



CHICAGO JOURNALS



The University of Chicago

Costs of Learning and the Evolution of Mimetic Signals

Author(s): David W. Kikuchi, Thomas N. Sherratt

Source: *The American Naturalist*, Vol. 186, No. 3 (September 2015), pp. 321-332

Published by: [The University of Chicago Press](#) for [The American Society of Naturalists](#)

Stable URL: <http://www.jstor.org/stable/10.1086/682371>

Accessed: 09/09/2015 10:39

Your use of the JSTOR archive indicates your acceptance of the Terms & Conditions of Use, available at
<http://www.jstor.org/page/info/about/policies/terms.jsp>

JSTOR is a not-for-profit service that helps scholars, researchers, and students discover, use, and build upon a wide range of content in a trusted digital archive. We use information technology and tools to increase productivity and facilitate new forms of scholarship. For more information about JSTOR, please contact support@jstor.org.



The University of Chicago Press, The American Society of Naturalists, The University of Chicago are collaborating with JSTOR to digitize, preserve and extend access to *The American Naturalist*.

<http://www.jstor.org>

Costs of Learning and the Evolution of Mimetic Signals

David W. Kikuchi* and Thomas N. Sherratt

Department of Biology, Carleton University, Ottawa, Ontario K1S 5B6, Canada

Submitted August 29, 2014; Accepted May 5, 2015; Electronically published July 10, 2015

Online enhancement: appendix.

ABSTRACT: Predators must use the appearance of their prey to decide whether it is likely to be defended. Most theory assumes that predators have completed learning about prey appearance, yet we do not understand how predators learn which aspects of appearance to use for classifying prey. If sampling prey can be risky, predators might forgo opportunities to learn about the relationship between prey appearance and defense. Using Bayesian inference and dynamic programming, we modeled how the immediate risks and future rewards of learning about prey appearance influence how predators learn. In addition, we explored how variation in predator learning affects the evolution of mimicry, which occurs when two prey evolve to share a common signal to predators. We found that when learning about prey with distinct appearances was expensive, optimal predators tended to lump them into the same category or exhibit an unwillingness to sample at all (neophobia). This resulted in a reduction in selection for defensive mimicry. However, the same predator behavior favored the evolution of aggressive mimicry, because in that case, mimics benefited from being sampled. When prey were very rare and costs of sampling them were high, predators exhibited neophobia, refusing to attack. This behavior could forestall the evolution of mimicry and instead select for polymorphism.

Keywords: animal communication, multicomponent signal, categorization, exploration-exploitation, bandit problem, Bayesian inference.

Introduction

Those who make many species are the “splitters,” and those who make few are the “lumpers.” (Darwin, 1857)

Deciding whether organisms of different appearances should be classified separately is a critical dilemma, not only for taxonomists, as Darwin's letter pithily points out, but also for predators (Bradbury and Vehrencamp 2011). A predator faced with variants of unfamiliar prey must decide whether to treat them all as the same or to learn about each separately. How predators decide this is of considerable importance, not just for revealing insights about cate-

gorization behavior, but also for shaping selection on prey. Consider the evolution of Batesian mimicry, where undefended prey (mimics) resemble defended prey (models) to deceive predators (Bates 1862). If, for example, a mimetic insect has a differently colored abdomen from its model, selection for improvement in mimicry would depend on whether predators lump model and mimic phenotypes together or split them apart. Indeed, there are many puzzling instances of imperfect mimicry in nature where a Batesian mimic shares only a few features in common with its model (Kikuchi and Pfennig 2013), and several experiments have shown that predators ignore potentially informative features (Beatty et al. 2004; Bain et al. 2007; Kikuchi and Pfennig 2010; Valkonen et al. 2011; Hossie and Sherratt 2013). Predators might well consider prey that share more features to be more likely to have similar defenses; nevertheless, if predators paid attention to those features not shared, they would be able to successfully separate rewarding prey from unrewarding prey when the prey types differ in profitability. Here, we seek a formal treatment of when predators should invest in learning about specific features of their prey. We assume the prey are already similar enough for predators to consider treating them as the same—for example, they may be varieties of snakes, beetles, or butterflies. Conversely, we would not expect a predator to do anything but treat prey with radically different appearances, such as beetles and snakes, independently.

Most theoretical studies of how predators categorize prey assume that learning is already complete (e.g., Oaten et al. 1975; Getty 1985; Sherratt 2002). Although these are highly relevant for understanding many behaviors, we are interested in what goes on during the learning process, which has long been argued to be important for the evolution of mimicry (e.g., Müller 1879). Indeed, learning behavior has been studied extensively in the context of foraging (e.g., Getty and Krebs 1985; Alatalo and Mappes 1996; Lindström et al. 1997, 2004). Here, we examine the optimal way that predators should categorize different prey as they learn about them when they are free to choose which and how many prey they sample and the implications that such categorization behavior has for selection on prey appearance.

* Corresponding author; e-mail: dwkikuchi@gmail.com.

Am. Nat. 2015. Vol. 186, pp. 321–332. © 2015 by The University of Chicago.
0003-0147/2015/18603-5573\$15.00. All rights reserved.
DOI: 10.1086/682371

We model a predator that learns about the profitability of two prey that can be perfectly discriminated on the basis of an obvious appearance cue (e.g., insects with black or red abdomens) but must decide whether to treat these distinct varieties as having different properties. We then analyze how the optimal predator's behavior will affect the evolution of prey appearance, specifically focusing on mimicry. Our model is very general, however, and can be applied to many phenomena where agents may or may not decide to invest in the necessary learning to finely categorize objects' properties, such as sexual selection, habitat selection, or even internet advertising (Pandey et al. 2007).

Predators are not concerned with classification for the sake of erudition but for survival, so they will seek the level of categorization that maximizes their fitness. Learning can be an expensive form of phenotypic plasticity (Snell-Rood 2013), so predators must economize how many prey they sample to gain information. They need to balance any future rewards of learning about unfamiliar prey against the likely immediate cost if a prey turns out to be defended. This is an exploration-exploitation problem, where an agent must sometimes choose between learning more and exploiting the knowledge it already has (Jones 1978; Krebs et al. 1978; Cohen et al. 2007; Dayan and Daw 2008). The general conundrum has also been presented as a multiarmed bandit problem, in reference to hypothetical gamblers who must choose between several slot machine arms, each with an unknown probability of returning a reward (Robbins 1952; Bellman 1956)—at each moment in time, gamblers are faced with the option of finding out more about the slot machines or exploiting their current knowledge to select what they believe to be the highest-paying arm.

Sherratt (2011) argued that predators encountering an unfamiliar prey type are faced with a two-armed bandit problem, with one arm's probability of reward fixed at zero (rejection of prey item) and the other arm (attack) having an unknown benefit (depending on the probability that prey are defended and how costly defended prey are). He showed that predators should sample more of a prey as it increases in abundance, because the future value of the information grows higher. Conversely, if the cost of sampling a defended prey item is high and/or these prey are rare, then predators should exhibit neophobia (reluctance to sample unknown prey), because the possible benefit and future value of information does not outweigh the likely immediate sampling cost. Here we extend the model of Sherratt (2011) to allow for prey with two different appearance phenotypes. Such an approach allows novel questions to be asked, such as under what conditions predators will incorporate a new prey appearance dimension into their categorization strategy and what implications this will have for the evolution of mimicry.

Model Description

We assume that two types of prey (1 and 2) are encountered with probabilities q and $1 - q$, respectively, in a random order across a finite T time steps, which may be interpreted very generally since the relevant scale for some predator-prey relationships may range from a few encounters across a lifespan to a single foraging bout in an unfamiliar habitat patch. In nature, predators would also learn about q , but this is a different problem from ours, because that information can be gleaned passively rather than through potentially costly, deliberate interactions with prey. The difference in the appearance of prey might be represented by any categorical trait, such as color. The optimal predator can use information gained by sampling prey to evaluate whether it should treat each distinct prey type as a separate bandit problem or lump them together as one.

For simplicity, we assume that payoffs comprise two levels (profitable and unprofitable), so that the predator gains a fixed benefit b for attacking a desired prey (e.g., undefended prey) and pays a cost c for attacking an undesired one (e.g., defended prey). Predators must infer any association between payoff and appearance (we are more interested in predators that must learn the probability of a prey being costly to attack than the magnitude of costs, because classes of organisms often have similar costs—think of most venomous snakes as deadly and most noxious butterflies as unpalatable).

Prey of appearance i are costly to attack with true probability p_i and beneficial to attack with true probability $1 - p_i$. Predators are ignorant of the true value of p_i and must estimate it as π_i . To perform their estimation, they keep track of the number of costly trials (r_i) out of their total number of trials (n_i) in which they have sampled prey of appearance i . Thus, (r_i, n_i) constitutes the data that the predator has about the prey it has attacked. As new information concerning the profitability of prey is accumulated, Bayesian inference converts prior beliefs (beliefs held before sampling) into posterior beliefs (beliefs held after conducting a learning trial; Courville et al. 2006; McNamara et al. 2006; Kruschke 2010). Moreover, if the prior belief distribution is chosen appropriately (i.e., a conjugate to the parameter being estimated), then the posterior conveniently follows the same distribution, albeit with different parameter values (Bolker 2008). The conjugate prior for a probability p with binomial outcomes (costly/beneficial) is the beta distribution. If our prior is a beta distribution with shape parameters α and β (i.e., $\text{Beta}(\alpha, \beta)$), then after r_i costly experiences from n_i trials, the posterior distribution will follow $\text{Beta}(\alpha + r_i, \beta + n_i - r_i)$. These priors α and β represent the beliefs of the predators and could be gained through inheritance, observational learning, or other mechanisms. In any beta distribution, the expectation is simply $\alpha/(\alpha + \beta)$,

so that at any stage, the best estimate of the proportion of prey with appearance i being costly is $\pi_i(r_i, n_i) = (\alpha_i + r_i) / (\alpha_i + \beta_i + n_i)$, where α_i and β_i are the original parameters of the beta distribution for predators with no experience at all with the prey. Note that as n_i increases, the expectation converges on r_i/n_i .

For a predator to accurately estimate the palatability of one prey appearance, it does not need to sample as extensively as it would to accurately estimate the palatability of two. When a predator lumps prey together, their probability of defense can be described by a single parameter $\pi_0(r_1, r_2, n_1, n_2) = (\alpha_0 + r_1 + r_2) / (\alpha_0 + \beta_0 + n_1 + n_2)$. On the other hand, if it splits prey by appearance i , then a separate estimate $\pi_i(r_i, n_i)$ must be made for each. The latter model may be advantageous if $p_1 \neq p_2$, but it may also be expensive to gather enough data for it to be accurate. Therefore, the probability that each model is true will affect the predator's decision to attack or defer. On encountering prey, the predator will calculate these probabilities given its prior expectations and the experience it has gained during previous foraging. We introduce Bayesian model selection as a means of calculating the probability of either a one-parameter or a two-parameter model, that is, $p_1 = p_2$ versus $p_1 \neq p_2$ (and so, our optimal predator chooses between two alternative models). We refer to the one- and two-parameter models with subscripts L and S, respectively, for conceptually lumping and splitting.

Bayesian model selection requires calculating each model's marginal likelihood, which is inferred from the data gathered while sampling, and each model's prior probability of being true. Initially, we assume equal prior probabilities for the one- and two-parameter models (i.e., $\Pr(L) = 1 - \Pr(S) = .5$), which amounts to an assumption about the exact learning rule the predator uses. Whenever $\Pr(S)$ is not set at 0 or 1, it is possible to consider both models, but if it is 0 or 1, then there is only one possible model (split or lump, respectively). The marginal likelihoods $\bar{\mathcal{L}}_j$ of each model are found by integrating each model's likelihood over the priors for its parameter(s), that is, the Beta(α, β) priors for π_i . Below are the binomial likelihoods of each model integrated over its prior(s):

$$\bar{\mathcal{L}}_L(r_1, r_2, n_1, n_2) = \prod_{i=1}^k \binom{n_i}{r_i} \frac{\Gamma(\alpha_0 + \beta_0)}{\Gamma(\alpha_0)\Gamma(\beta_0)} \frac{\Gamma(\sum_{i=1}^k r_i + \alpha_0) \Gamma(N - \sum_{i=1}^k r_i + \beta_0)}{\Gamma(N + \alpha_0 + \beta_0)}, \quad (1)$$

$$\bar{\mathcal{L}}_S(r_1, r_2, n_1, n_2) = \prod_{i=1}^k \binom{n_i}{r_i} \frac{\Gamma(\alpha_i + \beta_i)}{\Gamma(\alpha_i)\Gamma(\beta_i)} \frac{\Gamma(r_i + \alpha_i) \Gamma(n_i - r_i + \beta_i)}{\Gamma(n_i + \alpha_i + \beta_i)}, \quad (2)$$

where k is the total number of prey phenotypes and $N = \sum_{i=1}^k n_i$. Under the assumption of uniform Beta(1, 1) priors for all π_i , these expressions can be simplified to

$$\bar{\mathcal{L}}_L(r_1, r_2, n_1, n_2) = \prod_{i=1}^k \binom{n_i}{r_i} \frac{\Gamma(\sum_{i=1}^k (r_i + 1)) \Gamma(N - \sum_{i=1}^k (r_i + 1))}{\Gamma(N + 2)} \quad (3)$$

and

$$\bar{\mathcal{L}}_S(n_1, n_2) = \left[\prod_{i=1}^k (n_i + 1) \right]^{-1} \quad (4)$$

(cf. Ntzoufras 2009, pp. 399–400). In equation (4), r_i conveniently drops out. The marginal likelihoods in equations (3) and (4) represent the probabilities of the data given each model, $\Pr(\text{data}|\text{model} = \text{true})$. By applying Bayes's rule to the priors and marginal likelihoods, the posterior odds ratio is

$$\frac{\Pr(L = \text{true}|\text{data})}{\Pr(S = \text{true}|\text{data})} = \frac{\Pr(\text{data}|L = \text{true})}{\Pr(\text{data}|S = \text{true})} \cdot \frac{\Pr(L)}{\Pr(S)}. \quad (5)$$

If we assume $\Pr(S) = .5$, an even simpler expression can be used to calculate the posterior probability \mathcal{P}_j of each model j by dividing the marginal likelihood $\bar{\mathcal{L}}_j$ of each model by the sum of the marginal likelihoods of all m models:

$$\mathcal{P}_j(r_1, r_2, n_1, n_2) = \frac{\bar{\mathcal{L}}_j(\dots)}{\sum_{i=1}^m \bar{\mathcal{L}}_i(\dots)}. \quad (6)$$

At each time step, marginal likelihood of the one- and two-parameter models is recalculated to update \mathcal{P}_j (fig. A1; figs. A1–A3 available online).

Naturally, the predator's decisions will be guided by not only whether a one- or two-parameter model is most probable, but also the immediate costs and benefits of attacking or not and potential future benefits that it might reap from the information gained through attacking. To model the predator's entire decision-making process, we use stochastic dynamic programming (Mangel and Clark 1988; Clark and Mangel 2000). Dynamic programming explores all possible sequences of actions and outcomes to arrive at an optimal solution for behavior.

The predator's dynamic programming equation (DPE) is partitioned first between encountering prey of one appearance or another and then between attacking and deferring. The payoffs for attacking and deferring are each estimated according to one- and two-parameter models. The payoff for deferring is always the same under both models, because if a predator does not sample the prey it encounters, then it will simply advance to the next time step. We express this as $U(r_1, r_2, n_1, n_2, t + 1)$. Estimates of the payoff to attacking are, however, the averaged predictions of the one- and two-parameter models, each weighted by their posterior probabilities of being true. For example, on encountering prey type 1, the one-parameter model predicts the following payoff to attacking: $\pi_0(r_1, r_2, n_1, n_2)(U(r_1 + 1,$

$r_2, n_1 + 1, n_2, t + 1) - c) + (1 - \pi_0(r_1, r_2, n_1, n_2))(U(r_1, r_2, n_1 + 1, n_2, t + 1) + b)$. The term $\pi_0(r_1, r_2, n_1, n_2)$ is the estimate of the probability that the prey is defended. $U(r_1 + 1, r_2, n_1 + 1, n_2, t + 1)$ is the future expected payoff if the prey turns out to be defended, and c is the immediate cost the predator will pay if the prey is defended. Correspondingly, the estimated probability that the prey is beneficial is $(1 - \pi_0(r_1, r_2, n_1, n_2))$, and the future and immediate benefits are $(U(r_1, r_2, n_1 + 1, n_2, t + 1) + b)$. This entire expression is multiplied by the probability that the one-parameter model is true, $\mathcal{G}_L(r_1, r_2, n_1, n_2)$, and summed with the analogous prediction made by the two-parameter model to produce a final estimate of the payoff for attacking. The optimal predator does not distinguish between one-parameter and two-parameter models any more than is worthwhile—unless the future benefit to learning about prey of two appearances is greater than the cost of sampling, it will simply lump them together. We present a graphical summary in figure 1, and the full DPE is given below:

$$\begin{aligned}
 U(r_1, r_2, n_1, n_2, t) = & q \cdot \max \left\{ \mathcal{G}_L(r_1, r_2, n_1, n_2) \right. \\
 & \cdot [\pi_0(r_1, r_2, n_1, n_2)(U(r_1 + 1, r_2, n_1 + 1, n_2, t + 1) - c) \\
 & + (1 - \pi_0(r_1, r_2, n_1, n_2))(U(r_1, r_2, n_1 + 1, n_2, t + 1) + b)] \\
 & + \mathcal{G}_S(r_1, r_2, n_1, n_2) \\
 & \cdot [\pi_1(r_1, n_1)(U(r_1 + 1, r_2, n_1 + 1, n_2, t + 1) - c) \\
 & + (1 - \pi_1(r_1, n_1))(U(r_1, r_2, n_1 + 1, n_2, t + 1) + b)], \\
 & U(r_1, r_2, n_1, n_2, t + 1) \left. \right\} + (1 - q) \cdot \max \left\{ \mathcal{G}_L(r_1, r_2, n_1, n_2) \right. \\
 & \cdot [\pi_0(r_1, r_2, n_1, n_2)(U(r_1, r_2 + 1, n_1, n_2 + 1, t + 1) - c) \\
 & + (1 - \pi_0(r_1, r_2, n_1, n_2))(U(r_1, r_2, n_1, n_2 + 1, t + 1) + b)] \\
 & + \mathcal{G}_S(r_1, r_2, n_1, n_2) \\
 & \cdot [\pi_2(r_2, n_2)(U(r_1, r_2 + 1, n_1, n_2 + 1, t + 1) - c) \\
 & + (1 - \pi_2(r_2, n_2))(U(r_1, r_2, n_1, n_2 + 1, t + 1) + b)], \\
 & U(r_1, r_2, n_1, n_2, t + 1) \left. \right\}.
 \end{aligned}
 \tag{7}$$

To solve the DPE, one must work backward from the maximum value for time, T , in a process called backward iteration (Mangel and Clark 1988; Clark and Mangel 2000).

Model Application

In the exploration of our predator with prior $\Pr(S) = .5$, we refer to it as the mixer because it equally considers the one- and two-parameter models. We compare it to two extremes that we call the pure lumper and the pure splitter, which

consider only one or the other, that is, $\Pr(S) = 0$ or 1 , respectively. For our comparisons, we focus on three ecological scenarios. The first includes one costly and one beneficial prey, each with a different appearance. This can select for Batesian mimicry (Bates 1862; Pfennig et al. 2001; Ruxton et al. 2004) or aggressive mimicry (Dominey 1981; Haynes et al. 2002; Cheney 2010; Stoddard 2012). The second scenario is where both prey are costly. This could select for Müllerian mimicry, in which two defended species evolve to resemble one another so that predators sample fewer of each (Müller 1879; Sherratt 2008). We also include a third scenario where both prey are beneficial to further characterize predator behavior under ecological conditions that they are likely to encounter concurrently with mimicry systems. To represent our three scenarios, we assume that $p_i = (1, 0)$, $(1, 1)$, or $(0, 0)$. Below, we first characterize the behavior mixers, pure lumpers, and pure splitters when confronted with each scenario and then subsequently explore how their behavior affects selection for mimicry.

Model Analysis

Behavior and Performance of Predator Strategies

For all analyses, we fix the values of the priors for p_i (not to be confused with the prior for the one- or two-parameter models) at $\text{Beta}(1, 1)$ so predators have uniform priors (i.e., expectation 0.5, all values between 0 and 1 equally likely). This assumption greatly simplifies the calculation of the marginal likelihoods (see eqq. [3], [4]) and allows us to remain agnostic about the true state of nature. We fix our time horizon at $T = 40$ because it allows for a reasonable population size of prey with each appearance but is still computationally tractable. All of our results are obtained by forward iteration (i.e., Monte Carlo simulation; Clark and Mangel 2000). If we assume $q = 0.5$, then in all of our forward iterations, predators encounter an average of 20 prey of each appearance. Costs of interacting with undesired prey can influence the evolution of mimicry (Duncan and Sheppard 1963; Oaten et al. 1975; Sherratt 2002; Holen and Johnstone 2004, 2006; Franks et al. 2009; Lehmann et al. 2014), so we let the cost c vary between 0 and 5, which covers the entire spectrum of possible cost-mediated variation in behavior (because with b fixed at 1 and $T = 40$, the mixer, pure lumper, and pure splitter refuse to sample for $c > 4.2$; fig. 2).

Differences in the behavior of the mixer and its extreme variants are driven by costs of sampling and perspectives on available information. The pure lumper behaves as though 40 prey of one appearance are available, whereas the pure splitter behaves as though there are on average 20 prey of two different appearances. Therefore, the pure lumpers are more optimistic about the future benefits of learning and

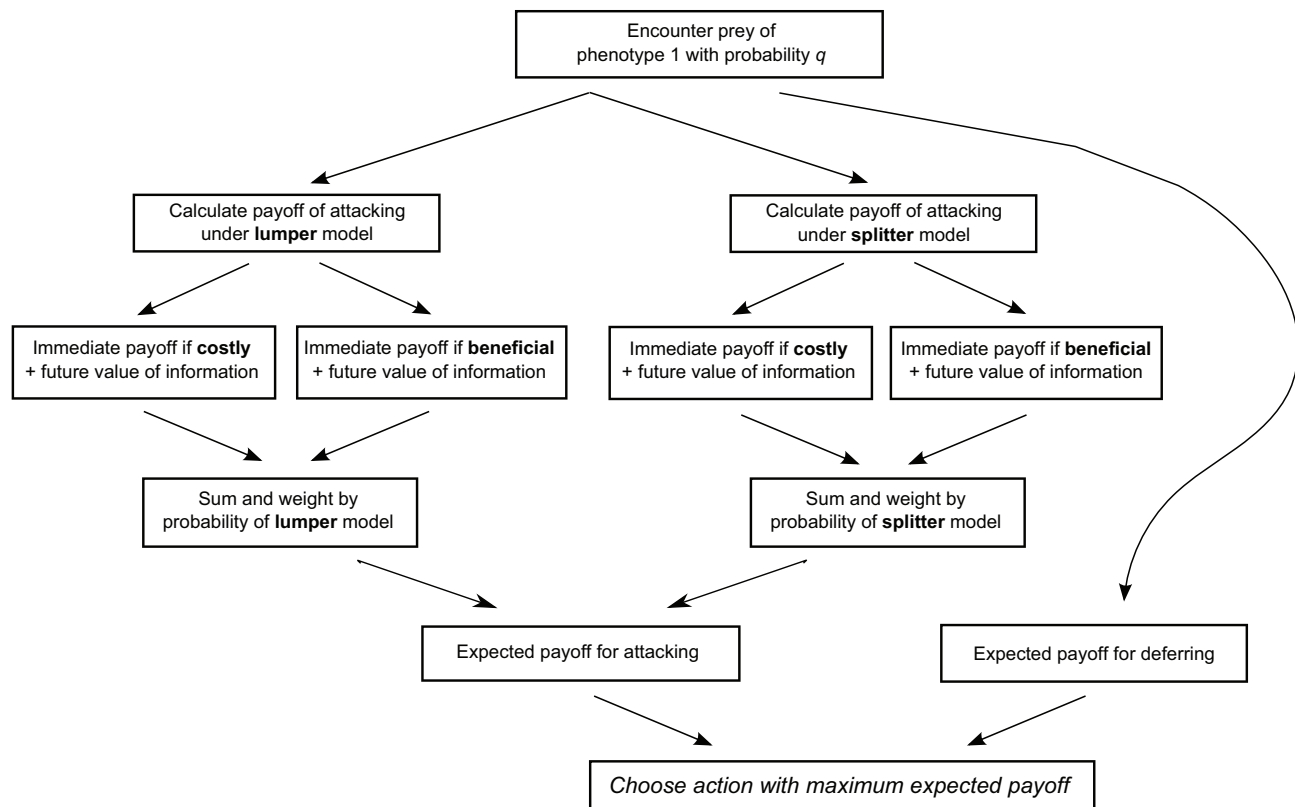


Figure 1: Schematic of the dynamic programming equation (7) for the predator's decision-making process. This diagram represents the process when a prey of appearance 1 is encountered, which is symbolically represented in the first half of equation (7). The process is the same for prey of appearance 2 but is not visualized here, for simplicity. The payoff from deferring (0) is the same whether the predator believes the two prey appearances indicate the same or different properties.

is willing to sample at higher values of c than the pure splitter (fig. 2). When a one-parameter model best represents the ecological scenario, for example, $p_i = (1, 1)$ or $(0, 0)$, the pure lumper outperforms the splitter by sampling fewer prey before ceasing if both prey are costly (fig. 2A) and by sampling at high values of c if both prey are beneficial (fig. 2B); all values of p_i are considered equally likely before sampling (uniform priors), so even when both prey are actually beneficial, high potential costs may deter sampling. However, when there is information contained about profitability in prey appearance, for example, $p_i = (1, 0)$, the pure splitter is able to take advantage of that information, consistently attacking all beneficial prey and quickly rejecting the costly ones (fig. 2C, 2D). The pure lumper does not distinguish between the two and thus fails to make any profit for $c > 1$ (fig. A2). In all scenarios, the mixer typically samples an intermediate number of prey relative to the pure lumper and pure splitter, which agrees with its constant ability to consider the one- and two-parameter models.

The most important result from the above analysis is that neither the pure lumper, the pure splitter, nor the mixer is

always best. Which prior has the highest fitness depends on how often each ecological scenario occurs. To find the best prior under different conditions, we estimated their performance when they have to learn about each of the three ecological scenarios with different frequencies (fig. 3). From this ternary plot, it is clear that over the whole region of cost space, the pure lumper generally does best when the mixture of scenarios encountered seldom contains one costly and one beneficial prey, and the pure splitter does best when the opposite is true. The mixer is superior when it encounters a mixture of both. In general, we suspect that if it were allowed to evolve, a predator's prior for the one-versus two-parameter models would come to reflect the frequency with which it has to learn about each ecological scenario during the course of its life (see table A1, available online, for an analysis of conditions to which the mixer is optimally adapted).

The pure lumper, pure splitter, and mixer can vary in their relative performance depending on costs, so we repeated our analysis on smaller cost intervals (fig. 3B–3E). This provides a nuanced perspective on how the costs of

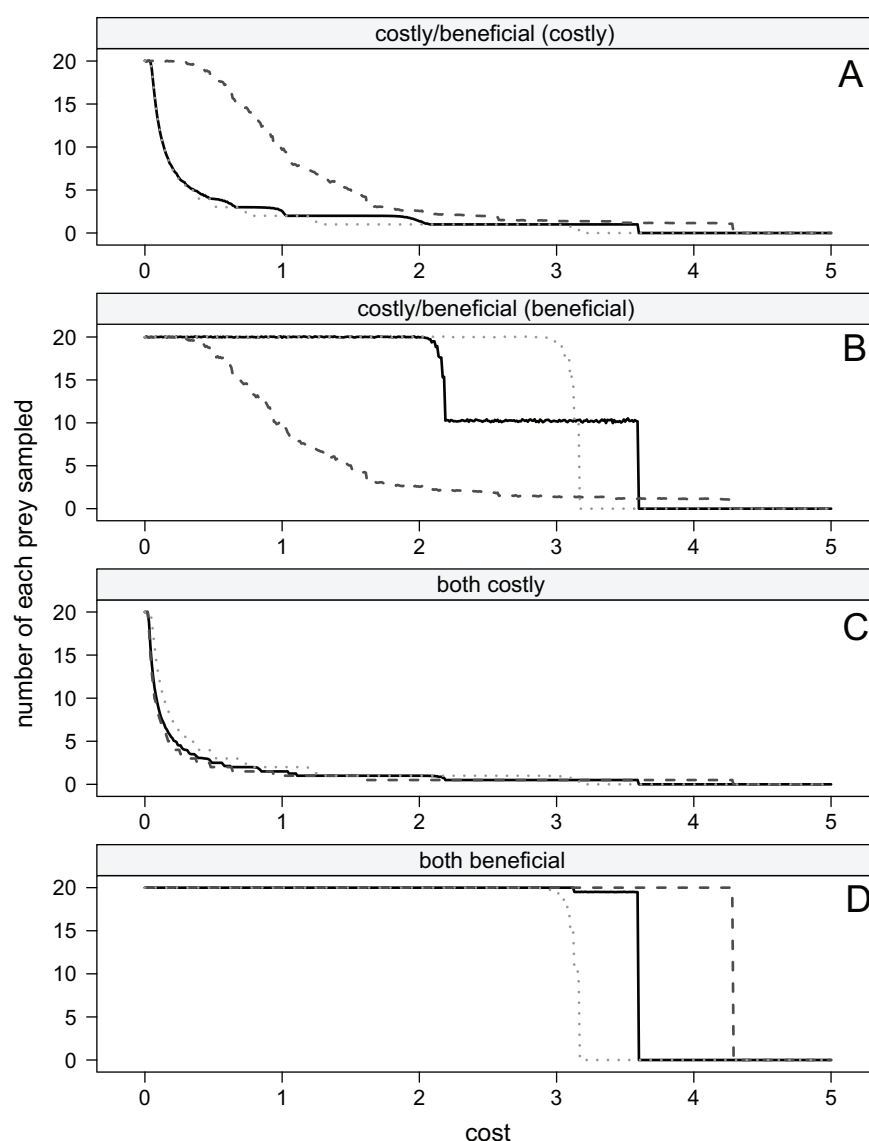


Figure 2: A–D, Number of prey sampled by the mixer, pure lumpers, and pure splitters under three different ecological scenarios, with cost varying between 0 and 5. The behavior of the mixer is represented by the solid line, the pure lumpers by the dashed line, and the pure splitters by the dotted line. Results are from 1,000 iterations. When prey were both costly and both beneficial, results for both appearances were averaged together.

sampling and the value of future information influence how advantageous it is to discriminate between prey. When c is low, the pure splitter's dominant performance shows that it pays to learn more about prey appearance (because sampling is cheap, and the more information, the better), but such behavior can be a liability when costs are high, so the mixer and pure lumpers predominate in such conditions. To explore sensitivity of our results to the number of time steps, we also performed an analysis with $T = 4$ and c ranging from 0 to 2. We found that, unsurprisingly, the pure

lumpers does much better when there are fewer time steps, because there is less value to future information (fig. A3).

Selection on Prey

In analyzing selection for traits to be incorporated into a signal, we must consider not only how predators behave but also how their actions impact prey. This produces a large number of potential relationships (Vanewright 1976; Endler 1981), so we analyze three of the most familiar. One

is Batesian mimicry, where mimics and models suffer a fitness cost from being sampled by a predator. We also explore classical (defensive) Müllerian mimicry, where both prey are defended and suffer a cost for being sampled. Finally, because it is also possible for sampling by predators to be beneficial, we included aggressive mimicry, in the sense that one prey may in fact be a predator or sexual rival and so benefits from being attacked by a predator.

In measuring fitness in Batesian mimicry and Müllerian mimicry, we assume that prey that are sampled by predators are killed, but if they go unsampled, they survive to reproduce. Under this assumption, the probability of surviving to reproduction is e^{-k} , where k is the average number of encounters with predators a prey must survive (Holen and Svenning 2012). For aggressive mimicry, we assume that each prey in the population needs to attract a predator only once to satisfy its needs. The probability of not being sampled is then $1 - e^{-k}$. For simplicity, we assume $k = 1$, which does not affect results qualitatively. Again, our results are obtained by forward iteration, and priors for p_i are fixed at Beta(1, 1). We assume that a novel mutant in the population makes negligible impact on q .

In Batesian mimicry, figure 4A shows that, unsurprisingly, the fitness of mutant perfect mimics and imperfect mimics is equal under predation by pure lumpers, since pure lumpers disregard variation. The pure splitter exerts selection as long as it continues to sample prey; however, it ceases sampling at high c (neophobia is optimal), and thus, the relative fitness of imperfect mimics becomes equal to that of perfect mimics. The mixer attacks all mimics while it favors the two-parameter model, but for a region of cost space between 2.2 and 3.17, it favors the one- or two-parameter models, depending on its first experience. Thus, it exerts weaker selection that nonetheless continues sampling at higher costs than the pure splitter does (fig. 4A). Of course, defended prey must be somewhat costly for there to be much benefit to mimicry in the first place, which explains why mutant fitness under both the pure splitter and the mixer rises rapidly as costs increase from zero.

Of fundamental importance is that at high costs, the mixer exerts little selection for improvement in the mimetic signal, because it increasingly follows the one-parameter model. Therefore, we expect that selection for mimicry will decrease with the cost of attacking defended prey. This is contrary to what we would expect if predators already knew the relationship between appearance and profitability, as in many signal detection models of mimicry.

In the case of aggressive mimicry, the relationship between cost and selection is reversed. When prey benefit from predator sampling, selection for the improvement of mimicry should increase with the cost of predator errors (fig. 4B), except in the case where predators are pure lumpers. This pattern emerges because beneficial prey—in this

case, the putative models—are always sampled completely if they are sampled at all, but a diminishing proportion of costly prey are sampled as costs increase.

Considering the evolution of Müllerian mimicry, we are interested not only in how selection varies with differences in costs but also in differences in the relative frequencies of prey of different appearances. Theory predicts that, all else being equal, the more abundant a prey appearance, the more advantageous it will be to resemble. We analyze relative fitness of mutants that resemble the nonfocal prey appearance relative to a focal prey that serves as a fitness reference. We assume all prey are defended and suffer from sampling. We analyze variation in both c and q (probability of a prey being type 1) to see under what conditions one prey should converge on the other (fig. 5). At low costs, when prey of each appearance are equally abundant, resembling the nonfocal prey does not increase fitness much, but at values that differ from parity, selection increasingly favors convergence by the rarer prey. This pattern disappears at higher costs. Neophobia is optimal at higher costs, so rare prey benefit from a complete lack of sampling (note values for selection less than 1, implying selection to remain distinctive). The result is that at moderate levels of cost, selection for convergence peaks at intermediate values of q , but at higher levels of cost, weak divergent selection often predominates, perhaps selecting for polymorphism as a result of neophobia.

Discussion

We have characterized how a predator should sample unfamiliar prey of two different appearances. The cost that predators pay for negative learning experiences is important, although the abundance of prey can also have an effect. Variation in the relative performance of the pure lumpers, pure splitter, and mixer priors suggests that predators' categorization behavior may evolve in response to the suite of ecological scenarios they typically encounter.

We have also explored when predator learning behavior will select for mimicry. Some outcomes are counterintuitive. For example, when predators encounter one beneficial prey type and one costly one, an analysis that neglects the costs of learning would predict the ready evolution of Batesian mimicry. However, under our model, at high potential costs, predators may not sample at all because neophobia is optimal, and their avoidance of all prey in a higher-level category (e.g., snakes or a coral snake and its imperfect mimic) would produce no selection for mimicry (or improvement in mimicry).

Our model may help explain why the behavior of predators toward mimics with complex signals—those that involve multiple traits, components, modalities, or displays—has hitherto been difficult to understand. Although many

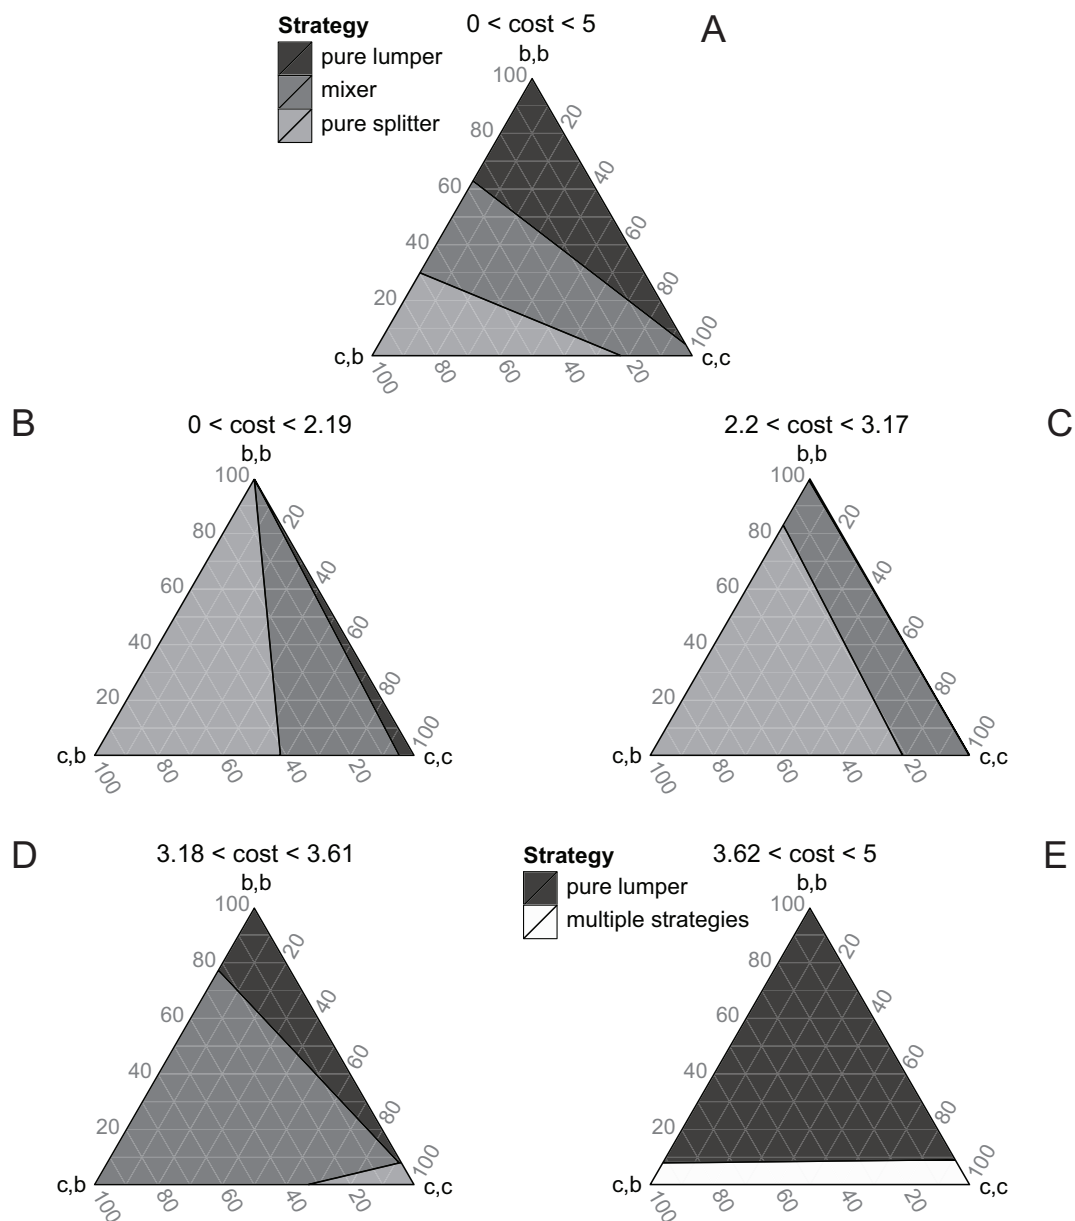


Figure 3: The prior (mixer, pure lumper, and pure splitter) with the highest fitness under conditions that vary in the proportion of time that predators encounter scenarios where both prey are beneficial, both are costly, or one is beneficial and one is costly. Each corner of these ternary plots represents a pure ecological scenario, but areas away from the corners represent mixtures of those scenarios. Areas where one prior outperforms the others are shaded. *A*, Over the whole range of costs considered, the mixer, pure lumper, and pure splitter each dominate for about one-third of possible conditions. The pure lumper does best when prey tend to provide the same payoff, and the pure splitter does best when prey tend to be different, but the mixer excels in the middle. *B–D*, When cost space is subdivided, the three priors show much variation in their degree of superiority over one another. At low costs, the pure splitter does relatively well because it is penalized less for sampling. However, as costs of sampling rise, the mixer dominates under a wide range of conditions because it can choose to treat prey of both appearances as the same. *E*, Above $c = 3.61$, only the pure lumper continues to sample and thus dramatically outpaces the pure splitter and mixer when it is possible to make a profit at all. However, the pure splitter and the mixer both outperform it when sampling is unprofitable (they have equal payoffs in the white region on the bottom of the triangle). The figure was produced by calculating predator performance (i.e., sum of costs + benefits earned over a learning session with 40 prey) in each of the three scenarios over each interval of c (we evaluated c at increments of 0.01 and averaged results within each interval) and then weighing those three performance measures by the frequency of encountering each scenario at each point in the ternary plot.

prey have evolved to mimic complex aposematic (warning) signals, predators often ignore some apparent components (e.g., Kikuchi and Pfennig 2010; Kazemi et al. 2014) or different combinations of components (e.g., Valkonen et al. 2011; Hossie and Sherratt 2013). Our theory suggests that predators may not always be willing to pay the sampling costs of learning exhaustively about differences between models and mimics. Each aspect of appearance that predators learn about will require further sampling to associate with a payoff, and the relative gain of learning about new traits will only diminish with the number of informative traits the predator has already learned about. Indeed, we think it will be prohibitively costly for predators to learn about all aspects that separate models and mimics, especially the numerous interactions between traits. Therefore, some degree of imperfection in mimicry is probable in all systems, although some mimics may be less imperfect than others.

Differing degrees of signal complexity have been documented in mimicry systems from various higher-level categories of prey. These differences may be linked to the stronger defenses of, for example, coral snakes compared to butterflies. The deadly New World coral snakes (*Micruroides*, *Leptomicrurus* spp., and *Micrurus* spp.) have simple patterns of colored rings on their bodies that advertise their deadly venom. They are mimicked by a host of nonvenomous snakes, but even imprecise mimics receive protection (Brodie 1993; Hinman et al. 1997; Harper and Pfennig 2007). In some regions, the aposematic signal that predators respond to may merely be the ratio of red:black (Kikuchi and Pfennig 2010). At the other extreme of signal complexity, unpalatable butterflies in the genus *Heliconius* have very complex color patterns that they employ in Müllerian mimicry. Many dimensions of their color patterns are under stringent selection by wild predators (Kapan 2001; Merrill et al. 2012; Finkbeiner et al. 2014). In light of the model that we have proposed here, predators may be much less willing to gather information about a potentially lethal snake than they are to gather information about a potentially distasteful butterfly. In fact, some predators have evolved innate aversions to prey that resemble coral snakes (Smith 1975, 1977). In general, our theory predicts that with increasing costs of interacting with undesired prey, the number of signal dimensions used in mimicry should decrease.

An unusual prediction of our model is that the fitness of prey with high potential costs of sampling should be highest when they are rare; the combination of prior expectation of defense and high costs results in neophobia. This prediction is borne out when novel undefended prey appear to benefit from neophobia (fig. 2D; Pfennig et al. 2007; Franks and Oxford 2009) and occurs because sampling rare prey provides low future benefits (see also Sherratt 2011). For




Figure 4: A, Relative fitness of undefended prey that mutate to resemble defended prey, assuming that prey suffer a fitness cost from being sampled by predators (e.g., Batesian mimicry). B, Fitness of aggressive mimics that mutate to resemble their models, assuming mimics benefit from being sampled by predators or other signal receivers (e.g., aggressive mimicry). Selection for resembling the alternative appearance is quite strong in both cases. Relative fitness for the mutant prey when the predator is a mixer is represented by the solid line; when it is a pure lumper, by the dashed line; and when it is a pure splitter, by the dotted line. Results are from 1,000 iterations; $q = 0.5$.

defended prey, however, this is contrary to the conventional wisdom that warning signals provide safety in numbers and the observation of uniformity in aposematic signals across some species' ranges (e.g., Harper and Pfennig 2007). The benefits of being rare according to our model may be mitigated if natural predators have uncertainty about the costs of sampling prey or their abundance. This could increase sampling, which would select for aposematism and Müllerian mimicry in defended species and perhaps crypsis rather than gaudy polymorphism in undefended species. In any case, more information about the incidence of polymorphism versus Müllerian mimicry in natural systems where there is geographic variation in abundance would be most welcome.

Models of mimicry based on signal detection theory (e.g., Duncan and Sheppard 1963; Oaten et al. 1975; Sherratt

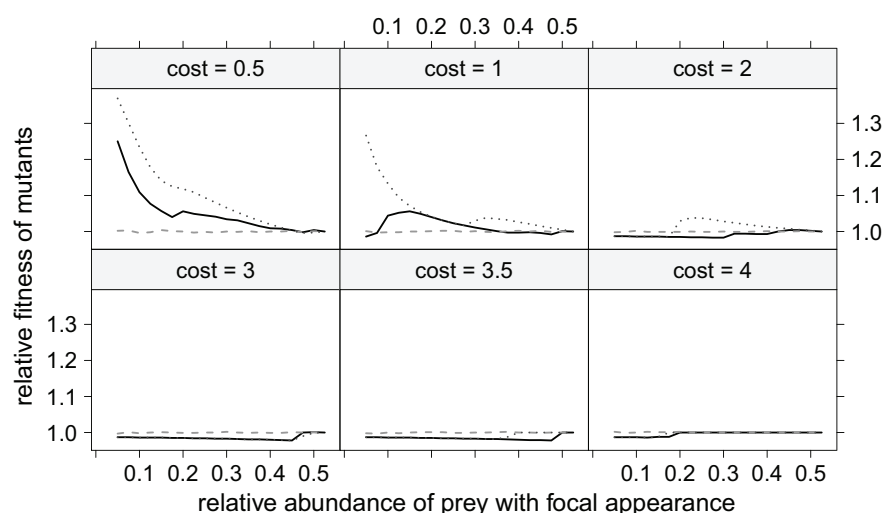


Figure 5: Müllerian mimicry. Relative fitness of a mutant defended prey that mutates to resemble another defended prey, assuming both are negatively affected by predator responses. At lower costs (defenses), there is more of an advantage to rarer prey in resembling the more abundant appearance phenotype. Higher costs erase much of this advantage, however, because rare prey can also benefit from neophobia on the part of predators. This indicates that selection for convergence is typically present at lower levels of cost and low abundance or at intermediate levels of cost and intermediate abundance. Relative fitness for the mutant prey when the predator is a mixer is represented by the solid line; when it is a pure lumper, by the dashed line; and when it is a pure splitter, by the dotted line. Results are from 1,000 iterations.

2002) have predicted that the precision of mimetic signals should generally decrease with increasing costs of errors, which is analogous to our prediction that an appearance trait should be less likely to be used with increasing costs. However, there is a distinction between these two patterns. Signal detection models describe the precision of mimicry within a dimension of appearance (or set of dimensions). Our analysis, on the other hand, assumes that predators can distinguish perfectly between prey of different appearances but must pay a cost of sampling to associate appearance with payoff. Thus, in contrast to signal detection theory, this approach addresses whether an aspect of appearance should be used at all.

In sum, we have found the optimal sampling model for a predator learning about prey with two different appearance phenotypes. Our analysis suggests that the selection imposed on prey by predators depends on the costs of learning. When prey suffer from predator sampling, higher costs should be less likely to select for mimicry, but when prey benefit from being sampled by predators, higher costs should increase selection for mimicry. Future studies should explicitly address the influence of costs of sampling on the number of appearance traits that predators use in categorizing prey.

Acknowledgments

We would like to thank K. Abbott, M. Servedio, and two anonymous reviewers for comments that greatly improved the manuscript. This work was supported by NSF-1306037

(D.W.K.) and an NSERC Discovery Grant (T.N.S.). We declare that we have no conflicts of interest.

Literature Cited

- Alatalo, R. V., and J. Mappes. 1996. Tracking the evolution of warning signals. *Nature* 382:708–710.
- Bain, R. S., A. Rashed, V. J. Cowper, F. S. Gilbert, and T. N. Sherratt. 2007. The key mimetic features of hoverflies through avian eyes. *Proceedings of the Royal Society B: Biological Sciences* 274:1949–1954.
- Bates, H. W. 1862. Contributions to an insect fauna of the Amazon valley (Lepidoptera: Heliconidae). *Transactions of the Linnean Society of London* 23:495–556.
- Beatty, C. D., K. Beirincx, and T. N. Sherratt. 2004. The evolution of Müllerian mimicry in multispecies communities. *Nature* 431:63–67.
- Bellman, R. 1956. A problem in the sequential design of experiments. *Sankhya* 16:221–229.
- Bolker, B. M. 2008. *Ecological models and data in R*. Princeton University Press, Princeton, NJ.
- Bradbury, J. W., and S. L. Vehrencamp. 2011. *Principles of animal communication*. Sinauer, Sunderland, MA.
- Brodie, E. D. 1993. Differential avoidance of coral snake banded patterns by free-ranging avian predators in Costa Rica. *Evolution* 47:227–235.
- Cheney, K. L. 2010. Multiple selective pressures apply to a coral reef fish mimic: a case of Batesian-aggressive mimicry. *Proceedings of the Royal Society B: Biological Sciences* 277:1849–1855.
- Clark, C. W., and M. Mangel. 2000. *Dynamic state variable models in ecology*. Oxford University Press, New York.
- Cohen, J. D., S. M. McClure, and A. J. Yu. 2007. Should I stay or should I go? how the human brain manages the trade-off between

- exploitation and exploration. *Philosophical Transactions of the Royal Society B: Biological Sciences* 362:933–942.
- Courville, A. C., N. D. Daw, and D. S. Touretzky. 2006. Bayesian theories of conditioning in a changing world. *Trends in Cognitive Sciences* 10:294–300.
- Darwin, C. R. 1857. Letter to J. D. Hooker on August 1. Darwin Correspondence Database. <http://www.darwinproject.ac.uk/entry-2130>.
- Dayan, P., and N. D. Daw. 2008. Decision theory, reinforcement learning, and the brain. *Cognitive, Affective, and Behavioral Neuroscience* 8:429–453.
- Dominey, W. J. 1981. Maintenance of female mimicry as a reproductive strategy in bluegill sunfish (*Lepomis macrochirus*). *Environmental Biology of Fishes* 6:59–64.
- Duncan, C. J., and P. M. Sheppard. 1963. Continuous and quantal theories of sensory discrimination. *Proceedings of the Royal Society B: Biological Sciences* 158:343–363.
- Endler, J. A. 1981. An overview of the relationships between mimicry and crypsis. *Biological Journal of the Linnean Society* 16:25–31.
- Finkbeiner, S. D., A. D. Briscoe, and R. D. Reed. 2014. Warning signals are seductive: relative contributions of color and pattern to predator avoidance and mate attraction in *Heliconius* butterflies. *Evolution* 68:3410–3420.
- Franks, D. W., and G. S. Oxford. 2009. The evolution of exuberant visible polymorphisms. *Evolution* 63:2697–2706.
- Franks, D. W., G. D. Ruxton, and T. N. Sherratt. 2009. Warning signals evolve to disengage Batesian mimics. *Evolution* 63:256–267.
- Getty, T. 1985. Discriminability and the sigmoid functional-response: how optimal foragers could stabilize model-mimic complexes. *American Naturalist* 125:239–256.
- Getty, T., and J. R. Krebs. 1985. Lagging partial preferences for cryptic prey: a signal detection analysis of great tit foraging. *American Naturalist* 125:39–60.
- Harper, G. R., and D. W. Pfennig. 2007. Mimicry on the edge: why do mimics vary in resemblance to their model in different parts of their geographical range? *Proceedings of the Royal Society B: Biological Sciences* 274:1955–1961.
- Haynes, K. F., C. Gemenio, K. V. Yeargan, 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? *Chemoecology* 12:99–105.
- Hinman, K. E., H. L. Throop, K. L. Adams, A. J. Dake, K. K. McLauchlan, and M. J. McKone. 1997. Predation by free-ranging birds on partial coral snake mimics: the importance of ring width and color. *Evolution* 51:1011–1014.
- Holen, Ø. H., and R. A. Johnstone. 2004. The evolution of mimicry under constraints. *American Naturalist* 164:598–613.
- . 2006. Context-dependent discrimination and the evolution of mimicry. *American Naturalist* 167:377–389.
- Holen, Ø. H., and T. O. Sæviing. 2012. Aposematism and the handicap principle. *American Naturalist* 180:629–641.
- Hossie, T. J., and T. N. Sherratt. 2013. Defensive posture and eyespots deter avian predators from attacking caterpillar models. *Animal Behaviour* 86:383–389.
- Jones, P. W. 1978. On the two armed bandit with one probability known. *Metrika* 25:235–239.
- Kapan, D. D. 2001. Three-butterfly system provides a field test of Müllerian mimicry. *Nature* 409:338–340.
- Kazemi, B., G. Gamberale-Stille, B. S. Tullberg, and O. Leimar. 2014. Stimulus salience as an explanation for imperfect mimicry. *Current Biology* 24:965–969.
- Kikuchi, D. W., and D. W. Pfennig. 2010. Predator cognition permits imperfect coral snake mimicry. *American Naturalist* 176:830–834.
- . 2013. Imperfect mimicry and the limits of natural selection. *Quarterly Review of Biology* 88:279–315.
- Krebs, J. R., A. Kacelnik, and P. Taylor. 1978. Test of optimal sampling by foraging great tits. *Nature* 275:27–31.
- Kruschke, J. K. 2010. *Doing Bayesian data analysis: a tutorial with R and WinBUGS*. Elsevier, Burlington, MA.
- Lehmann, K. D. S., B. W. Goldman, I. Dworkin, D. M. Bryson, and A. P. Wagner. 2014. From cues to signals: evolution of interspecific communication via aposematism and mimicry in a predator-prey system. *PLoS ONE* 9:e91783.
- Lindström, L., R. V. Alatalo, A. Lyytinen, and J. Mappes. 2004. The effect of alternative prey on the dynamics of imperfect Batesian and Müllerian mimics. *Evolution* 58:1294–1302.
- 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 B: Biological Sciences* 264:149–153.
- Mangel, M., and C. W. Clark. 1988. *Dynamic modeling in behavioral ecology*. Princeton University Press, Princeton, NJ.
- McNamara, J. M., R. F. Green, and O. Olsson. 2006. Bayes' theorem and its applications in animal behaviour. *Oikos* 112:243–251.
- Merrill, R. M., R. W. R. Wallbank, V. Bull, P. C. A. Salazar, J. Mallet, M. Stevens, and C. D. Jiggins. 2012. Disruptive ecological selection on a mating cue. *Proceedings of the Royal Society B: Biological Sciences* 279:4907–4913.
- Müller, F. 1879. *Ituna and Thyridia: a remarkable case of mimicry in butterflies*. Transactions of the Entomological Society of London 1879:xx–xxix.
- Ntzoufras, I. 2009. *Bayesian modeling using WinBUGS*. Wiley, Hoboken, NJ.
- 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.
- Pandey, S., D. Chakrabarti, and D. Agarwal. 2007. Multi-armed bandit problems with dependent arms. 24th International Conference on Machine Learning (ICML), pp. 721–728.
- Pfennig, D. W., W. R. Harcombe, and K. S. Pfennig. 2001. Frequency-dependent Batesian mimicry. *Nature* 410:323.
- Pfennig, D. W., G. R. Harper, A. F. Brumo, W. R. Harcombe, and K. S. Pfennig. 2007. Population differences in predation on Batesian mimics in allopatry with their model: selection against mimics is strongest when they are common. *Behavioral Ecology and Sociobiology* 61:505–511.
- Robbins, H. 1952. Some aspects of the sequential design of experiments. *Bulletin of the American Mathematical Society* 58:527–536.
- Ruxton, G. D., T. N. Sherratt, and M. P. Speed. 2004. *Avoiding attack*. Oxford University Press, New York.
- Sherratt, T. N. 2002. The evolution of imperfect mimicry. *Behavioral Ecology* 13:821–826.
- . 2008. The evolution of Müllerian mimicry. *Naturwissenschaften* 95:681–695.
- . 2011. The optimal sampling strategy for unfamiliar prey. *Evolution* 65:2014–2025.
- Smith, S. M. 1975. Innate recognition of coral snake pattern by a possible avian predator. *Science* 187:759–760.

- . 1977. Coral-snake pattern recognition and stimulus generalization by naive great kiskadees (Aves: Tyrannidae). *Nature* 265:535–536.
- Snell-Rood, E. C. 2013. An overview of the evolutionary causes and consequences of behavioural plasticity. *Animal Behaviour* 85: 1004–1011.
- Stoddard, M. C. 2012. Mimicry and masquerade from the avian visual perspective. *Current Zoology* 58:630–648.
- Valkonen, J. K., O. Nokelainen, and J. Mappes. 2011. Antipredatory function of head shape for vipers and their mimics. *PLoS ONE* 6:e22272.
- Vanewright, R. I. 1976. A unified classification of mimetic resemblances. *Biological Journal of the Linnean Society* 8:25–56.

Associate Editor: Vlastimil Křivan
Editor: Judith L. Bronstein



The harmless Arizona mountain kingsnake (*Lampropeltis pyromelana*) is a putative Batesian mimic of the deadly Arizona coral snake (*Micruroides euryxanthus*). The two have similar colors but a different order to their colored rings. Photo credit: David W. Kikuchi.