The Evolution of Mimicry under Constraints

Øistein Haugsten Holen[1,*] **and Rufus A. Johnstone**[2,†]

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 March 24, 2004; Accepted July 21, 2004;
Electronically published September 29, 2004

Online enhancement: appendix.

ABSTRACT: The resemblance between mimetic organisms and their models varies from near perfect to very crude. One possible explanation, which has received surprisingly little attention, is that evolution can improve mimicry only at some cost to the mimetic organism. In this article, an evolutionary game theory model of mimicry is presented that incorporates such constraints. The model generates novel and testable predictions. First, Batesian mimics that are very common and/or mimic very weakly defended models should evolve either inaccurate mimicry (by stabilizing selection) or mimetic polymorphism. Second, Batesian mimics that are very common and/or mimic very weakly defended models are more likely to evolve mimetic polymorphism if they encounter predators at high rates and/or are bad at evading predator attacks. The model also examines how cognitive constraints acting on signal receivers may help determine evolutionarily stable levels of mimicry. Surprisingly, improved discrimination abilities among signal receivers may sometimes select for less accurate mimicry.

Keywords: arms race, evolutionarily stable strategy (ESS), signal detection, polymorphism, aggressive mimicry, Batesian mimicry.

Many predators and parasites use aggressive mimicry to deceive their victims. Some examples are avian brood parasites that lay mimetic eggs in the nests of other birds (Wickler 1968; Rothstein and Robinson 1998*b*; Davies 2000), spiders that vibrate the webs of others in imitation of captured prey, thus luring the owners toward them to be consumed (Jackson and Pollard 1996; Tarsitano et al. 2000), and the *Ophrys* orchids, which have rewardless flowers that mimic female wasps and bees in order to attract the males as pollinators (Kullenberg 1961; Wickler 1968; Nilsson 1992). All of these aggressive mimics gain by exploiting behavioral patterns that the victims usually direct toward the model.

Prey organisms may also use mimicry in order to deceive their predators. In classical Batesian mimicry, palatable prey mimic poisonous or otherwise defended model prey in order to deter attack. In this group we find hoverflies that mimic wasps and bees, palatable butterflies that mimic unpalatable butterflies (e.g., Wickler 1968), and many myrmecomorphic (ant-mimicking) spiders (McIver and Stonedahl 1993; Cushing 1997).

Systems of aggressive mimicry and Batesian mimicry may be understood as evolutionary arms races (Dawkins and Krebs 1979) in which signal receivers (operators) are under selection to improve discrimination between models and mimics and the mimics are under selection to appear more similar to the models and thus more difficult to discriminate against. Note that classical Müllerian mimicry is quite different because both model and mimic prey are defended and benefit from sharing a warning signal, and there is no selection on the predators to discriminate between models and mimics.

The degree to which mimetic organisms resemble their models varies quite a lot. Among avian brood parasites, the level of egg mimicry ranges from near perfect to nonmimetic (Wickler 1968; Rothstein and Robinson 1998*a*; Davies 2000). Many common hoverflies that are Batesian mimics show only a crude resemblance to their models, while other, less common hoverflies resemble their models very closely (Azmeh et al. 1998; Edmunds 2000; Howarth and Edmunds 2000; Gilbert in press). Edmunds (2000) suggests that a similar pattern may be found among myrmecomorphic spiders.

Several hypotheses have been put forward to explain the existence of inaccurate mimics. For instance, it has been proposed that inaccurate mimics may be "general mimics" that resemble several model species crudely but none very closely (e.g., Edmunds 2000; Pekár and Král 2002; Sherratt 2002), or that mimics may be kept from evolving accurate mimicry by kin selection (Johnstone 2002). Other expla-

* E-mail: o.h.holen@bio.uio.no.

† E-mail: raj1003@hermes.cam.ac.uk.

Am. Nat. 2004. Vol. 164, pp. 598–613. 0003-0147/2004/16405-40367$15.00.

nations have focused on the cognitive apparatus of the operator: Dittrich et al. (1993) suggested that what is only a crude resemblance to human observers may appear quite convincing to a predator. Moreover, the combination of cognitive constraints and environmental noise can make it impossible for an operator to simultaneously maximize its chances of correctly identifying both model signals and mimic signals, and a trade-off between the two may exist. Thus, if it is very costly to erroneously classify a model as a mimic, and there is only a small benefit to be gained from correctly identifying a mimic, operators may benefit from being susceptible to deceit by mimics (Wiley 1994), which again may weaken selection for accurate mimicry. For instance, predators may avoid attacking inaccurate Batesian mimics if they resemble very noxious models (e.g., Edmunds 2000; Holloway et al. 2002; Sherratt 2002). Similarly, it could benefit male bees and wasps to be susceptible to deceit by sexually deceptive orchids if this reduced their chance of erroneously rejecting real females; being duped by an orchid may only cost some time and energy (Kullenberg 1961; Peakall 1990), while the cost of missing a real mating opportunity could be greater. If discrimination per se carries a cost, for instance, related to increased vigilance or to the maintenance of a specialized cognitive apparatus, then operators with only a little to gain from discrimination may be selected to give up discrimination entirely. Thus, there will be no selection for accurate mimicry.

The evolutionary lag hypothesis provides a nonadaptive explanation for inaccurate mimicry (e.g., Rothstein and Robinson 1998*a*; Edmunds 2000; Holloway et al. 2002). It suggests that the mimics simply lag behind in the evolutionary arms race; more accurate mimicry is adaptive but has not yet appeared and increased in frequency in the population due to lack of suitable genetic variation or time for natural selection to work. The amount of evolutionary lag may also be affected by evolution of the model population. In systems of Batesian mimicry, the presence of mimics is generally thought to weaken the protection from predation given by the model signal. Therefore, the model organisms may in fact be under selection to become more easily distinguishable from the mimics, giving rise to a separate coevolutionary race (Vane-Wright 1976; Gavrilets and Hastings 1998; Holmgren and Enquist 1998). However, if models benefit from the presence of mimics, as when an aggressive mimic simulates the signal properties of a prey item that the operator consumes, the model may in fact be under selection to become more similar to the mimic (Vane-Wright 1976).

Typically, predictions about equilibrium levels of mimicry are made under the implicit assumption that evolution may improve mimicry at virtually no cost, that is, without affecting other activities or functions of the mimic in a negative way (but see Servedio and Lande 2003). This may be a realistic assumption in some cases, as when mimics trick operators by using emitted signals of short duration, such as the light flashes used by *Photuris* fireflies (e.g., Haynes and Yeargan 1999) or the mechanosensory stimuli used by web-invading spiders. Although the mimic might pay an energetic cost per emitted signal, and there may be costs related to the maintenance of the signaling apparatus, the difference in costs between a well-tuned mimetic signal and a badly matched signal may still be negligible. It is hard to see why a small evolutionary tweak of, say, the temporal properties of a signal must be costly to the mimic even if it could greatly improve the chances of deceiving the operator.

However, there are also many ways in which traits that improve mimicry may confer significant costs on their carriers. Selection for mimicry can entail radical phenotypic changes in the mimic, including changes in body size and shape, modifications and loss of body parts, and changes in color, scent, and behavior (e.g., Wickler 1968; McIver and Stonedahl 1993), all of which may potentially interfere with other activities or functions of the mimetic organism. For instance, selection for mimicry has lead to a reduction or loss of wings in many myrmecomorphic insects and to a more narrow body shape and reduced fecundity in many myrmecomorphic spiders (McIver and Stonedahl 1993; Cushing 1997). Mimetic traits may potentially reduce mating success, as in butterflies, where color patterns are important both for courtship and mimicry (e.g., Turner 1978). Eggs of the common cuckoo *Cuculus canorus* often closely match the colorations and markings of host eggs but nevertheless tend to be larger than host eggs (Davies and Brooke 1988; Rothstein 1990). Good mimicry of egg size may be costly to the parasite, because reductions of egg size in birds may reduce hatching probability, size at hatching, and survival in the early nestling stage (Wilson 1991; Williams 1994). In small organisms with a high surface-to-volume ratio, changes in body color and brightness may significantly affect the absorption of heat radiation from the environment. Body color and brightness have been found to influence thermoregulation in spiders (Oxford and Gillespie 1998), butterflies (Shreeve 1992), beetles (Brakefield 1985), and hoverflies (Holloway et al. 1997; Ottenheim et al. 1999), all of which are groups well known for containing many mimetic species.

Costs of mimicry are not necessarily restricted to changes in morphological traits. Behavioral mimicry may require a mimic to spend more time performing certain tasks than it otherwise would. For instance, some hoverflies mimic the flight behavior of their models while foraging, possibly compromising their superior flight abilities (Golding and Edmunds 2000; Golding et al. 2001).

Myrmecomorphic spiders often walk with a more erratic, ant-like gait, which they abandon when disturbed, probably in order to escape predators after being discovered (Cushing 1997).

In this article, evolutionary game theory (Maynard Smith 1982) is used to explore the equilibrium level of mimetic resemblance in several different contexts (including both aggressive mimicry and Batesian mimicry), assuming that closer similarity to the model entails fitness costs for the mimic.

The Model

We present a general model of mimetic evolution applicable to both aggressive and Batesian mimicry. Several different model variants will be considered for each type of mimicry, characterized by different assumptions about how the fitness of a mimic depends on its ability to deceive signal receivers (operators). We start by outlining the basic assumptions of the model that are common to all variants.

The discrimination task of the operators is modeled using signal detection theory (e.g., Egan 1975; Wiley 1994). Operators encounter models and mimics at fixed relative frequencies $(1 - P)$ and P, respectively, and attempt to discriminate between them based on a continuous trait such as size, width, or brightness. Perception is assumed to be imperfect due to environmental noise and limits to the cognitive abilities (e.g., resolution of the sense organs) of the operators. Formally, it is assumed that operator estimates of individual model and mimic trait values follow a normal distribution centered on the true values with standard deviation σ, which represents the degree of perceptual error. The operators classify all individuals with perceived trait values higher than some threshold t as mimics and the rest as models (fig. 1). Setting a high threshold increases the chance of correctly classifying models but reduces the chance of correctly classifying mimics. A low threshold has the opposite effect (fig. 1).

The trait value of the model organisms is standardized to 0 (with no loss of generality) and is assumed fixed in the model. The latter assumption, although common in models of mimicry (e.g., Rodríguez-Gironés and Lotem 1999; Johnstone 2002; Sherratt 2002), is simplistic. Nevertheless, there is some biological justification for it; the model organism may be constrained by its need to be correctly identified by operators (Nur 1970).

Mimic trait values (denoted by m) are by convention positive and are measured relative to the trait value of the models. Thus, the closer a mimic trait value is to 0, the better the mimicry is. However, the evolution of better mimicry is assumed to interfere with other functions of the mimetic organism; a more accurate mimic will be better at deception but will have reduced fitness in other contexts. The exact relation between mimetic trait and cost will become clear later.

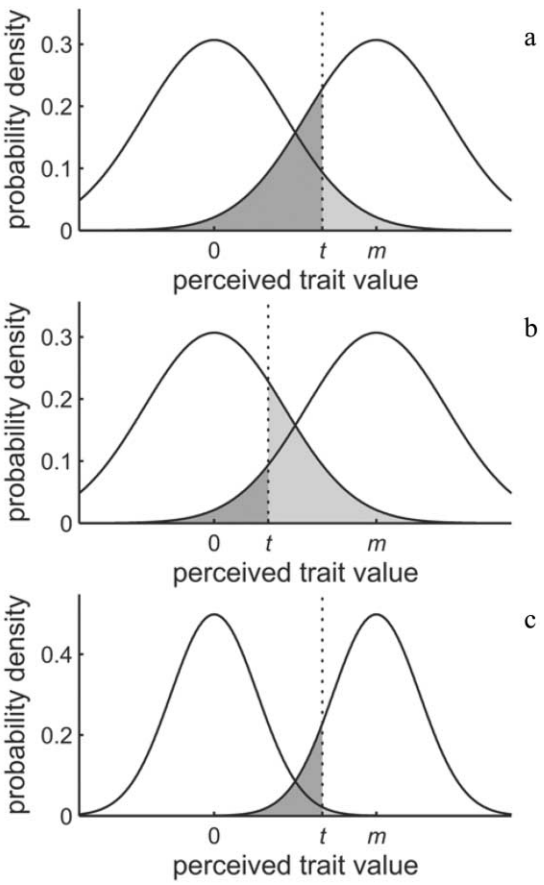

Figure 1: Operators estimate the trait values of encountered models and mimics and classify all those with perceived trait values above their response threshold t as mimics and all below as models. The true trait value of the models is standardized to 0, and the true trait value of the mimics is denoted m. Mimic trait values closer to 0 represent more accurate mimicry. Trait value estimates follow a normal distribution centered on the true value (as shown), with a standard deviation σ that reflects perceptual errors. The dark gray areas give the probability of erroneously classifying a mimic as a model, while the light gray areas give the probability of erroneously classifying a model as a mimic. As illustrated by *a* and *b*, setting a lower threshold increases the probability of correctly identifying mimics but simultaneously decreases the probability of correctly identifying models. As illustrated by *a* and *c*, if operators evolve a better cognitive machinery (i.e., if σ decreases), they may be able to increase the probability of correctly identifying both models and mimics.

The mimic trait value is assumed to be determined by genotype; thus, the population level of mimicry changes on an evolutionary timescale. In contrast, operators are assumed to adjust their discrimination strategies on a much faster timescale; they adjust their thresholds behaviorally in an optimal manner given the currently typical

level of mimicry. Such strong separation of timescales is appropriate when operator discrimination is based on learning from direct encounters. However, if predatory mimics are very efficient at capturing duped victims or Batesian mimics simulate models that are extremely dangerous to a predator (e.g., poisonous coral snakes, Wickler 1968), learning may be difficult, and innate recognition may evolve instead (e.g., Smith 1975), which falls outside the scope of this article.

Operator Fitness Functions: Aggressive versus Batesian Mimicry

Our model variants fall into two main categories, aggressive mimicry and Batesian mimicry, which are characterized by different operator fitness functions.

Consider, first, the case of aggressive mimicry. Aggressive mimics take advantage of operators by simulating model organisms that the operators benefit from interacting with. Hence, we assume that interacting with a model organism yields a benefit B to the operator and that interacting with a mimic entails a fitness cost C. If a model or mimic is rejected, the payoff is 0. The operator interacts with all models and mimics whose perceived trait value falls below t, and the others are rejected. The probabilities that a model and that a mimic are perceived to have a trait value below t are equal to $Z(t)$ and $Z(t - m)$, respectively, where Z denotes the cumulative distribution of $N(0, \sigma^2)$. We follow earlier approaches (e.g., Greenwood 1986; Johnstone 2002; Sherratt 2002) and assume that the operator maximizes expected fitness gain per encounter. This may be a reasonable assumption when search time is long relative to handling time or when future benefits are strongly discounted (shortsighted decision rules may also be more psychologically realistic; Stephens [2002] suggests that shortsighted rules may be optimal when animals suffer information-processing constraints). The expected fitness gain per encounter (our measure of fitness) is then given by

$$w_{\text{op}} = (1 - P)Z(t)B - PZ(t - m)C. \tag{1a}$$

By setting $\partial w_{\text{op}}/\partial t = 0$, checking second-order conditions, and introducing the "mimetic load" K, the optimal threshold t^* as a function of mimic trait value m is found as

$$t^* = \frac{m}{2} - \frac{\ln (K)\sigma^2}{m}, \text{ where } K = \frac{PC}{(1 - P)B}. \tag{1b}$$

The mimetic load is 0 in the absence of mimics and increases with mimic encounter rate P and the cost C of responding to mimics. It may be interpreted as follows: if we think of model organisms as emitting—possibly unintentionally—a signal that informs operators of the benefits to be gained by responding, then the mimetic load represents the extent to which the operators should distrust that information.

Now consider the case of Batesian mimicry. Batesian mimics simulate noxious prey in order to deter attack by predators. We thus assume that attacking a mimic yields a fitness benefit B to the operator and attacking a model yields a fitness cost C, while ignoring an encountered prey item yields a zero payoff. The predator attacks all model and mimic prey with a perceived trait value $>t$ (fig. 1), and the rest are left in peace. The probabilities of attacking encountered model prey and mimic prey are, respectively, $1 - Z(t)$ and $1 - Z(t - m)$. The expected fitness gain to the operator per prey encounter is

$$w_{\text{op}} = P[1 - Z(t - m)]B - (1 - P)[1 - Z(t)]C. \tag{2a}$$

The optimal threshold t^* as a function of mimic trait value m can now be found as

$$t^* = \frac{m}{2} - \frac{\ln (K)\sigma^2}{m}, \text{ where } K = \frac{PB}{(1 - P)C}. \tag{2b}$$

This equation is identical to (1b) except that B and C have exchanged places in the expression for the mimetic load K. Here, K increases with the mimic encounter rate P and the benefit B of attacking mimics and decreases when the cost of attacking models increases. Again, the mimetic load may be understood as the extent to which the information provided by the model's signal (here a warning signal) should be trusted.

Mimic Fitness Functions

In all model variants, constraints on the evolution of mimicry are incorporated in the following way. A function $S(m)$ gives the fraction of the maximum possible fitness gain, given by the constant b, that can be obtained by a mimic with trait value m (fig. 2). When the mimic trait value equals m_{opt}, which, for instance, could be a physiological optimum, the mimic may obtain the full fitness gain b (i.e., $S(m_{\text{opt}}) = 1$). Formally, the trait value m_{opt} is a unique global maximum point of $S(m)$, and $S(m)$ falls off on each side of m_{opt}, enabling only smaller portions of the benefit b to be realized. For trait values lower than or equal to some threshold $m_{\min}$, no benefit at all may be realized. It is assumed that $S(m)$ is differentiable on $(m_{\min}, \infty)$ and log concave on $(m_{\min}, m_{\text{opt}})$. Furthermore, in our exploration of the model, we have restricted the range of parameters so that $-2\sigma \leq m_{\min} < m_{\text{opt}} \leq 6\sigma$ and $(m_{\text{opt}} - m_{\min}) \leq 6\sigma$.

Additional assumptions about the mimic fitness function differ between model variants. We consider three ver-

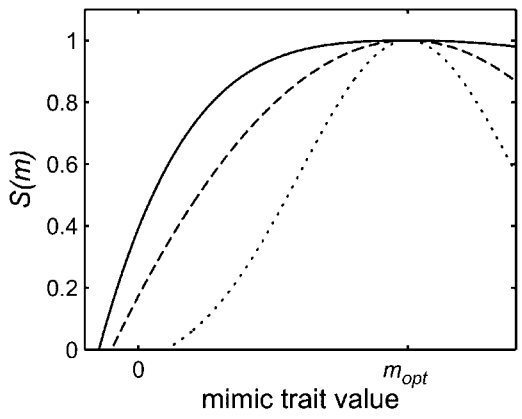

Figure 2: The mimics are assumed to pay a fitness cost that increases the more closely the trait value of the models is mimicked. Mimic trait values closer to 0 represent more accurate mimicry. In the absence of operator discrimination, the trait value m_{opt} is optimal for the mimics. The function $S(m)$ represents the fraction of some benefit b that may be obtained by a mimic with trait value m. If $m < m_{min}$, no benefit may be obtained. Some possible shapes of the function $S(m)$ are shown, and m_{min} may be either smaller or larger than the model trait value (which is 0). The dashed line shows a second-order polynomial; we have used second-order polynomials for the function $S(m)$ in figures 3–6.

sions of the model that deal with aggressive mimicry: Scenario A1 (multiple victims) focuses on an obligate aggressive mimic that obtains an additive fitness benefit from each operator it deceives. The resources obtained through deception are necessary for survival and/or reproduction and cannot be gained by other means. Scenario A2 (single victim) focuses on an obligate aggressive mimic that need only trick one victim before some fixed period of time has passed in order to obtain a fitness benefit. The average number of encounters with operators in the period is denoted n. This scenario is an illustrative contrast to scenario A1 and is motivated by the fact that in many cases it would be unrealistic to assume that the total fitness gain obtained increases proportionally with the number of victims deceived. For example, the male function (i.e., pollinia removal) of the flower of a deceptive orchid may be completed after one or two pollinator visits (e.g., O'Connell and Johnston 1998; Johnson and Edwards 2000); similarly, the orchid may gain little from receiving pollen more than once on a single flower (Nilsson 1992; O'Connell and Johnston 1998; Schiestl and Ayasse 2001). Finally, scenario A3 (facultative) focuses on a facultative aggressive mimic that faces the same challenge as in scenario A1 but that is sometimes able to obtain fitness benefits without deceiving victims. For instance, it may obtain the benefit by force in a proportion r of encounters, or it may exploit a second victim species that does not discriminate between models and mimics and which has a relative abundance equal to r (scenario A1 is thus a special case of A3, with r set equal to 0; we shall treat the two separately, however, because they yield very different results).

We consider two versions of the model that deal with Batesian mimicry. Scenario B1 (single reproductive bout) focuses on a mimic that must avoid being killed by predators for some fixed period of time before it obtains a fitness benefit. This may apply to mimics that must survive until the end of the season or until reaching another life stage before starting to reproduce. The average number of predators encountered in the period is denoted n. Scenario B2 (multiple reproductive bouts) focuses on a mimic that gains time-discounted benefits as long as it continues to escape predators (the discount factor is denoted λ). This may apply to mimics that reproduce in many small bouts under continuous risk of predation. In both scenarios, a mimic has a probability $(1 - q)$ of evading capture even if it fails to dupe a predator.

The precise mimic fitness functions for each scenario are summarized in table 1; further explanation of the derivations are given in the appendix in the online edition of the *American Naturalist*.

Analysis

First, consider the optimal discrimination strategy of the operators. The optimal rejection threshold t^* is given by (1b) and (2b) for aggressive and Batesian mimicry, respectively. For any fixed level of mimetic resemblance, the threshold t^* decreases as mimetic load K increases. More-

Table 1: Mimic fitness functions

Scenario	Function
Aggressive mimicry:	
A1 (obligate aggressive mimicry, multiple victims)	$w_{mA1} = Z(t-m)S(m)b$
A2 (obligate aggressive mimicry, single victim)	$w_{mA2} = [1 - e^{-nZ(t-m)}]S(m)b$
A3 (facultative aggressive mimicry)	$w_{mA3} = (1-r)Z(t-m)S(m)b + rS(m)b$
Batesian mimicry:	
B1 (single reproductive bout)	$w_{mB1} = e^{-nq[1-Z(t-m)]}S(m)b$
B2 (many reproductive bouts)	$w_{mB2} = \frac{S(m)b}{1-\lambda+\lambda q[1-Z(t-m)]}$

over, if $K>1$, then t^* will also decrease if mimicry improves. If $K<1$, then t^* will decrease as mimicry improves until $m = [-2\sigma^2 \ln(K)]^{1/2}$, below which t^* will start to increase as mimicry improves further.

Now consider the mimics. A mimic trait value is an evolutionarily stable strategy (ESS) if it cannot be invaded by a mutant with another trait value when adopted by all the members of a population (Maynard Smith 1982). We assume that a mutant will have a negligible effect on the optimal operator threshold t^* when rare. Thus, for a mimic trait m^* to be an ESS, it is necessary (but not sufficient) that it satisfies

$$\frac{\partial w_m}{\partial m}(m, t^*(m^*))\bigg|_{m=m^*} = 0, \tag{3}$$

where w_m is the mimic fitness function in question (table 1). Any strategy pair (m^*, t^*) that simultaneously solves $\partial w_m/\partial m = 0$ and $\partial w_{op}/\partial t = 0$ also satisfies equation (3). In order to be an ESS for a given mimetic load, a mimic trait m^* must both satisfy equation (3) and be a global maximum point of $w_m(m)$ when operators set the threshold $t^*(m^*)$. Since $S(m) = 0$ for all $m \le m_{min}$, $S(m)$ decreases for $m > m_{opt}$, and $Z(t - m)$ strictly decreases for all m, it is easy to see from the fitness functions that any mimic ESS m^* must satisfy $m_{min} < m^* < m_{opt}$.

In the following analyses, we will identify the endpoints of the evolutionary process (i.e., the mimic ESSs) for the different model variants. We will check evolutionary attainability by assuming adaptive dynamics, that is, that the rate of change in a trait is proportional to the slope of the associated fitness function (Hofbauer and Sigmund 1998). In the (t, m) plane, we will refer to $\partial w_{op}/\partial t = 0$ as the operator nullcline (i.e., t nullcline) and $\partial w_m/\partial m = 0$ as the mimic nullcline (i.e., m nullcline) because they define where the rate of change in the relevant trait is 0. Because operators are assumed to adjust their threshold optimally on a very fast timescale, we can simplify the dynamics and assume that the pair of population strategies (t, m) will always be positioned on the operator nullcline. Assuming that mutations in the mimic population have only small phenotypic effects and that K is constant, we may think of the populations of operators and mimics as moving along the operator nullcline, with the operators constantly adjusting their thresholds as the mimic trait value evolves.

For mathematical convenience, obligate aggressive mimicry (scenarios A1 and A2) will now be analyzed separately from facultative aggressive mimicry and Batesian mimicry (scenarios A3, B1, and B2).

Scenarios A1 and A2: Obligate Aggressive Mimicry. In

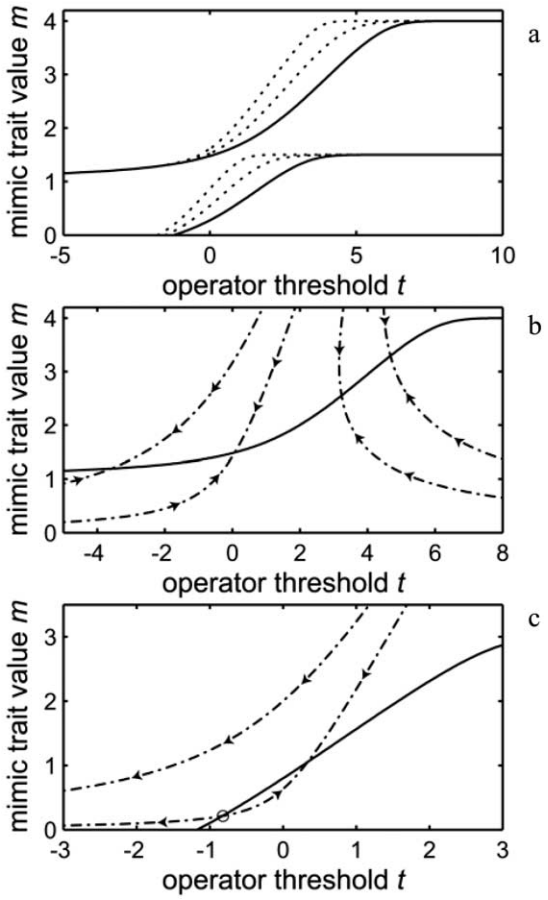

Figure 3: Mimic nullclines (*a–c*) and operator nullclines (*b*, *c*) are shown for different choices of parameters, detailed below. The mimic nullclines give the optimal trait value m^* corresponding to each operator threshold t, while the operator nullclines give the optimal threshold t^* corresponding to each mimic trait value m. *a*, Mimic nullclines in scenario A1 (*solid lines*) and scenario A2 (*dotted lines*) are shown for comparison. *b*, *c*, Operator nullclines (*dashed-dotted lines*) are shown superimposed on mimic nullclines (*solid lines*). Due to strong separation of timescales (evolutionary versus behavioral timescale), the dynamics are simplified, and the trajectories follow the operator nullcline (see main text for further explanation). The mimic trait values at the intersections of operator and mimic nullclines are monomorphic evolutionarily stable strategies, and all except the one marked with a circle in *c* are also attainable by evolution. Parameter values: $\sigma = 1$. *a*, *Solid lines* (scenario A1), top to bottom: $m_{min} = 1$, $m_{opt} = 4$; $m_{min} = -0.5$, $m_{opt} = 1.5$. *Dotted lines* (scenario A2), top to bottom: $m_{min} = 1$, $m_{opt} = 4$, $n = 12$; $m_{min} = 1$, $m_{opt} = 4$, $n = 4$; $m_{min} = -0.5$, $m_{opt} = 1.5$, $n = 12$; $m_{min} = -0.5$, $m_{opt} = 1.5$, $n = 4$. *b*, Scenario A1: $m_{min} = 1$, $m_{opt} = 4$. *Dash-dotted lines*, left to right: $\ln(K) = 5$, $\ln(K) = 1$, $\ln(K) = -5$, $\ln(K) = -10$. *c*, Scenario A2: $m_{min} = -1$, $m_{opt} = 3$, $n = 10$. *Dash-dotted lines*, left to right:$\ln(K) = 2$, $\ln(K) = 0.2$.

scenarios A1 and A2, the mimic nullclines (i.e., $\partial w_{m\mathrm{A1}}/\partial m = 0$ and $\partial w_{m\mathrm{A2}}/\partial m = 0$) give the unique optimal mimic strategy m^* for each operator threshold t, and it can be shown that m^* in both scenarios strictly increases with t (see appendix). Figure 3*a* plots m^* as a function of t for some illustrative parameter values. As we should expect, as n tends to 0, the nullcline $\partial w_{m\mathrm{A2}}/\partial m = 0$ becomes increasingly similar to the nullcline $\partial w_{m\mathrm{A1}}/\partial m = 0$, implying that the optimal mimic strategies for the two scenarios converge. This makes sense because the mimics in scenario A2 ("single victim") are under increasingly strong selection pressure to trick each operator as n decreases and encounters with operators become rarer.

In figure 3*b* and 3*c*, operator nullclines $\partial w_{\mathrm{op}}/\partial t = 0$ given by equation (1b) are plotted for several different mimetic loads K superimposed on a plot of mimic nullclines $\partial w_{m\mathrm{A1}}/\partial m = 0$ (fig. 3*b*) and $\partial w_{m\mathrm{A2}}/\partial m = 0$ (fig. 3*c*). Each intersection of the nullclines represents a monomorphic ESS for the mimic population at the corresponding mimetic load. Above the mimic nullcline, there is selection for lower mimic trait values (more accurate mimicry), so the populations will move downward along the operator nullcline. Below the mimic nullcline there is selection for higher trait values (less accurate mimicry), and the populations will move upward along the operator nullcline. Thus, all the ESSs in fig. 3*b* and 3*c* are attainable by natural selection except the one marked by a circle. When $m_{\min} < 0$ and $K > 1$, there will often be one such unattainable ESS in addition to an attainable ESS. However, if the mimic and operator nullcline do not intersect for any positive m, the mimics will instead race toward perfect mimicry, and the operators will respond by lowering their thresholds toward minus infinity. The biological relevance and interpretation of this outcome will depend on the biological context: For instance, if operators cannot raise offspring without interacting with models, the population will crash if t becomes too low.

Scenarios A3, B1, and B2: Facultative Aggressive Mimicry and Batesian Mimicry. In scenarios A3, B1, and B2, the mimic nullclines (i.e., $\partial w_{m\mathrm{A3}}/\partial m = 0$, $\partial w_{m\mathrm{B1}}/\partial m = 0$, and $\partial w_{\mathrm{m}B2}/\partial m = 0$) may represent both local fitness maxima and local fitness minima. It can be shown that there is a unique optimal mimic strategy m^* close to m_{opt} if the operator threshold t is sufficiently high or low (see appendix). The level of mimicry is most accurate when intermediate operator thresholds are adopted.

We will refer to the expressions $(1 - r)/r$ (scenario A3), nq (scenario B1), and $q\lambda/(1 - \lambda)$ (scenario B2) as the "incentive for deception," because for any fixed operator threshold, the optimal level of mimicry is more accurate when these expressions are high. This is illustrated in figure 4, where the mimic nullclines $\partial w_{m\mathrm{A3}}/\partial m = 0$, $\partial w_{m\mathrm{B1}}/\partial m = 0$, and $\partial w_{m\mathrm{B2}}/\partial m = 0$ are plotted in the (t, m) plane for some different parameter values. The incentive for deception is a measure of how tightly the mimics are locked into the ecological interaction with the operators and plays a key role in determining the qualitative behavior of the model. When the incentive for deception is low, the qualitative results are easily summarized. There is a unique optimal mimic strategy corresponding to each operator threshold (fig. 4), and each intersection of mimic and operator nullclines corresponds to a monomorphic ESS, as in scenarios A1 and A2.

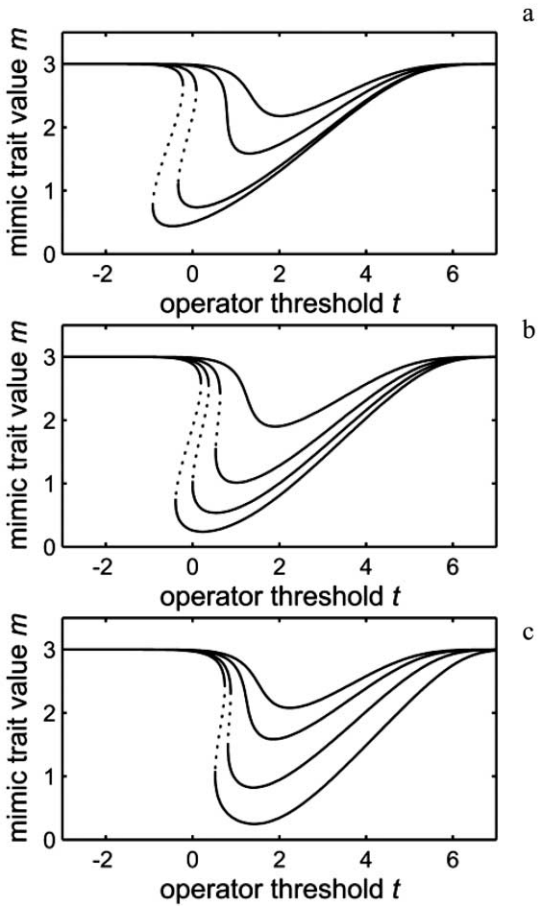

Figure 4: The mimic nullclines of scenarios A3, B1, and B2 are shown for different parameter choices. The solid line sections represent mimic trait values m^* that (locally) maximize fitness at different operator thresholds t, while dotted sections represent trait values that (locally) minimize fitness. Selection for accurate mimicry is strongest at intermediate rejection thresholds and when the "incentive for deception" is high. Parameter values: $\sigma = 1$, $m_{\min} = -0.5$, $m_{\mathrm{opt}} = 3$. *a*, Scenario A3, top to bottom: $(1 - r)/r = 3/7$, $(1 - r)/r = 1$, $(1 - r)/r = 4$, $(1 - r)/r = 9$. *b*, Scenario B1, top to bottom: $nq = 0.5$, $nq = 1.2$, $nq = 2$, $nq = 3$. *c*, Scenario B2, top to bottom: $q\lambda/(1 - \lambda) = 0.5$, $q\lambda/(1 - \lambda) = 1$, $q\lambda/(1 - \lambda) = 10/3$, $q\lambda/(1 - \lambda) = 20$.

A high incentive for deception, by contrast, may lead to an S-shaped mimic nullcline, which for intermediate operator thresholds has two local fitness maxima with a local fitness minimum somewhere between them (fig. 4). Such S-shaped mimic nullclines may allow for dimorphic ESSs (featuring two morphs that differ in their mimetic accuracy), which will be investigated in the next section. (Exactly how strong the incentive for deception must be before the mimic nullcline becomes S-shaped depends on the function $S(m)$. If $m_{\min}$ is close to m_{opt}, so that fitness falls off very fast as mimicry improves, then the mimic nullcline becomes S-shaped only at a very high incentive for deception, or it may never do so; scenario B2 is particularly sensitive to this effect. If the distance between $m_{\min}$ and m_{opt} is greater, then an S-shaped nullcline arises more easily.) Our implementation of signal discrimination in the model presupposes positive mimic trait values; thus, our explorations of dimorphic ESSs are restricted to S-shaped nullclines that feature only positive mimic trait values.

Dimorphic ESSs. In the following, we will use m_1 and m_2 to denote any two mimic types in a dimorphic population (by convention, we will let $m_1 < m_2$). We will assume that evolution takes place on two timescales: selection causes genotype frequencies to change fast in the population, while mutations arise only rarely. Thus, we may assume that the genotype frequencies in a population always are at equilibrium when new mutations are introduced. Consequently, in order for a pair of mimic strategies (m_1, m_2) to constitute a dimorphic ESS, we require that the mimic types should coexist at a stable equilibrium and be uninvadable by mutant strategies at that equilibrium.

Let ρ denote the proportion of m_1 mimics and $(1 - \rho)$ denote the proportion of m_2 mimics. We will use $t^*(m_1, m_2, \rho)$ to denote the optimal operator threshold against a dimorphic mimic population (as before, $t^*(m)$ will denote optimal operator thresholds in monomorphic mimic populations). The optimal threshold $t^*(m_1, m_2, \rho)$ changes monotonically from $t^*(m_2)$ to $t^*(m_1)$ as ρ changes from 0 to 1 (see appendix). In the following, we will refer to a population state (m_1, m_2, ρ) as stationary if m_1 and m_2 have equal payoffs when operators adopt the optimal threshold $t^*(m_1, m_2, \rho)$. A stationary population state (m_1, m_2, ρ^*) is stable if it also satisfies

$$\left[\frac{\partial}{\partial\rho} w_m(m_1, t^*(m_1, m_2, \rho))\right]_{\rho=\rho^*} < \left[\frac{\partial}{\partial\rho} w_m(m_2, t^*(m_1, m_2, \rho))\right]_{\rho=\rho^*}. \quad (4)$$

This condition says that the fitness of the two morphs is frequency dependent in the sense that each morph has a slightly higher fitness than the other when present in a marginally smaller proportion than at the equilibrium ρ^*.

In order to constitute a dimorphic ESS, a dimorphic strategy pair must coexist at a stable population state and be uninvadable by mutants at this population state. If the strategy pair has a unique stable population state, the dimorphic ESS can be characterized solely by its constituent strategies m_1 and m_2. In scenarios B1 and B2, a stationary population state is stable and unique if $t^*(m_1) < t^*(m_2)$ (see appendix). In scenario A3, a stationary population state is stable and unique if $t^*(m_1) < t^*(m_2)$ and m_1 and m_2 are mutually invasible (see appendix). The two strategies m_1 and m_2 are mutually invasible when both $w_m(m_1, t^*(m_2)) > w_m(m_2, t^*(m_2))$ and $w_m(m_2, t^*(m_1)) > w_m(m_1, t^*(m_1))$.

Existence of Dimorphic ESSs. In the appendix, it is shown that each S-shaped mimic nullcline features at least one operator threshold t' at which the corresponding pair of mimic strategies (m_1^*, m_2^*) lying on the upper and lower parts of the nullcline obtain equal payoffs. Since we are not able to solve $\partial w_m/\partial m = 0$ for m analytically, such strategy pairs must be found numerically; our explorations of scenarios A3, B1, and B2 always resulted in a unique solution. Given that the upper and lower part of the mimic nullcline represent fitness maxima, such pairs are, when stable, also uninvadable by mutants. Because $t^*(m_1, m_2, \rho)$ changes monotonically from $t^*(m_2)$ to $t^*(m_1)$ as ρ changes from 0 to 1, a necessary condition for the stability of (m_1^*, m_2^*) is that t' must lie between $t^*(m_1^*)$ and $t^*(m_2^*)$. By applying the additional stability conditions given above for scenario B1 and B2, it is easy to see that (m_1^*, m_2^*) is a dimorphic ESS if $t^*(m_1^*) < t' < t^*(m_2^*)$; in scenario A3 (m_1^*, m_2^*) is a dimorphic ESS if $t^*(m_1^*) < t' < t^*(m_2^*)$ and m_1^* and m_2^* are mutually invasible.

In our exploration of the model, we found that S-shaped mimic nullclines intersected with operator nullclines once (the most typical case) or three times (this happened most often when both $(m_{\text{opt}} - m_{\min})$ and m_{opt} were high). Due to the shape of the nullclines, it is clear that the requirement $t^*(m_1^*) < t' < t^*(m_2^*)$ must always hold when the operator nullcline intersects with the S-shaped mimic nullcline only at a single point, and this point lies on the section of the mimic nullcline that runs from (t', m_1^*) to (t', m_2^*). Moreover, in each case where the nullclines intersected three times and t' was between $t^*(m_1^*)$ and $t^*(m_2^*)$, we found that $t^*(m_1^*) < t' < t^*(m_2^*)$. In our explorations of scenario A3, the strategy pair (m_1^*, m_2^*) was almost always mutually invasible when $t^*(m_1^*) < t' < t^*(m_2^*)$. In the very few cases where mutual invasibility did not hold, stability was confirmed using condition (20) in the appendix.

The Attainability of Dimorphic ESSs. In many cases, operator nullclines may intersect with S-shaped mimic null-

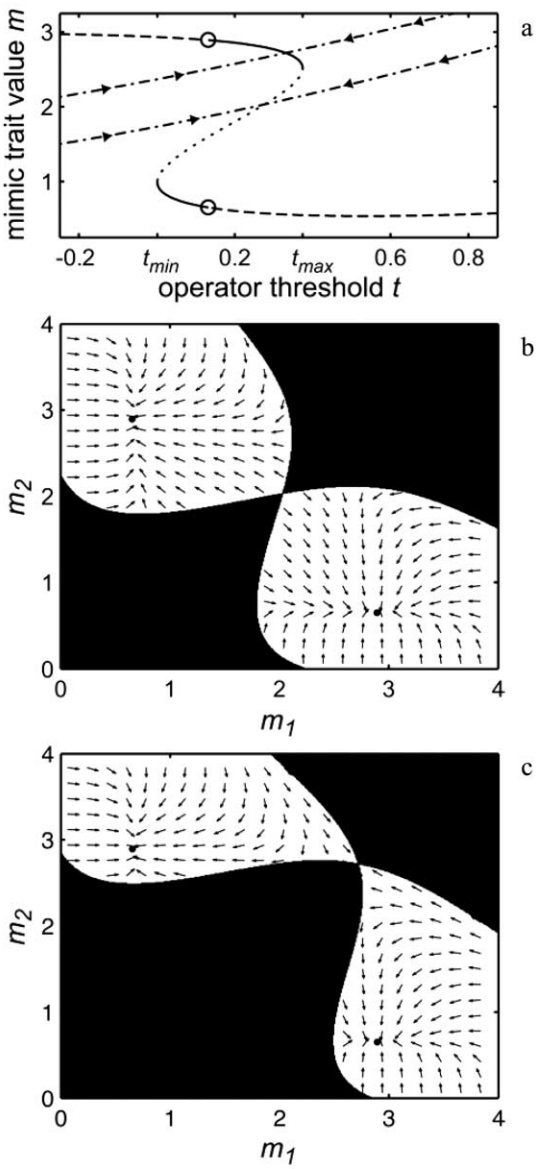

Figure 5: *a*, An S-shaped mimic nullcline is shown; dashed line sections give mimic trait values m^* that are global fitness maxima for different operator thresholds *t*, solid line sections give trait values that are local fitness maxima, and dotted line sections give trait values that are local fitness minima. Two operator nullclines are also shown (*dashed-dotted lines*); the upper intersects the mimic nullcline at a local fitness maximum, while the lower intersects at a local fitness minimum, representing an evolutionary branching point. Due to strong separation of timescales (evolutionary vs. behavioral timescale), the dynamics are simplified, and the trajectories follow the operator nullcline (see main text for further explanation). The two circles represent the two mimic types in the dimorphic evolutionarily stable strategy (ESS). *b*, *c*, The white areas represent pairs of trait values (m_1, m_2) that are mutually invasible. The superimposed vectors give the relative rate and direction of evolution in the two morphs. *b*, The dimorphic ESS is attainable from a monomorphic state using small mutational steps. This corresponds to the case in *a* where the operator nullcline intersects the mimic nullcline at a local fitness minimum. *c*, The dimorphic ESS is not attainable from a monomorphic state using small mutational steps but requires an initial mutation with a large phenotypic effect. This corresponds to the case in *a* where the operator nullcline intersects the mimic nullcline at a local fitness maximum. Parameter values: $nq = 2$, $m_{min} = -0.5$, $m_{opt} = 3$, $\sigma = 1$. *a*, Upper operator nullcline: $\ln(K) = 3.2$. Lower operator nullcline: $\ln(K) = 1.5$. *b*, $\ln(K) = 1.5$. *c*, $\ln(K) = 3.2$.

clines only at a point (*t*, *m*) that is a local fitness minimum for the mimic population, that is, where $\partial^2 w_m/\partial m^2 > 0$ (fig. 5*a*). Monomorphic mimic populations will evolve closer and closer to this point, which is an evolutionary branching point sensu Geritz et al. (1997, 1998). If the resident population is close to the branching point and a mutant type arises with a phenotypic value just on the other side of the branching point, the mutant will invade but cannot replace the resident population. The two mimic types will be mutually invasible, each having higher fitness than the other when sufficiently rare, and will coexist in a "protected dimorphism" sensu Geritz et al. (1998). When $t^*(m_1) < t^*(m_2)$, each protected dimorphism has a unique stable population state.

The evolutionary attainability of dimorphic ESSs has been explored (numerically) for different sets of parameters by using adaptive dynamics in which each morph in a protected dimorphism evolves at a rate proportional to its fitness gradient at the unique stable population state. The results when nullclines intersect only once may be summarized as follows. Dimorphic ESSs are attractors for sufficiently close strategy pairs and are thus attainable by natural selection. The size of the basin of attraction depends on the mimetic load *K*. If the operator nullcline intersects the mimic nullcline at a fitness minimum, a monomorphic population will first evolve toward the evolutionary branching point, at which the population will diverge into a protected dimorphism. The protected dimorphism will then evolve toward the dimorphic ESS. This is illustrated in figure 5*b*, which show mutually invasible mimic strategy pairs for the same parameter choices as in figure 5*a*, with vectors superimposed that give the relative rate and direction of evolution in the two morphs. If the operator nullcline instead intersects the mimic nullcline at a local fitness maximum, the monomorphic population will evolve toward it and remain there. In this case, the dimorphic ESS has a smaller basin of attraction, and only mutations with large phenotypic effects can take the population away from the local fitness maximum and give rise to a protected dimorphism that evolves toward the dimorphic ESS (fig. 5*c*).

By choosing the right parameter values and mimetic loads, one can find operator and mimic nullclines that intersect more than once (not shown). Partial exploration of such cases indicates that dimorphic ESSs will still be

attainable, in the sense that strategy pairs that are sufficiently close to a dimorphic ESS will evolve closer to it.

The Relationship between Cognitive Abilities and Mimetic Accuracy. Operators may in the long run evolve a fundamentally better cognitive machinery that makes them perceive model and mimic trait values as more distinct. In the model, we may represent the evolution of a better cognitive apparatus by reducing the standard deviation σ of the operators' trait value estimates. When σ decreases, the probability of correct model classification will increase for a given probability of correct mimic classification. Thus, a better cognitive apparatus will always be beneficial for an operator.

If operators evolve a fundamentally better cognitive machinery, the optimal level of mimicry m^* may change in the mimic population. In the appendix, we show (for scenario B1) that $\partial m^*/\partial\sigma > 0$ for positive m^*, satisfying

$$m^* < \sqrt{2\sigma^2 + 2\sqrt{\sigma^4(\ln(K))^2 + 1}}. \quad (5)$$

For higher m^*, by contrast, $\partial m^*/\partial\sigma < 0$. Thus, at any given mimetic load K, a lower σ will select for better mimicry only if m^* is sufficiently low.

Results

The main results of the model are summarized in figure 6, which shows evolutionarily stable levels of mimicry corresponding to different mimetic loads K. In the case of aggressive mimicry (scenarios A1–A3), the mimetic load is low when mimics are rare relative to models and when mimics impose low costs on operators. Under these circumstances, the operators adopt a high threshold and will often interact with crude mimics ("adaptive gullibility"; Wiley 1994). This favors inaccurate mimicry, which has low cost. When the mimetic load is high, by contrast, the operators adopt a low threshold and tend to be deceived only by very good mimicry ("adaptive fastidiousness"; Wiley 1994).

In the case of obligate aggressive mimicry (scenarios A1 and A2), the only way that an aggressive mimic can obtain a fitness gain is by deceiving an operator. As a result, higher mimetic loads and adaptive fastidiousness on the part of operators always select for better mimicry (fig. 6*a*). When $m_{\min} < 0$ and the mimetic load is sufficiently high, no ESS exists, and the mimics will race toward better and better mimicry. However, such steady improvements in mimicry will force the operator population to adopt lower and lower rejection thresholds (fig. 3*c*), and the mimicry system will at some point break down, because virtually no operators will respond to either models or mimics.

In the case of Batesian mimicry (scenarios B1 and B2),

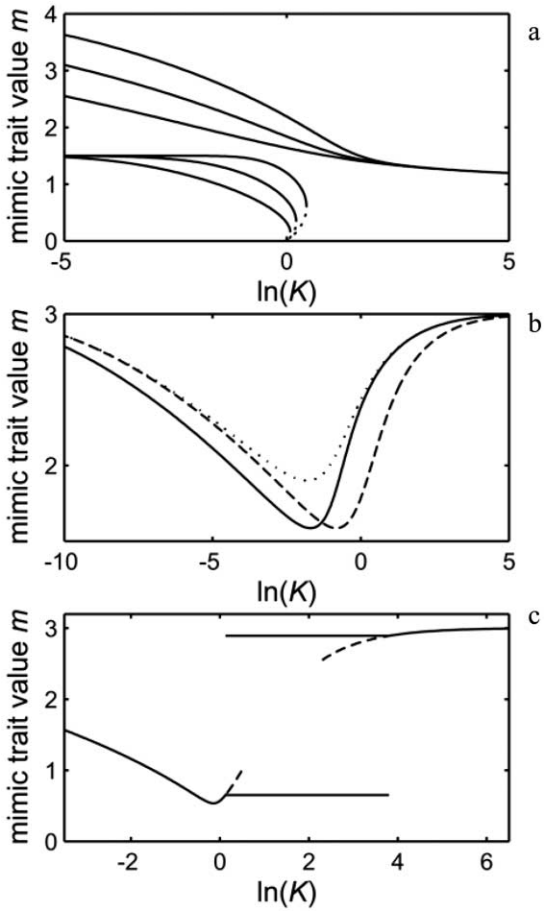

Figure 6: Evolutionarily stable levels of mimicry are plotted against $\ln(K)$, where K is the mimetic load. *a*, Scenarios A1 and A2 (aggressive mimicry): Solid line sections represent evolutionarily stable strategies (ESSs) that are attainable by natural selection, while dotted line sections represent unattainable ESSs. Increasing mimetic loads select for better mimicry. If $m_{\min}$ is negative (the three bottom lines) and K is sufficiently high, no ESS exists, and the mimics will race toward better and better mimicry, forcing the operators to set lower and lower thresholds. *b*, Scenario A3 (aggressive mimicry) and scenario B1 and B2 (Batesian mimicry), low incentive for deception: The lines represent ESSs that are attainable by natural selection. Intermediate mimetic loads select for more accurate mimicry, while high and low mimetic loads select for inaccurate mimicry. c, Scenario B1, high incentive for deception: Solid, curved line sections represent mimic trait values that are monomorphic ESSs, while the solid, straight line sections represent the two mimic types in the dimorphic ESS. All ESSs are attainable by natural selection. The dashed line sections represent mimic trait values that locally maximize fitness. Parameter values: $\sigma = 1$. *a*, Top to bottom: scenario A2, $m_{\min} = 1$, $m_{\text{opt}} = 4$, $n = 12$; scenario A2, $m_{\min} = 1$, $m_{\text{opt}} = 4$, $n = 4$; scenario A1, $m_{\min} = 1$, $m_{\text{opt}} = 4$; scenario A2, $m_{\min} = -0.5$, $m_{\text{opt}} = 1.5$, $n = 12$; scenario A2, $m_{\min} = -0.5$, $m_{\text{opt}} = 1.5$, $n = 4$; scenario A1, $m_{\min} = -0.5$, $m_{\text{opt}} = 1.5$. *b*, $m_{\min} = -0.5$, $m_{\text{opt}} = 3$. *Dashed line*, scenario A3, $(1 - r)/r = 1$. *Dotted line*, scenario B1, $nq = 0.5$; *solid line*, scenario B2, $q\lambda/(1 - \lambda) = 1$. *c*, Scenario B1, $nq = 2$, $m_{\min} = -0.5$, $m_{\text{opt}} = 3$, $\sigma = 1$ (as in fig. 5).

the mimetic load is low when mimic prey is rare relative to model prey and when attacks on mimics yield a benefit that is small relative to the cost of attacking defended models. Under these circumstances, the operators set a high attack threshold so that few prey will be attacked. Thus, costly mimicry is not necessary, and crude mimics are favored (as in scenarios A1, A2, and A3). As the mimetic load increases, there is at first selection for more accurate mimicry (fig. 6*b*), but at sufficiently high mimetic loads, crude mimics have greater fitness again. At these high mimetic loads, predators are very likely to attack, and it may simply become too costly to evolve the level of mimicry that is necessary for protection. Instead, the mimics profit from having an inaccurate but cheap level of mimicry that will enable them, if by chance they escape the attention of predators, to realize as much fitness benefit as possible.

The fact that the aggressive mimics in scenarios A1 and A2 must encounter and deceive operators in order to obtain fitness gains and the Batesian mimics in scenarios B1 and B2 may obtain fitness gains without encountering a single operator leads to opposite predictions for the level of mimicry at high mimetic loads. However, if we allow for the possibility that aggressive mimics can sometimes receive a fitness benefit without deceiving the focal operator species, as in scenario A3, high mimetic loads will favor crude mimics, just as in the case of Batesian mimicry (fig. 6*b*).

When mimics may potentially receive a fitness benefit without deceiving any operators, as in scenarios A3, B1, and B2, but the incentive for deception is high, the mimic population may also evolve mimetic dimorphism at some intermediate mimetic loads. One mimic morph will be inaccurate (trait values close to m_{opt}), while the other will be much more accurate. Figure 6*c* summarizes the most typical case (when the nullclines only intersect once for any given K) for scenario B1: When the mimetic load is low or high, only a monomorphic ESS exists in the mimic population (*curved solid lines*). When the mimetic load is intermediate, a unique dimorphic ESS exists (*straight solid lines*) but no monomorphic ESS. However, at some intermediate mimetic loads, the mimic population may also reach a locally stable monomorphic level of mimicry (*dashed lines*); this can only be invaded by mutants with large phenotypic effects. Scenarios A3 and B2 lead to the same qualitative results as shown for scenario B1 in figure 6*c*.

For a given S-shaped mimic nullcline, the mimic trait values at the dimorphic ESSs are exactly the same over the whole range of mimetic loads that give rise to dimorphism (fig. 6*c*). Thus, the trait values at a dimorphic ESS are stable against small perturbations of mimetic load; no mutant may replace one of the pure strategies in the support of the dimorphic ESS. However, the proportion of inaccurate mimics that will be present at a dimorphic ESS will change when mimetic load is perturbed; the proportion rises from 0 to 1 as mimetic load increases. The fitness values of the two morphs are constant over this range, which follows from the fact that the operators always adopt the same threshold at the dimorphic ESSs. The range of mimetic loads that give rise to dimorphic ESSs may be very wide for a given mimic nullcline; in figure 6*c*, K varies over an order of magnitude.

In order to illustrate what goes on at a dimorphic ESS, we have done some calculations for the dimorphic ESSs associated with each of the seven S-shaped nullclines shown in figure 4: The mimics' expected payoffs are in the range of 5%–31% of the expected payoff obtained by a hypothetical mimic that pays no costs of mimicry and always succeeds in deception. Moreover, an accurate mimic is in each encounter much more likely to deceive an operator than an inaccurate mimic (eight to 1,009 times more likely). However, in the absence of discrimination (if all mimics were classified as models), the accurate morph would, due to costs of mimicry, only obtain between 42%–75% of the expected benefit obtained by the inaccurate morph.

Evolutionary improvements in cognitive abilities among the operators will affect the degree of mimicry at equilibrium. This has been investigated in scenario B1. Suppose that a mimic population is monomorphic for a trait value m^* that is a local fitness maximum. If the operators become marginally better at discriminating between models and mimics (i.e., if σ decreases slightly) and m^* is sufficiently low, the mimics will come under selection for more accurate mimicry. In other words, better discrimination abilities select for better mimicry at equilibrium. On the other hand, if m^* is high, better discrimination abilities may in fact select for less accurate mimicry at equilibrium.

Discussion

We have used evolutionary game theory and signal detection theory to analyze the evolution of mimicry under the assumption that mimicry-enhancing traits confer costs on their carriers. Two factors have turned out to be crucial for predicting the degree of mimicry at equilibrium; both reflect ecological conditions and emerge naturally as aggregated parameters in the model. The first predictive factor, which we have called "mimetic load," reflects the selection pressure imposed by mimics on operators. The second predictive factor, called the "incentive for deception," reflects the extent to which the fitness of the mimics depends on their ability to deceive operators.

In Batesian mimicry, the mimetic load is low when mimic prey are unprofitable and/or rare and when model

prey are very nasty or dangerous. In aggressive mimicry, the mimetic load is low when mimics are rare and/or impose only low costs on their victims. The importance of mimetic load for understanding mimicry systems is generally acknowledged. To be sure, the consensus view is that Batesian mimics should be more successful at deception (and be attacked less) at low mimetic loads (e.g., Sheppard 1959; Nur 1970; Turner 1978; Huheey 1988; Lindström et al. 1997; Sherratt 2002). Likewise, aggressive mimics should be more successful at duping victims at low mimetic loads (e.g., Dawkins and Krebs 1979; Wiley 1994; Davies et al. 1996; Rodríguez-Gironés and Lotem 1999; Holen et al. 2001). As mimetic load increases, however, operators should become less gullible, which is in general thought to increase selection pressure for better mimicry, at least as long as the mimetic load does not reach a critical level (see below). In accordance with this, our model predicts inaccurate mimicry as the only evolutionary equilibrium at low mimetic loads, due to the assumed trade-offs between mimicry-enhancing traits and other aspects of fitness. As the mimetic load increases from low to intermediate, operators become more discriminating, resulting in increasingly accurate mimicry at evolutionary equilibrium. However, mimicry will never be perfect at evolutionary equilibrium: as in the model of Servedio and Lande (2003), the optimal mimic trait value will represent a compromise between the need to deceive operators and other aspects of fitness.

The second predictive factor—the incentive for deception—has received less attention in mimicry theory but may be crucial for understanding the selection pressures acting on mimics at high mimetic loads. In Batesian mimicry, the incentive for deception is high if the mimics often encounter operators and/or are easily captured by them. Thus, this factor is related to the amount of enemy-free space that the mimetic prey has. In aggressive mimicry, the incentive for deception is high when mimics are strongly dependent on deceiving victims in order to obtain resources necessary for survival and/or reproduction. For instance, a specialist avian brood parasite may be completely locked into a coevolutionary arms race with its host, thus having a high incentive for deception, while a spider-hunting spider may use mimicry as only one of several hunting techniques (e.g., Jackson and Pollard 1996), resulting in a lower incentive for deception. The model shows that the level of mimicry at equilibrium increases when the incentive for deception increases.

Batesian Mimicry at High Mimetic Loads

In the context of Batesian mimicry, it is thought that a high abundance of mimics relative to models may make it difficult for predators to associate the warning pattern with unpalatability, which could cause the protection from predation to disappear completely both for models and mimics (e.g., Sheppard 1959; Nur 1970; Holloway et al. 2002). It has therefore been proposed that a high abundance of mimics may lead to the evolution of mimetic polymorphism because this will spread the parasitic load over several defended prey species, making each warning signal less diluted; alternatively, a high abundance may cause relaxed selection on mimetic perfectionism, leading to inaccurate mimicry and/or high levels of phenotypic variation in the mimetic trait (Sheppard 1959; Nur 1970; Huheey 1988; Holloway et al. 2002). Our model enables us to predict the conditions that would favor one or the other of these two alternatives.

When the incentive for deception is low (i.e., mimics rarely encounter predators and/or are difficult to capture), high mimetic loads should select for monomorphic inaccurate mimicry. When the incentive for deception is high (i.e., mimics often encounter predators and/or are easily captured), there may be selection for mimetic dimorphism when mimetic loads are high and monomorphic inaccurate mimicry when mimetic loads are very high. A mimetic dimorphism will consist of one inaccurate and one accurate morph (a so-called mimic-nonmimic polymorphism; Turner 1978). The accurate morph maximizes its probability of deceiving predators, while the inaccurate morph maximizes its fitness if by chance it should escape predators. This result is in accordance with Turner (1978), who argued that stability of mimic-nonmimic polymorphisms would depend on nonmimics being fitter than mimics in the absence of models or predators, for instance, due to having better flight abilities or advantages in mate choice. Mimic-nonmimic polymorphisms are well-known among butterflies but typically in a sex-limited form with mimetic females and nonmimetic males (e.g., Wickler 1968; Turner 1978). Sex-limited cases with mimetic males and nonmimetic females have been reported in moths (Turner 1978) and beetles (Hespenheide 1975). Both female-limited and male-limited mimicry are known among bumblebee mimics (Gilbert, in press).

To summarize, our model predicts that (1) Batesian mimics that are very common and/or mimic very weakly defended models should evolve either inaccurate mimicry (by stabilizing selection) or mimetic polymorphism, and (2) Batesian mimics that are very common and/or mimic very weakly defended models are more likely to evolve mimetic polymorphism if they encounter predators at high rates and/or are bad at evading predator attacks.

Note that these predictions concern patterns of mimicry that are expected at evolutionary equilibrium but do not imply anything about the origin of mimetic interactions. Moreover, our model does not explicitly incorporate switching to mimicry of other models than the focal

model. Nevertheless, there are good reasons to believe that conditions that select for mimic-nonmimic polymorphisms could also lead to mimicry of different models. Because the inaccurate morph in a sense has "given up" mimicry of the focal model, it may benefit strongly from evolving similarity to other defended model species. It would be very interesting to extend our modelling approach to multimodel mimicry; we suspect that prediction (1) and (2) may also hold for other types of mimetic polymorphisms.

It has recently been pointed out that many common Batesian mimetic hoverflies are rather poor mimics, while rare mimics tend to have a closer resemblance to their models (Azmeh et al. 1998; Edmunds 2000; Howarth and Edmunds 2000). Proposed explanations are that common mimics are less precise because they are "generalist" mimics that resemble several model species but none very closely (Edmunds 2000; Sherratt 2002) and that common mimics may be less precise due to kin selection effects (Johnstone 2002). Our model provides an alternative and quite general explanation for why common Batesian mimics should be inaccurate, though the different explanations are not necessarily mutually exclusive.

A Batesian mimic that mimics an aposematic model may be more conspicuous and be more often detected by predators. Consequently, it will have a higher incentive for deception. Therefore, we may expect polymorphisms to be more widespread in Batesian systems where aposematic patterns are mimicked (e.g., many butterflies and hoverflies) than in other Batesian systems (ant-mimicking spiders). Moreover, if accurate mimicry interferes with escape abilities, for instance due to behavioral mimicry (e.g., Golding and Edmunds 2000) or changes in body shape (e.g., Srygley 1994), polymorphisms may be more likely. A very intriguing pattern is present in mimetic hoverflies: bumblebee mimics are often polymorphic, while bee and wasp mimics are not (Gilbert, in press). Could it be that bumblebee mimics are worse at evading predators when attacked? Bumblebee mimics are also larger than bee and wasp mimics (Gilbert, in press); perhaps bumblebee mimics, due to their large size, are easier to detect and thus need to deceive predators more often, which again may select for polymorphism.

Aggressive Mimicry at High Mimetic Loads

Obligate aggressive mimics cannot gain benefits in any way other than through the deception of an operator. The predictions for obligate aggressive mimicry are easily summarized: the mimetic accuracy at equilibrium always improves as mimetic load increases. However, note that evolutionary equilibria are sometimes lacking when mimetic load is sufficiently high; this gives rise to evolutionary chases toward better mimicry and stronger discrimination.

In the facultative case, aggressive mimics may sometimes obtain benefits without deceiving operators. This leads to essentially the same predictions at high mimetic loads as for Batesian mimicry: a low incentive for deception selects for inaccurate mimicry, and a high incentive for deception may select for mimetic dimorphism when mimetic loads are high and monomorphic inaccurate mimicry when mimetic loads are very high.

At high mimetic loads, aggressive mimics impose a very strong ecological pressure on operators. Many systems of aggressive mimicry could probably not be maintained at a constant mimetic load sufficiently high to cause operators to adopt the strict thresholds that make inaccurate mimicry and dimorphisms possible. Consequently, the model predictions associated with high mimetic loads may seem to be of limited value for aggressive mimicry. This line of argument, however, overlooks the possibility that the relative rate at which operators encounter mimics may vary temporally. Operators would then benefit from accepting most signalers (and reducing the risk of failing to respond to model organisms) when mimetic load is low and discriminating more strongly (and reducing the risk of parasitism or predation) when mimetic load is high. Such flexible discrimination may cause mimics to mostly encounter operators that adopt strict thresholds, thus making inaccurate mimicry and dimorphisms feasible even when the mimetic load averaged over all periods is low. Flexible discrimination may easily follow from a continuous learning process or could be due to a phenotypically plastic response induced by environmental cues. Evidence for cue-induced flexible discrimination is known from avian brood parasitism: some hosts of the common cuckoo *Cuculus canorus* show stronger egg discrimination when adult cuckoos have been spotted near the nest (e.g., Davies and Brooke 1988; Moksnes et al. 1993).

Cognitive Abilities and Mimetic Accuracy

Although our model does not explicitly predict evolutionary changes in cognitive abilities in the operator population, it does predict some effects that this would have on the equilibrium level of mimicry. From an evolutionary arms race perspective, one would expect improved discrimination abilities (e.g., better visual acuity) to select for better mimicry. Our model confirms this for Batesian mimics (scenario B1) that reside at an accurate evolutionary equilibrium, which certainly will be selected for even better mimicry when discrimination abilities improve. However, improved discrimination abilities select for less accurate mimicry in populations that already reside at a sufficiently inaccurate evolutionary equilibrium. For

strongly constrained mimics, it may simply be too costly to improve mimicry further when the discrimination abilities of the operators improve; instead, they benefit from having an even cheaper level of mimicry. Thus, the operators may "win" this round of the evolutionary arms race, forcing the mimics to shift to other strategies than mimicry. This kind of collapse in investment by one side as the other side's investment increases has been predicted in several models of coevolutionary arms races (e.g., Abrams 1986; Greeff and Parker 2000).

Conclusion

The explicit incorporation of costs of mimicry in our model leads to several new predictions that have not been anticipated beforehand. Most surprisingly, the model indicates that costs of mimicry may sometimes lead to the evolution of mimetic dimorphism. Thus, the effects of such costs are not trivial. Our analysis shows that the evolution of mimicry under high mimetic loads depends very strongly on how tightly the mimics are locked into the ecological interaction with the operators (i.e., on the "incentive for deception"). It suggests that enemy-free space and escape abilities of mimetic prey are crucial factors for understanding the evolution of inaccurate mimicry and polymorphisms among common Batesian mimics.

Acknowledgments

We thank C. Brinch, N. Davies, M. Enquist, G.-P. Saetre, T. Slagsvold, N. C. Stenseth, and J. O. Vik for helpful discussions and comments on the manuscript; F. Schiestl for valuable advice regarding deceptive orchids; and two anonymous reviewers for their constructive comments and suggestions. Financial support was provided by the Norwegian Research Council (to Ø.H.H.).

Literature Cited

Abrams, P. A. 1986. Adaptive responses of predators to prey and prey to predators: the failure of the arms-race analogy. Evolution 40:1229–1247.

Azmeh, S., J. Owen, K. Sørensen, D. Grewcock, and F. Gilbert. 1998. Mimicry profiles are affected by human-induced habitat changes. Proceedings of the Royal Society of London B 265:2285–2290.

Brakefield, P. M. 1985. Polymorphic Müllerian mimicry and interactions with thermal melanism in ladybirds and a soldier beetle: a hypothesis. Biological Journal of the Linnean Society 26:243–267.

Cushing, P. E. 1997. Myrmecomorphy and myrmecophily in spiders: a review. Florida Entomologist 80:165–193.

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., M. de L. Brooke, and A. Kacelnik. 1996. Recognition errors and probability of parasitism determine whether reed warblers should accept or reject mimetic cuckoo eggs. Proceedings of the Royal Society of London B 263:925–931.

Dawkins, R., and J. R. Krebs. 1979. Arms races between and within species. Proceedings of the Royal Society of London B 205:489–511.

Dittrich, W., F. Gilbert, P. Green, P. McGregor, and D. Grewcock. 1993. Imperfect mimicry: a pigeon's perspective. Proceedings of the Royal Society of London B 251:195–200.

Edmunds, M. 2000. Why are there good and poor mimics? Biological Journal of the Linnean Society 70:459–466.

Egan, J. P. 1975. Signal detection theory and ROC analysis. Academic Press series in cognition and perception. Academic Press, New York.

Gavrilets, S., and A. Hastings. 1998. Coevolutionary chase in two-species systems with applications to mimicry. Journal of Theoretical Biology 191:415–427.

Geritz, S. A. H., J. A. J. Metz, É. Kisdi, and G. Meszéna. 1997. Dynamics of adaptation and evolutionary branching. Physical Review Letters 78:2024–2027.

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.

Gilbert, F. In press. The evolution of imperfect mimicry. *In* M. D. E. Fellowes, G. J. Holloway, and J. Rolff, eds. Insect evolutionary ecology. CABI, Wallingford.

Golding, Y. C., and M. Edmunds. 2000. Behavioural mimicry of honeybees (*Apis mellifera)* by droneflies (Diptera: Syrphidae: *Eristalis* spp.). Proceedings of the Royal Society of London B 267:903–909.

Golding, Y. C., A. R. Ennos, and M. Edmunds. 2001. Similarity in flight behaviour between the honeybee *Apis mellifera* (Hymenoptera: Apidae) and its presumed mimic, the dronefly *Eristalis tenax* (Diptera: Syrphidae). Journal of Experimental Biology 204:139–145.

Greeff, J. M., and G. A. Parker. 2000. Spermicide by females: what should males do? Proceedings of the Royal Society of London B 267:1759–1763.

Greenwood, J. J. D. 1986. Crypsis, mimicry, and switching by optimal foragers. American Naturalist 128:294–300.

Haynes, K. F., and K. V. Yeargan. 1999. Exploitation of intraspecific communication systems: illicit signalers and receivers. Annals of the Entomological Society of America 92:960–970.

Hespenheide, H. A. 1975. Reversed sex-limited mimicry in a beetle. Evolution 29:780–783.

Hofbauer, J., and K. Sigmund. 1998. Evolutionary games and population dynamics. Cambridge University Press, Cambridge.

Holen, Ø. H., G.-P. Saetre, T. Slagsvold, and N. C. Stenseth. 2001. Parasites and supernormal manipulation. Proceedings of the Royal Society of London B 268:2551–2558.

Holloway, G., F. Gilbert, and A. Brandt. 2002. The relationship between mimetic imperfection and phenotypic variation in insect colour patterns. Proceedings of the Royal Society of London B 269:411–416.

Holloway, G. J., C. G. Marriott, and H. J. Crocker. 1997. Phenotypic plasticity in hoverflies: the relationship between colour pattern and season in *Episyrphus balteatus* and other Syrphidae. Ecological Entomology 22:425–432.

Holmgren, N. M. A., and M. Enquist. 1998. Dynamics of mimicry evolution. Biological Journal of the Linnean Society 66:145–158.

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.

Huheey, J. E. 1988. Mathematical models of mimicry. American Naturalist 131(suppl.):S22–S41.

Jackson, R. R., and S. D. Pollard. 1996. Predatory behavior of jumping spiders. Annual Review of Entomology 41: 287–308.

Johnson, S. D., and T. J. Edwards. 2000. The structure and function of orchid pollinaria. Plant Systematics and Evolution 222:243–269.

Johnstone, R. A. 2002. The evolution of inaccurate mimics. Nature 418:524–526.

Kullenberg, B. 1961. Studies in *Ophrys* pollination. Zoologiska Bidrag från Uppsala 34:1–340.

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.

Maynard Smith, J. 1982. Evolution and the theory of games. Cambridge University Press, Cambridge.

McIver, J. D., and G. Stonedahl. 1993. Myrmecomorphy: morphological and behavioral mimicry of ants. Annual Review of Entomology 38:351–379.

Moksnes, A., E. Røskaft, 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.

Nur, U. 1970. Evolutionary rates of models and mimics in Batesian mimicry. American Naturalist 104:477–486.

O'Connell, L. M., and M. O. Johnston. 1998. Male and female pollination success in a deceptive orchid: a selection study. Ecology 79:1246–1260.

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.

Oxford, G. S., and R. G. Gillespie. 1998. Evolution and ecology of spider coloration. Annual Review of Entomology 43:619–643.

Peakall, R. 1990. Responses of male *Zaspilothynnus trilobatus* Turner wasps to females and the sexually deceptive orchid it pollinates. Functional Ecology 4:159–167.

Pekár, S., and J. Král. 2002. Mimicry complex in two central European zodariid spiders (Araneae: Zodariidae): how *Zodarion* deceives ants. Biological Journal of the Linnean Society 75:517–532.

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. 1990. A model system for coevolution: avian brood parasitism. Annual Review of Ecology and Systematics 21:481–508.

Rothstein, S. I., and S. K. Robinson. 1998*a*. 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.

———. 1998*b*. Parasitic birds and their hosts: studies in coevolution. Oxford Ornithology Series. Oxford University Press, New York.

Schiestl, F. P., and M. Ayasse. 2001. Post-pollination emission of a repellent compound in a sexually deceptive orchid: a new mechanism for maximising reproductive success? Oecologia (Berlin) 126:531–534.

Servedio, M. R., and R. Lande. 2003. Coevolution of an avian host and its parasitic cuckoo. Evolution 57:1164–1175.

Sheppard, P. M. 1959. The evolution of mimicry: a problem in ecology and genetics. Cold Spring Harbor Symposia on Quantitative Biology 24:131–140.

Sherratt, T. N. 2002. The evolution of imperfect mimicry. Behavioral Ecology 13:821–826.

Shreeve, T. G. 1992. Adult behaviour. Pages 22–45 *in* R. L. H. Dennis, ed. The ecology of butterflies in Britain. Oxford University Press, Oxford.

Smith, S. M. 1975. Innate recognition of coral snake pattern by a possible avian predator. Science 187:759–760.

Srygley, R. B. 1994. Locomotor mimicry in butterflies? the associations of positions of centres of mass among groups of mimetic, unprofitable prey. Philosophical

Transactions of the Royal Society of London B 343:145–155.

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.

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.

Turner, J. R. G. 1978. Why male butterflies are non-mimetic: natural selection, sexual selection, group selection, modification and sieving. Biological Journal of the Linnean Society 10:385–432.

Vane-Wright, R. I. 1976. A unified classification of mimetic resemblances. Biological Journal of the Linnean Society 8:25–56.

Wickler, W. 1968. Mimicry in plants and animals. Weidenfeld & Nicolson, London.

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.

Wilson, H. R. 1991. Interrelationships of egg size, chick size, posthatching growth and hatchability. World's Poultry Science Journal 47:5–20.

Associate Editor: Allen J. Moore