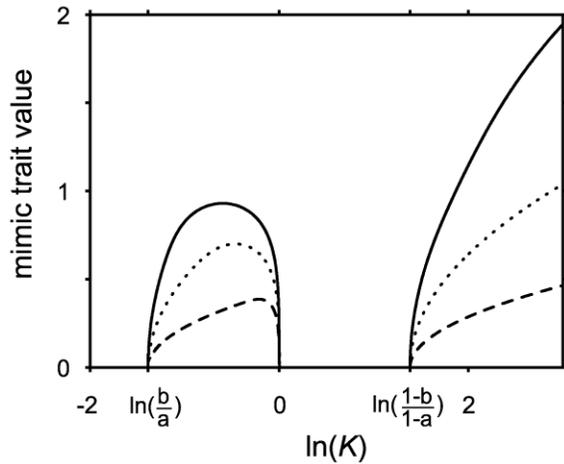
In this section, we follow Johnstone (2002) and assume that individual signal receivers encounter signalers sequentially and randomly only within local neighborhoods (e.g., within their own territory). Individual mimics tend to be related to others in the local neighborhood. Since kin are likely to share the same phenotype, more accurate mimics are then likely to be surrounded by more accurately mimetic neighbors. We assume that mutations affecting the mimetic trait have very small phenotypic effects and that signal receivers adopt response thresholds in presence and absence of the external cue that are optimal given the locally average mimic trait value. The probability that a mimic deceives the signal receiver (i.e., is misclassified as a model) as a function of individual trait value  $m'$  and locally average trait value  $\bar{m}$  can then be written

$$u(m', \bar{m}) = aZ_c^*(\bar{m}) - m' + (1 - a)Z_N^*(\bar{m}) - m'. \quad (5)$$

Although the probability that an individual mimic deceives the signal receiver strictly increases as its trait value  $m'$  changes in the direction of the model trait value (i.e.,  $\partial u(m', \bar{m})/\partial m' < 0$ ), it may sometimes decrease as the locally average mimic trait value changes toward the model trait value (i.e., we may have  $\partial u(m', \bar{m})/\partial \bar{m} > 0$ ; see also

Johnstone 2002). We assume that mimic fitness is a continuously differentiable, strictly increasing function of the probability of deceiving the signal receiver. Employing the “direct fitness” method (Taylor and Frank 1996; Day and Taylor 1998; Frank 1998), the combined effects of individual selection and kin selection can be shown to lead to less accurate mimicry if  $\partial u(m', \bar{m})/\partial m' + r\partial u(m', \bar{m})/\partial \bar{m} > 0$ , evaluated at  $m' = \bar{m} = m$ , where  $m$  (as before) is the population-wide resident phenotype and  $r$  is a relatedness coefficient that specifies the regression of local average trait value on the trait value of the focal mimic (see Johnstone 2002). It is straightforward to show that this condition can only be satisfied if a slightly more accurate mimic population has a lower probability of deceiving signal receivers (i.e., if  $du(m, m)/dm > 0$ ). As discussed earlier, when informative external cues are unavailable or ignored by signal receivers, this happens only if the mimetic load  $K > 1$ ; therefore, kin selection can only hinder the evolution of more accurate mimicry if  $K > 1$ , as found by Johnstone (2002).

When signal receivers attend to informative cues, by contrast, the effective mimetic load in the presence of the cue,  $K_c$ , may be  $>1$  even if the overall mimetic load  $K < 1$ , which may cause the probability of deceiving signal receivers to decrease as the trait value of the mimic population changes toward the model trait value (fig. 1, 3). Kin selection may then lead to evolution of less accurate mimicry also when the overall mimetic load  $K < 1$ ; an illustrative case is shown in fig. 3. On the other hand, when signal receivers attend to cues, a mimetic load  $K > 1$  does not anymore guarantee that the probability of de-



**Figure 3:** When signal receivers attend to informative background cues, kin selection can lead to evolution of less accurate mimicry both when the mimetic load  $K$  is above and below 1. The probability of deceiving signal receivers may increase (*above the solid lines*) or decrease (*below the solid lines*) if the mimic population evolves a closer resemblance to the model population (i.e., a lower trait value). When it decreases, kin selection may potentially lead to evolution of less accurate mimicry, depending on the local relatedness  $r$  among mimics. If  $r = 1/2$ , mimic populations below the dotted lines will evolve less accurate mimicry; if  $r = 1/8$ , mimic populations below the dashed lines will evolve less accurate mimicry. Note that the effective mimetic load in presence of the cue,  $K_C$ , is higher than 1 when  $K > b/a$  and that the effective mimetic load in absence of the cue,  $K_N$ , is higher than 1 when  $K > (1 - b)/(1 - a)$ . Parameters:  $a = 0.8$ ,  $b = 0.2$ .

ceiving signal receivers should decrease as the mimic population closely approaches the model trait value. If  $K_C$  is very high, signal receivers alerted by the cue tend to adopt very strict (low) thresholds, and the overall probability of being classified as a model may be most sensitive to the response threshold adopted by unalert signal receivers. Thus, if the effective mimetic load in the absence of the cue,  $K_N$ , is  $< 1$ , the probability of deceiving signal receivers may strictly increase as the mimic population approaches the model trait value, even if the overall mimetic load is  $> 1$ . This is illustrated in the range  $1 < K < (1 - b)/(1 - a)$  in figure 3.

It should be noted that the precise kin selection condition that we have used assumes negligible competition between relatives, that is, that population regulation is global. In general, competition among relatives will decrease the kin selection advantage of altruism (Frank 1998; West et al. 2002).

#### Costs of Mimicry

Holen and Johnstone (2004) investigated the evolution of mimicry under the assumption that mimetic traits may

interfere negatively with other functions of a mimetic organism. For instance, changes in body size or shape may possibly affect fecundity (Cushing 1997) and flight performance (Srygley 2004), while changes in brightness and color may affect thermoregulation (Ottenheim et al. 1999). Gilbert (2005) discusses possible costs of mimicry in hoverflies. As in scenario A1 in Holen and Johnstone (2004), suppose that an obligate aggressive mimic obtains a resource of potential value  $b$  each time it deceives a signal receiver. Costs of mimicry are incorporated via the function  $S(m)$ , which represents the efficacy with which the resource can be converted into fitness equivalents by a mimic with trait value  $m$ . When the mimic trait value equals  $m_{\text{opt}}$ , the mimic's ability to utilize the resource is maximized. Formally, the trait value  $m_{\text{opt}}$  is a global maximum point of  $S(m)$ , and  $S(m)$  falls off on each side of  $m_{\text{opt}}$ , enabling only smaller portions of the benefit to be converted into fitness equivalents. For trait values lower than or equal to some threshold  $m_{\text{min}}$ , no benefit at all may be realized. We will use a second-order polynomial with a maximum in  $m_{\text{opt}}$  as our choice for  $S(m)$  on the interval  $(m_{\text{min}}, m_{\text{opt}})$ ; however, similar results have been obtained with other log-concave functional forms. The mimic maximizes expected fitness gain per encounter. Assuming that a rare mutant mimic will have a negligible effect on the optimal signal receiver threshold, the expected payoff per encounter for a rare mutant  $m'$  in a population with trait value  $m$  is

$$w_m(m', m) = [aZ(t_C^*(m) - m') + (1 - a)Z(t_N^*(m) - m')]S(m')b. \quad (6a)$$

We are concerned with identifying evolutionarily stable strategies (ESSs). Candidate strategies  $m^*$  are found by (numerically) solving

$$\left. \frac{\partial}{\partial m'} w_m(m', m^*) \right|_{m'=m^*} = 0 \quad (6b)$$

with respect to  $m^*$ . Since  $Z[t_C^*(m) - m']$  and  $Z[t_N^*(m) - m']$  are decreasing functions of  $m'$  and  $S(m')$  falls off on each side of  $m_{\text{opt}}$ , it can be seen from equation (6a) that any optimal mimic strategy  $m^*$  must satisfy  $m_{\text{min}} < m^* < m_{\text{opt}}$ . Having found a candidate strategy  $m^*$ , we check (numerically) that  $m^*$  is uninvadable by mutants and convergently stable. The condition for  $m^*$  to be uninvadable by nearby mutants is

$$\left. \frac{\partial^2}{\partial m'^2} w_m(m', m^*) \right|_{m'=m^*} < 0, \quad (6c)$$

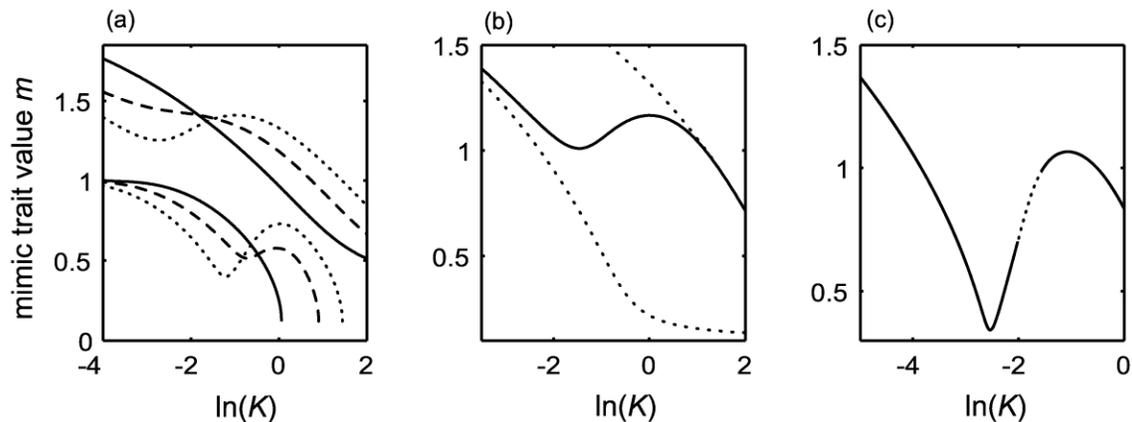
and a sufficient condition for  $m^*$  to be an ESS is that  $w_m(m^*, m^*) > w_m(m', m^*)$  for all  $m' \neq m^*$ . A condition for  $m^*$  to be convergence stable is

$$\frac{\partial}{\partial m^*} \left\{ \left[ \frac{\partial}{\partial m'} w_m(m', m^*) \right]_{m'=m^*} \right\} < 0. \quad (6d)$$

If the cue is uninformative ( $a = b$ ), mimicry scenario A1 from Holen and Johnstone (2004) is recovered. In this case, stronger discrimination on behalf of the signal receivers will select for better mimicry and the mimetic accuracy at equilibrium will increase with mimetic load; moreover, if  $m_{\min} < 0$  and  $K$  is sufficiently high, no stable equilibrium may exist, leading to an evolutionary chase toward better mimicry and stronger discrimination (fig. 4a). Introduction of a weakly informative cue does not qualitatively change the outcome. However, when the cue is highly informative ( $a$  is high and  $b$  low), then mimetic accuracy at equilibrium may in fact decrease with mimetic load over an intermediate range (fig. 4a). The explanation is quite simple. Alert signal receivers always adopt stricter (lower) thresholds than unalert signal receivers and are thus more difficult to deceive. When mimetic load is low, alert signal receivers set relatively lenient (high) rejection thresholds, and mimics obtain substantial benefits from deceiving them. However, as mimetic load increases, it

becomes increasingly costly for mimics to have the level of mimicry that is effective for deceiving alert signal receivers. The mimics instead profit from having a less accurate, cheaper level of mimicry that enables them to realize larger fitness benefits when deceiving unalert signal receivers. If a mimic could flexibly adjust its level of mimicry (and the associated cost) according to context, it should be more accurate when encountering alert signal receivers (fig. 4b, *dotted lines*). However, when mimetic accuracy is fixed, the optimal mimic trait reflects the relative importance of deceiving alert and unalert signal receivers, which vary with mimetic load (fig. 4b, *solid line*). Just as a predator in a predator-prey arms race could be selected to specialize on young and vulnerable prey, a mimic could be selected to specialize on the most easily duped victims.

For some choices of parameter values, we find convergently stable fitness minima (fig. 4c), which represents evolutionary branching points (Geritz et al. 1998). Note that such branching points only arise when signal receivers pay attention to informative cues and were not found in scenario A1 in Holen and Johnstone (2004). Thus, flexible discrimination may potentially change the evolutionary outcome from monomorphism to dimorphism (a full exploration of mimetic dimorphism is outside the scope of this article). In three other mimicry scenarios in Holen



**Figure 4:** *a*, Evolutionarily stable levels of mimicry are shown for various mimetic loads  $K$ . Assuming costs to mimicry, the evolutionarily stable level of mimicry represents a compromise between the need to deceive signal receivers and other aspects of fitness. Only evolutionarily stable strategies (ESSs) that are attainable by evolution (i.e., that are convergence stable) are shown. The external background cue varies from uninformative (*solid lines*) to informative (*dashed lines*) to highly informative (*dotted lines*) of the relative risk of encountering mimics. Mimics may be strongly constrained by costs (i.e., *upper three lines*:  $m_{\min} = 0.3$ ,  $m_{\text{opt}} = 2$ ; *solid line*:  $a = b$ ; *dashed line*:  $a = 0.6$ ,  $b = 0.05$ ; *dotted line*:  $a = 0.8$ ,  $b = 0.05$ ) or less strongly constrained by costs (i.e., *lower three lines*:  $m_{\min} = -0.5$ ,  $m_{\text{opt}} = 1$ ; *solid line*:  $a = b$ ; *dashed line*:  $a = 0.7$ ,  $b = 0.3$ ; *dotted line*:  $a = 0.8$ ,  $b = 0.2$ ). *b*, The optimal trait value against unalert signal receivers (*upper dotted line*) and the optimal trait value against alert signal receivers (*lower dotted line*) both differ from the evolutionarily stable level of mimicry (*solid line*). Parameters:  $m_{\min} = 0.1$ ,  $m_{\text{opt}} = 2$ ,  $a = 0.85$ ,  $b = 0.15$ . *c*, When signal receivers pay attention to background cues, convergent stable branching points (*dotted line segment*) may arise. Solid line segments represent convergent stable fitness maxima; these are globally stable ESSs, except in the range  $-2.38 < \ln(K) < -1.57$ , where stability is only local. Parameters:  $a = 0.7$ ,  $b = 0.05$ ,  $m_{\min} = -1.2$ ,  $m_{\text{opt}} = 2.5$ .

and Johnstone (2004), branching points were indeed found, but only for mimetic loads  $>1$ . Flexible discrimination makes branching possible at mimetic loads well below 1 (fig. 4c).

Note that the results regarding evolutionarily stable levels of mimicry (but not those regarding branching) may be applicable also to situations where the discrimination strategies of the signal receivers are genetically determined. If  $m^*$  is an ESS, it must be the unique best answer to  $t_C^*(m^*)$  and  $t_N^*(m^*)$ , which are unique and (by definition) the best answers to  $m^*$ . Thus, the strategy coalition  $(m^*, t_C^*(m^*), t_N^*(m^*))$  will be evolutionarily stable under the standard assumption that mutants only arise one at a time in the coalition. The question of convergence stability is more difficult in coevolutionary scenarios and cannot be answered without making particular assumptions about the evolutionary dynamics and possibly also the genetic correlations between the two threshold traits.

### Discussion

In models of signal discrimination, the key factors determining the optimal discrimination strategies of signal receivers are the similarity of the signals, the relative rates at which desirable and undesirable signalers are encountered, and the costs of the two types of discrimination error. Our mimicry model captures the two latter factors in the “mimetic load,” an aggregated parameter that increases with the mimic-to-model ratio and with the ratio of the cost of misclassifying mimics to the cost of misclassifying models. In Batesian mimicry, the mimetic load is low when mimic prey are unprofitable and/or rare and model prey are strongly defended. In aggressive mimicry, it is low when mimics are rare and/or impose only small costs on their victims. It is usually held that deceptive signalers should be more successful when it is more difficult to distinguish between the signals. This is supported in many experimental studies of Batesian mimicry, in which attack rates on mimics decrease when the mimics are more similar to the models (e.g., Duncan and Sheppard 1965; Alcock 1970; Goodale and Sneddon 1977). According to long-standing predictions, however, when the mimetic load is high, populations of slightly inaccurate mimics should be attacked less often than accurate ones (Oaten et al. 1975; Johnstone 2002). Our novel results show that the same may happen at low mimetic loads and therefore in a much wider range of mimicry systems than previously thought. Our results apply to both aggressive and Batesian mimicry but require that mimics and models are differently associated with an external cue (e.g., because they are differently distributed among patch types) and that signal receivers pay attention to this cue when they discriminate. There are still insufficient data available to test

these predictions, since relevant experimental studies consider neither high mimetic loads nor external cues. However, in other biological settings, it is well established that signal receivers may adopt different discrimination strategies in different contexts (e.g., Jennions and Petrie 1997; Sherman et al. 1997; Liebert and Starks 2004), and there is no reason to believe that mimicry should be any different. Therefore, if we ignore the potential role of context and focus only on overall mimic and model abundances, we may easily come to the wrong conclusions regarding the typical selection pressures acting in mimicry systems.

External cues can provide signal receivers with additional information about the identities of signalers. At first sight, one could be tempted to think that signal receivers should use such information to reduce both kinds of discrimination error. However, our optimality analysis showed that this should happen only over an intermediate range of mimetic loads (the exact range depends on mimetic accuracy and cue reliability). For mimetic loads below this range, a signal receiver paying attention to external cues will in fact be more liable to misclassify models, and above this range, it will be more liable to misclassify mimics. A specific example may aid our intuitive understanding. In Batesian mimicry, the mimetic load may be thought of as a partial measure of the “profitability” of the mimic-model complex. If the mimetic load is very low and no informative cues are present, attack rates may be very low. Background cues that are informative of the mimic-to-model ratio may then make the mimicry complex profitable in at least some contexts, leading to increased attack rates on both prey types. On the other hand, if the mimetic load is very high and no informative cues are present, predators may attack almost all mimics and models they encounter. Informative background cues may then make the mimicry complex unprofitable in at least some contexts, thus leading to reduced attack rates on both prey types. It should be possible to test experimentally whether predators can associate different contextual cues with different mimic-to-model ratios and whether contextual cues can both protect and harm mimics, as predicted by the model.

When signal receivers are flexible and adopt different response thresholds in different contexts, selection pressures acting on signalers may change. If there are no opposing selection pressures favoring less accurate mimicry, individual mimics will always benefit from a closer resemblance to the model. In this case, flexible discrimination only acts to change the strength of selection on the mimic. Nevertheless, this may have important evolutionary consequences. A signal detection model of Batesian mimicry showed that when the mimetic load is sufficiently low, selection to improve mimicry will be very weak beyond a certain point, since the models will lend almost full pro-

tection also to slightly inaccurate mimics (Sherratt 2002). A similar point could be made for aggressive mimicry. However, the presence of informative external cues may remove much of the protection (fig. 1a), suggesting that significant selection for accurate mimicry may be maintained at much lower mimetic loads than otherwise expected.

When selection pressures opposing the evolution of accurate mimicry are present, flexible discrimination may alter the outcome of mimetic evolution more drastically. It has been shown that kin selection may oppose individual selection and favor inaccurate mimicry (Johnstone 2002). However, this result was limited to high mimetic loads (i.e.,  $K > 1$ ). We have shown that when signal receivers adjust their response thresholds according to external cues, kin selection may lead to evolution of inaccurate mimicry also at low mimetic loads. This result has important ramifications. For Batesian mimicry, it means that kin selection could be an important factor when mimics are rare and/or models are well defended, as probably is the case in the majority of systems. Inaccurate mimicry is then favored because accurate mimicry leads to less cautious attack by predators in those contexts where the effective mimetic load is high. Moreover, the result suggests a stronger role for kin selection in the evolution of aggressive mimicry, where high mimetic loads easily amount to unrealistically strong ecological pressures. Examples may be found in systems of aggressive mimicry where the discrimination task involves learning. Learning is important, for instance, in food-deceptive pollination systems (Dafni 1984).

We also investigated the evolution of aggressive mimicry under the assumption of a trade-off between mimetic accuracy and other functions of the mimetic organism. In aggressive mimicry, flexible discrimination could be seen as a form of inducible defence (Tollrian and Harvell 1999), where victims maintain effective resistance against deception at low cost by adjusting the response threshold according to risk. An earlier model of obligate aggressive mimicry showed that the level of mimicry at equilibrium (when it exists) should always be monomorphic and that the mimetic accuracy should improve as the mimetic load increases (scenario A1, Holen and Johnstone 2004). The latter result is due to the fact that signal receivers discriminate more strongly against mimics at higher mimetic loads. We incorporated an informative external cue into this model and found two important differences. First, if the cue is fairly reliable, mimetic accuracy at equilibrium may in fact decrease with mimetic load over an intermediate range. This is because the level of mimicry needed to effectively deceive victims that are alerted by the presence of the cue may become too costly to the mimic, so that it pays to instead have a cheaper level of mimicry that maximizes the return from deceiving those victims that

are encountered in absence of the cue. Second, very reliable external cues may give rise to branching points (Geritz et al. 1998), which indicate disruptive selection on the mimics and possible evolution of mimetic dimorphism (involving morphs with different levels of accuracy). Similar results were obtained earlier in a model of facultative aggressive mimicry in which a fraction of the resources could be obtained regardless of mimic trait value (e.g., from nondiscriminating victims) while the remaining resources could only be obtained through deception (scenario A3, Holen and Johnstone 2004). In this earlier model, however, branching was possible only at high mimetic loads, making it unclear whether it was an ecologically realistic outcome. The crucial point in both models is that a mimic that encounters different levels of discrimination may specialize on exploiting the least discriminating subset of signal receivers. This suggests that when other selection pressures oppose the evolution of accurate mimicry, the equilibrium level of mimicry may be sensitive both to the average and to the variation in discrimination level.

In general, we should expect mimics to encounter different levels of discrimination whenever different classes (or species) of signal receivers vary in factors such as perceptual acuity, the relative rate at which they encounter each signaler types, or in the exact costs they suffer for the two types of discrimination error. In Batesian mimicry, predators may adopt different attack strategies according to hunger state or the availability of alternative prey (e.g., Kokko et al. 2003; Sherratt 2003), and it would be interesting to investigate how this would influence the selection acting on the mimics. In avian brood parasitism, it has been suggested that the higher tolerance of dissimilar eggs among young hosts (e.g., Rothstein 1978; Lotem et al. 1992; Lotem et al. 1995) could contribute to a slowing down of the evolutionary arms race between host and parasite and weaken selection for better egg mimicry (Grim 2002). However, if the evolution of more accurate egg mimicry also entailed a cost to the parasite, a general increase in egg rejection among hosts could even reverse the selection for better egg mimicry, since the parasites might be better off by specializing on young hosts. To our knowledge, no study has been undertaken to investigate costs of egg mimicry; so this example is merely illustrative. However, it is not entirely implausible: smaller egg size is associated with lower offspring size at hatching and with lower survival in the early nestling stage in many birds (Williams 1994; Christians 2002), suggesting that good mimicry of egg size could potentially be costly.

If a mimetic trait is costly to display but behaviorally and/or phenotypically plastic, a mimic could potentially adjust its level of mimicry according to the response threshold it is likely to encounter, so that costs of accurate mimicry are only paid in the presence of certain external

cues. There are several examples of mimics that flexibly change their deceptive signals to mimic different models (Tarsitano et al. 2000; Norman et al. 2001; Haynes et al. 2002), but we are not familiar with any cases where mimetic accuracy has been demonstrated to vary between contexts to minimize signal costs. However, many forms of behavioral mimicry may be costly, such as the imitation of movement patterns and waving of forelegs to mimic antennae that are found in many ant-mimicking spiders and mimetic hoverflies (McIver and Stonedahl 1993; Gilbert 2005). It should be possible to test whether such behavioral mimicry is more likely to be performed in patch types in which the mimic-to-model ratio is high.

When signal receivers discriminate more strongly against mimics in the presence of an external cue, individual mimics will always benefit from encountering signal receivers that have not been alerted by the cue. If the external cue is emitted by the mimic, we may expect strong selection for suppression of the cue. Many avian brood parasites approach the host nest stealthily before laying; this reduces the chances of receiving host aggression and may be effective against subsequent flexible egg rejection strategies (Rothstein and Robinson 1998). If the cue is environmental, mimics may be under strong selection to shift their activity to circumstances where the cue is not present, which in effect could act to synchronize their activity temporally or spatially with those of model organisms (e.g., Uésugi 1996). However, due to physiological and ecological constraints, it could be quite difficult to suppress cues or converge to model activity patterns. The ecological theory of limiting similarity says that two species cannot coexist ecologically if their niches are too similar, suggesting that there usually would be some differences between models and mimics in resource and habitat use. Our model predicts that the response thresholds adopted by signal receivers in the presence and absence of the cue should differ more widely when mimetic accuracy is high, which suggests that selection for cue suppression or activity shifts may be stronger when mimetic accuracy is high. However, if the cue as a result becomes less informative, this may in turn weaken selection for mimetic accuracy. It is thus hard to predict how selection for mimetic accuracy will interact with selection for cue suppression/activity shifts; it will depend on which is the most labile trait. Finally, we found that cue utilization among signal receivers may benefit mimics at the population level. This suggests that kin selection could potentially oppose individual selection for cue suppression or activity synchronization in mimics.

There is some evidence for activity shifts in Batesian mimicry. A study of mimetic hoverflies found positive correlations between daily activity patterns of mimics and their hymenopteran models in 16 of 17 tested species, of

which nine were significant (Howarth et al. 2004). No effect of mimetic accuracy was found, but common mimics were more likely to have similar activity patterns to models than rare mimics were. Selection for temporal synchrony on such a short timescale is easily explained if predators have a very short memory or if they frequently sample the mimicry complex to track changes in relative abundance so that mimics are only protected by their resemblance for a short time after the predator attacks a model. However, it has been argued that birds retain learned aversions against models for much longer time periods (Waldbauer 1988). An equally plausible explanation is that predators pay attention to context when they discriminate between mimic and model prey, causing individual mimics to benefit from being encountered in contexts with lower mimic-to-model ratios, which could select for small-scale activity synchronization. It should be noted that evidence for synchrony on a seasonal scale is equivocal (Waldbauer 1988; Howarth and Edmunds 2000); perhaps seasonal activity patterns are more strongly constrained. Alternatively, it could be beneficial for mimics to emerge before most young birds have fledged, or after the defended models have emerged, to minimize the risk of encountering naive predators (Waldbauer 1988). The optimal level of synchrony may depend on the scale.

To our knowledge, few studies outside the field of avian brood parasitism have systematically looked at the role of external cues in mimetic discrimination. Our model suggests that use of external cues may have important implications for mimetic evolution. Many factors have been left out in our model (as in any model), and we see scope for some future extensions. For instance, one may add costs to acquiring information. It may not always pay to assess all available cues, since a predator that momentarily hesitates before attacking a prey may lose the opportunity. A more interesting extension to the model would perhaps be to include the evolution of model and mimic activity patterns, which would give us a better understanding of the degree of spatial and temporal synchrony we should expect in mimicry systems. It is not only mimics that might benefit from activity shifts; models could also be selected to avoid contexts with high mimic-to-model ratios.

### Conclusion

We have explored the ecological and evolutionary consequences of flexible (i.e., context-dependent) discrimination using aggressive mimicry and Batesian mimicry as examples. We found that depending on circumstances, flexible discrimination can be beneficial for models and detrimental to mimics, and perhaps more surprisingly, detrimental for both, or beneficial for both. If there are opposing selection pressures that favor less accurate mimicry,

flexible discrimination may qualitatively change the long-term outcome of mimetic evolution.

### Acknowledgments

We thank C. Brinch, N. Davies, M. Edmunds, A. Lotem, S. Rands, G.-P. Sætre, T. Slagsvold, N. C. Stenseth, and T. O. Sæviung for valuable discussions and helpful comments on various versions of the manuscript. We also thank two anonymous reviewers for their comments. Financial support was provided by the Norwegian Research Council (to Ø.H.H.).

### Literature Cited

- Alcock, J. 1970. Punishment levels and the response of white-throated sparrows (*Zonotrichia albicollis*) to three kinds of artificial models and mimics. *Animal Behaviour* 18:733–739.
- Alvarez, F. 1996. Model cuckoo *Cuculus canorus* eggs accepted by rufous bush chats *Cercotrichas galactotes* during the parasite's absence from the breeding area. *Ibis* 138:340–342.
- Bártol, I., Z. Karcza, C. Moskát, E. Røskaft, and T. Kisbenedek. 2002. Responses of great reed warblers *Acrocephalus arundinaceus* to experimental brood parasitism: the effects of a cuckoo *Cuculus canorus* dummy and egg mimicry. *Journal of Avian Biology* 33:420–425.
- Bouton, M. E. 1993. Context, time, and memory retrieval in the interference paradigms of Pavlovian learning. *Psychological Bulletin* 114:80–99.
- . 1994. Conditioning, remembering, and forgetting. *Journal of Experimental Psychology Animal Behavior Processes* 20:219–231.
- Brooke, M. de L., N. B. Davies, and D. G. Noble. 1998. Rapid decline of host defences in response to reduced cuckoo parasitism: behavioural flexibility of reed warblers in a changing world. *Proceedings of the Royal Society of London B* 265:1277–1282.
- Brower, J. V. Z. 1960. Experimental studies of mimicry. IV. The reactions of starlings to different proportions of models and mimics. *American Naturalist* 94:271–282.
- Christians, J. K. 2002. Avian egg size: variation within species and inflexibility within individuals. *Biological Reviews of the Cambridge Philosophical Society* 77:1–26.
- Cushing, P. E. 1997. Myrmecomorphy and myrmecophily in spiders: a review. *Florida Entomologist* 80:165–193.
- Dafni, A. 1984. Mimicry and deception in pollination. *Annual Review of Ecology and Systematics* 15:259–278.
- Davies, N. B. 2000. Cuckoos, cowbirds and other cheats. Poyser, London.
- Davies, N. B., and M. de L. Brooke. 1988. Cuckoos versus reed warblers: adaptations and counteradaptations. *Animal Behaviour* 36:262–284.
- Davies, N. B., S. H. M. Butchart, T. A. Burke, N. Chaline, and I. R. K. Stewart. 2003. Reed warblers guard against cuckoos and cuckoldry. *Animal Behaviour* 65:285–295.
- Day, T., and P. D. Taylor. 1998. Unifying genetic and game theoretic models of kin selection for continuous traits. *Journal of Theoretical Biology* 194:391–407.
- Duncan, C. J., and P. M. Sheppard. 1965. Sensory discrimination and its role in the evolution of Batesian mimicry. *Behaviour* 24:269–282.
- Egan, J. P. 1975. Signal detection theory and ROC analysis. Academic Press Series in Cognition and Perception. Academic Press, New York.
- Frank, S. A. 1998. Foundations of social evolution. Monographs in Behavior and Ecology. Princeton University Press, Princeton, NJ.
- Geritz, S. A. H., É. Kisdi, G. Meszéna, and J. A. J. Metz. 1998. Evolutionarily singular strategies and the adaptive growth and branching of the evolutionary tree. *Evolutionary Ecology* 12:35–57.
- Getty, T. 1985. Discriminability and the sigmoid functional response: how optimal foragers could stabilize model-mimic complexes. *American Naturalist* 125:239–256.
- Gilbert, F. 2005. The evolution of imperfect mimicry. Pages 231–288 in M. D. E. Fellowes, G. J. Holloway, and J. Rolff, eds. *Insect evolutionary ecology*. CABI, Wallingford.
- Goodale, M. A., and I. Sneddon. 1977. The effect of distastefulness of the model on the predation of artificial Batesian mimics. *Animal Behaviour* 25:660–665.
- Grim, T. 2002. Why is mimicry in cuckoo eggs sometimes so poor? *Journal of Avian Biology* 33:302–305.
- Haynes, K. F., C. Gemeno, K. V. Yeagan, J. G. Millar, and K. M. Johnson. 2002. Aggressive chemical mimicry of moth pheromones by a bolas spider: how does this specialist predator attract more than one species of prey? *Chemoeology* 12:99–105.
- Holen, Ø. H., and R. A. Johnstone. 2004. The evolution of mimicry under constraints. *American Naturalist* 164:598–613.
- Howarth, B., and M. Edmunds. 2000. The phenology of Syrphidae (Diptera): are they Batesian mimics of Hymenoptera? *Biological Journal of the Linnean Society* 71:437–457.
- Howarth, B., M. Edmunds, and F. Gilbert. 2004. Does the abundance of hoverfly (Syrphidae) mimics depend on the numbers of their hymenopteran models? *Evolution* 58:367–375.
- Jennions, M. D., and M. Petrie. 1997. Variation in mate choice and mating preferences: a review of causes and consequences. *Biological Reviews of the Cambridge Philosophical Society* 72:283–327.
- Johnstone, R. A. 2002. The evolution of inaccurate mimics. *Nature* 418:524–526.
- Kagel, J. H., L. Green, and T. Caraco. 1986. When foragers discount the future: constraint or adaptation? *Animal Behaviour* 34:271–283.
- Kokko, H., J. Mappes, and L. Lindström. 2003. Alternative prey can change model-mimic dynamics between parasitism and mutualism. *Ecology Letters* 6:1068–1076.
- Liebert, A. E., and P. T. Starks. 2004. The action component of recognition systems: a focus on the response. *Annales Zoologici Fennici* 41:747–764.
- Lindström, L., R. V. Alatalo, and J. Mappes. 1997. Imperfect Batesian mimicry: the effects of the frequency and the distastefulness of the model. *Proceedings of the Royal Society of London B* 264:149–153.
- Lotem, A., H. Nakamura, and A. Zahavi. 1992. Rejection of cuckoo eggs in relation to host age: a possible evolutionary equilibrium. *Behavioral Ecology* 3:128–132.
- . 1995. Constraints on egg discrimination and cuckoo-host co-evolution. *Animal Behaviour* 49:1185–1209.
- Marchetti, K. 1992. Costs to host defence and the persistence of parasitic cuckoos. *Proceedings of the Royal Society of London B* 248:41–45.

- McIver, J. D., and G. Stonedahl. 1993. Myrmecomorphy: morphological and behavioral mimicry of ants. *Annual Review of Entomology* 38:351–379.
- Moksnes, A., E. Røskoft, and L. Korsnes. 1993. Rejection of cuckoo (*Cuculus canorus*) eggs by meadow pipits (*Anthus pratensis*). *Behavioral Ecology* 4:120–127.
- Nilsson, L. A. 1992. Orchid pollination biology. *Trends in Ecology & Evolution* 7:255–259.
- Nonacs, P. 1985. Foraging in a dynamic mimicry complex. *American Naturalist* 126:165–180.
- Norman, M. D., J. Finn, and T. Tregenza. 2001. Dynamic mimicry in an Indo-Malayan octopus. *Proceedings of the Royal Society of London B* 268:1755–1758.
- Oaten, A., C. E. M. Pearce, and M. E. B. Smyth. 1975. Batesian mimicry and signal detection theory. *Bulletin of Mathematical Biology* 37:367–387.
- Ottenheim, M. M., B. Wertheim, G. J. Holloway, and P. M. Brakefield. 1999. Survival of colour-polymorphic *Eristalis arbustorum* hoverflies in semi-field conditions. *Functional Ecology* 13:72–77.
- Page, R. A., and M. J. Ryan. 2005. Flexibility in assessment of prey cues: frog-eating bats and frog calls. *Proceedings of the Royal Society of London B* 272:841–847.
- Reeve, H. K. 1989. The evolution of conspecific acceptance thresholds. *American Naturalist* 133:407–435.
- Robert, M., and G. Sorci. 1999. Rapid increase of host defence against brood parasites in a recently parasitized area: the case of village weavers in Hispaniola. *Proceedings of the Royal Society of London B* 266:941–946.
- Rodríguez-Gironés, M. A., and A. Lotem. 1999. How to detect a cuckoo egg: a signal-detection theory model for recognition and learning. *American Naturalist* 153:633–648.
- Rothstein, S. I. 1978. Mechanisms of avian egg-recognition: additional evidence for learned components. *Animal Behaviour* 26:671–677.
- Rothstein, S. I., and S. K. Robinson. 1998. The evolution and ecology of avian brood parasitism: an overview. Pages 3–56 in S. I. Rothstein, and S. K. Robinson, eds. *Parasitic birds and their hosts: studies in coevolution*. Oxford Ornithology Series. Oxford University Press, New York.
- Ruxton, G. D., T. N. Sherratt, and M. P. Speed. 2004. *Avoiding attack: the evolutionary ecology of crypsis, warning signals, and mimicry*. Oxford University Press, Oxford.
- Schiestl, F. P., R. Peakall, J. G. Mant, F. Ibarra, C. Schulz, S. Franke, and W. Francke. 2003. The chemistry of sexual deception in an orchid-wasp pollination system. *Science* 302:437–438.
- Sherman, P. W., H. K. Reeve, and D. W. Pfennig. 1997. Recognition systems. Pages 69–96 in J. R. Krebs, and N. B. Davies, eds. *Behavioural ecology: an evolutionary approach*. Blackwell Science, Oxford.
- Sherratt, T. N. 2002. The evolution of imperfect mimicry. *Behavioral Ecology* 13:821–826.
- . 2003. State-dependent risk-taking by predators in systems with defended prey. *Oikos* 103:93–100.
- Srygley, R. B. 2004. The aerodynamic costs of warning signals in palatable mimetic butterflies and their distasteful models. *Proceedings of the Royal Society of London B* 271:589–594.
- Stephens, D. W. 2002. Discrimination, discounting and impulsivity: a role for an informational constraint. *Philosophical Transactions of the Royal Society of London B* 357:1527–1537.
- Stephens, D. W., and D. Anderson. 2001. The adaptive value of preference for immediacy: when shortsighted rules have farsighted consequences. *Behavioral Ecology* 12:330–339.
- Tarsitano, M., R. R. Jackson, and W. H. Kirchner. 2000. Signals and signal choices made by the araneophagic jumping spider *Portia fimbriata* while hunting the orb-weaving web spiders *Zygiella x-notata* and *Zosis geniculatus*. *Ethology* 106:595–615.
- Taylor, P. D., and S. A. Frank. 1996. How to make a kin selection model. *Journal of Theoretical Biology* 180:27–37.
- Tollrian, R., and C. D. Harvell, eds. 1999. *The ecology and evolution of inducible defences*. Princeton University Press, Princeton, NJ.
- Uésugi, K. 1996. The adaptive significance of Batesian mimicry in the swallowtail butterfly, *Papilio polytes* (Insecta, Papilionidae): associative learning in a predator. *Ethology* 102:762–775.
- Waldbauer, G. P. 1988. Asynchrony between Batesian mimics and their models. *American Naturalist* 131(suppl.):S103–S121.
- Welbergen, J., J. Komdeur, R. Kats, and M. Berg. 2001. Egg discrimination in the Australian reed warbler (*Acrocephalus australis*): rejection response toward model and conspecific eggs depending on timing and mode of artificial parasitism. *Behavioral Ecology* 12:8–15.
- West, S. A., I. Pen, and A. S. Griffin. 2002. Cooperation and competition between relatives. *Science* 296:72–75.
- Wiley, R. H. 1994. Errors, exaggeration, and deception in animal communication. Pages 157–189 in L. A. Real, ed. *Behavioral mechanisms in evolutionary ecology*. University of Chicago Press, Chicago.
- Williams, T. D. 1994. Intraspecific variation in egg size and egg composition in birds: effects on offspring fitness. *Biological Reviews of the Cambridge Philosophical Society* 68:35–59.
- Wong, B. B. M., and F. P. Schiestl. 2002. How an orchid harms its pollinator. *Proceedings of the Royal Society of London B* 269:1529–1532.
- Wong, B. B. M., C. Salzmann, and F. P. Schiestl. 2004. Pollinator attractiveness increases with distance from flowering orchids. *Proceedings of the Royal Society of London B* 271(suppl.):S212–S214.

Associate Editor: Peter D. Taylor  
 Editor: Michael C. Whitlock