Coping with Danger and Deception:
Lessons from Signal Detection Theory

Øistein Haugsten Holen[1],* and Tom N. Sherratt[2]

1. Centre for Ecological and Evolutionary Synthesis (CEES), Department of Biosciences, University of Oslo, PO Box 1066 Blindern, NO-0316 Oslo, Norway;   2. Department of Biology, Carleton University, Ottawa, Ontario K1S 5B6, Canada

ABSTRACT: Signal detection theory (SDT) has been used to model optimal stimulus discrimination for more than four decades in evolutionary ecology. A popular standard model that maximizes payoff per encounter was recently criticized for being too simplistic, leading to erroneous predictions. We review a number of SDT models that have received less attention but have explicitly taken repeated encounters into account, focusing on prey choice, mate search, aggressive mimicry, and the aiding of kin. We show how these models can be seen as variants of a second standard model that can be analyzed in a unified framework. In contrast to the simpler model, in this second model a higher probability of an undesirable or dangerous event occurring may either decrease or increase the receiver's acceptance rates. In each instance, the latter outcome requires undesirable events to be undesirable in a relative rather than an absolute sense. Increasing the abundance of desirable signalers or the payoff from accepting them may also either raise or reduce acceptance rates. Our synthesis highlights fundamental similarities among models previously studied on a case-by-case basis and challenges some long-held beliefs. For example, some classic predictions of Batesian mimicry can be reversed when model prey are protected by low profitability rather than harmful defense.

Keywords: signal detection theory, sequential encounters, mate search, optimal foraging, Batesian mimicry, aggressive mimicry.

Introduction

Signal detection theory (SDT) is a widely used framework for analyzing stimulus discrimination tasks (Egan 1975; Green and Swets 1988) that has served as a modeling tool in behavioral ecology and evolutionary ecology for more than 40 years (Oaten et al. 1975; Treisman 1975). Appli-

* Corresponding author; email: o.h.holen@ibv.uio.no.
  ORCIDs: Holen, https://orcid.org/0000-0003-2530-007X.

cations of SDT in these fields include the modeling of Batesian mimicry (e.g., Oaten et al. 1975; Getty 1985; Johnstone 2002; Sherratt 2002; Holen and Johnstone 2004, 2018; McGuire et al. 2006; Speed and Ruxton 2010), aggressive mimicry (Holen and Johnstone 2004, 2006; Kloock and Getty 2019), mate search (Reeve 1989; Wiley 1994; Getty 1995, 1996b; Shizuka and Hudson 2020), honest signaling (Johnstone 1998b; Getty 2002; Holen and Svennungsen 2012), and many other phenomena (e.g., Staddon and Gendron 1983; Getty 1987, 1996a; Johnstone 1998a; Sherratt 2001; Trimmer et al. 2008; Abbott 2010; Abbott and Sherratt 2013; Bogaardt and Johnstone 2016; Metcalf et al. 2017; Trimmer et al. 2017a, 2017b; Sherratt and Holen 2018; Ehlman et al. 2019; McNamara and Trimmer 2019). SDT has also long motivated empirical studies and provided conceptual insights (e.g., Davies et al. 1996; Dor et al. 2007; Laubach et al. 2018).

At its simplest, SDT involves an individual that faces the task of discriminating between two types of circumstances or events, each characterized by different stimuli, and must choose between two different actions. We refer to this individual as the signal receiver (or "receiver" for short). The receiver benefits from accepting desirable events ("hits," or correct acceptances) and rejecting undesirable events (correct rejections). The receiver may, for example, be a predator that must discriminate between undefended prey, which it should accept, and defended prey, which it should reject. Because of noise and inherent variation in stimuli, the receiver cannot with certainty know which situation it faces when choosing the action, and it may therefore make two types of decision error: it may reject desirable events ("misses") or accept undesirable events ("false alarms," or incorrect acceptances). The four outcomes are summarized in table 1, using the convention that the desirable event is denoted A and the less desirable one B. The probabilities of making the two errors are inversely related, so that a reduced rejection rate of desirable events comes at the cost of an increased acceptance rate of undesirable events (Egan 1975; Green and Swets 1988).

Table 1: The four outcomes of the binary decision

	Event A: desirable	Event B: harmful/dangerous or undesirable
Accept	Correct acceptance (CA) or hit	False alarm (FA) or incorrect acceptance
Reject	Miss (MI) or incorrect rejection	Correct rejection (CR)

Note: The receiver experiences an event that may be of type A or B and in response decides between the actions accept and reject, leading to four possible outcomes.

The defining feature of SDT is the receiver operating characteristic (ROC), which gives the maximum "hit rate" (the probability of making correct acceptances in circumstance A) that the receiver is able to obtain for a given "false alarm rate" (the probability of making incorrect acceptances in circumstance B). In the prey example, for instance, the ROC gives the maximum probability of attacking undefended prey that is achievable by the predator as a function of the probability by which it (erroneously) will attack the defended prey type. The exact shape of the ROC depends on the discrimination task and can be recorded empirically by varying the receiver's motivation to accept (Blough 2001; Macmillan and Creelman 2005; Sumner and Sumner 2020). In a model, the ROC is typically derived from assumptions about the probability distributions of the observations associated with desirable and undesirable events (appendix A; appendixes A–E are available online). An SDT model also describes what the receiver attempts to achieve by discriminating—that is, a decision goal that defines a function to be maximized—and the ROC simply acts as a constraint on the maximization problem. In mathematical psychology, the functions to be maximized are usually quite simple, such as maximizing the fraction of correct responses or maximizing expected payoff per encounter (Egan 1975). In evolutionary ecology, we are interested in fitness consequences, and the payoff function to be maximized may be more complex (e.g., Getty 1985; Holen 2013; Trimmer et al. 2017*a*). Nevertheless, the most commonly used economic formulation simply maximizes the receiver's expected payoff per encounter. This single-encounter formulation has been applied to a wide range of ecological scenarios and is mathematically convenient, since the optimal strategy can be determined from the shape of the ROC and a single aggregated parameter (see below). Moreover, each of the four outcomes (table 1) have fixed payoffs, so that the optimal decision in an encounter can be determined without taking other encounters into account.

Despite its popularity and versatility, SDT has recently come in for stark criticism (Trimmer et al. 2017*a*; McNamara and Trimmer 2019). The focus for criticism in these articles was the single-encounter formulation, referred to as "standard" (Trimmer et al. 2017*a*) or "classical" (McNamara and Trimmer 2019) SDT. The single-encounter formulation makes the reasonable prediction that as the probability of a given event increases, individuals should be more likely to make the decision that is appropriate for that event. For example, if the underlying probability of danger or harm increases, individuals should be more cautious and likely to take steps to avoid that danger. However, despite its intuitive appeal, this insight is described as "thoroughly misleading" when applied to ecological scenarios that involve sequential encounters (Trimmer et al. 2017*a*). By modeling the building up of energy reserves under risk of predation, Trimmer et al. (2017*a*) and McNamara and Trimmer (2019) have shown (somewhat counterintuitively) that it can be optimal to respond to a higher overall probability of danger by being *less* cautious. They attribute this finding to the incorporation of sequential encounters in their models and question the soundness of results obtained using the single-encounter formulation. Contrasting their approach to the latter, Trimmer et al. (2017*a*, p. 5) "see the sequential approach as the more useful route ahead in behavioural ecology," while McNamara and Trimmer (2019, p. 19) conclude that the behavioral sciences may "benefit from taking a fresh look at the predictions of models that have simplified scenarios down to a single decision."

Here, we take a fresh look at the use of SDT in evolutionary ecology. We concentrate on economic formulations and contrast the predictions of the single-encounter formulation with previously published models that involve repeated encounters. To this end, we have analyzed signal detection models originally developed to study prey choice (Getty 1985; Stephens and Krebs 1986, pp. 66–72; Holen 2013), the aiding of kin (Reeve 1989), the search for mates and other objects (Reeve 1989; Wiley 1994; Getty 1995, 1996*b*; Shizuka and Hudson 2020), and aggressive mimicry (Kloock and Getty 2019). In addition, we consider a simple model of repeated search under mortality risk. These seven models share enough mathematical similarities to be analyzed in a unified framework, and we see them as variants of a second "standard" model that we have called the "repeated-encounter" (REP) formulation. We characterize the rich behavior of these models and identify broad conditions under which the predicted response to increased danger qualitatively matches that of the novel models (Trimmer et al. 2017*a*; McNamara and Trimmer 2019) but also broad conditions under which it matches that of the single-encounter formulation, suggesting that there

are circumstances under which the latter may be a reasonable approach. The analysis not only brings together several disparate applications under one roof but also provides novel insights into how organisms should respond to increased risk.

The ROC

Let R be the hit rate (the probability of accepting a desirable event A) and y be the false alarm rate (the probability of accepting an undesirable event B). We use the function $\widetilde{R}(y)$, defined on the unit interval, to depict the relationship between the receiver's hit and false alarm rate, that is, its ROC. The ROC graph is always nondecreasing from the point $(0, 0)$ to $(1, 1)$ in the (y, R) plane. The discriminability of two stimuli is a measure of the ability of a receiver to distinguish between them. The measure of discriminability depends on the type of ROC assumed, but the following holds generally: when discrimination is not possible, the ROC is the diagonal line connecting $(0, 0)$ and $(1, 1)$. As discriminability increases, the signal receiver will typically achieve a higher hit rate for a given false alarm rate, and a high-discriminability ROC will therefore lie above a low-discriminability ROC except at its end points (fig. 1). Unless otherwise stated, in our analysis we will make only three assumptions about the ROC, namely, that $\widetilde{R}(y)$ is strictly increasing from $(0, 0)$ to $(1, 1)$, is everywhere differentiable on $(0, 1)$, and is strictly concave (except when discriminability is zero, in which case $\widetilde{R}(y)$ is linear). These three assumptions hold for many ROC types, including the normal-normal equal variance ROC (the standard choice for evolutionary ecology models and for many empirical applications) and the power law ROC (also common in models). We refer to a short introduction in appendix A and the standard literature (Egan 1975) for the derivation of these and other ROCs. Unless otherwise stated, we use normal-normal equal variance ROCs when illustrating results in figures, meaning that observations drawn from desirable and undesirable signalers are assumed to follow two normal distributions with different means but the same variance.

The Single-Encounter Formulation

The single-encounter formulation assumes that the signal receiver maximizes expected payoff per encounter. Let V_{CA}, V_{MI}, V_{FA}, and V_{CR} be the payoffs associated with a correct acceptance (hit), a miss, a false alarm, and a correct rejection, respectively. By definition, correct responses generate higher payoff than errors, so we have $V_{CA} > V_{MI}$ and $V_{CR} > V_{FA}$. Let $\Pr(A) = p$ and $\Pr(B) = (1 - p)$ be the probabilities of events A and B, respectively, assuming $0 < p < 1$.

Indifference curves are isocontours of the payoff function to the receiver and show equally profitable combinations of the hit rate R and the false alarm rate y treated as free variables. Our analysis here and henceforth relies on the shape of these indifference curves and how these intersect with the ROC (Sperling 1984; Wiley 1994; Lynn and Barrett 2014), an approach that allows many results to be derived using only general assumptions about ROC shape. The single-encounter formulation has been analyzed in this way (Sperling 1984; Wiley 1994; Lynn and Barrett 2014), and we follow the same standard approach.
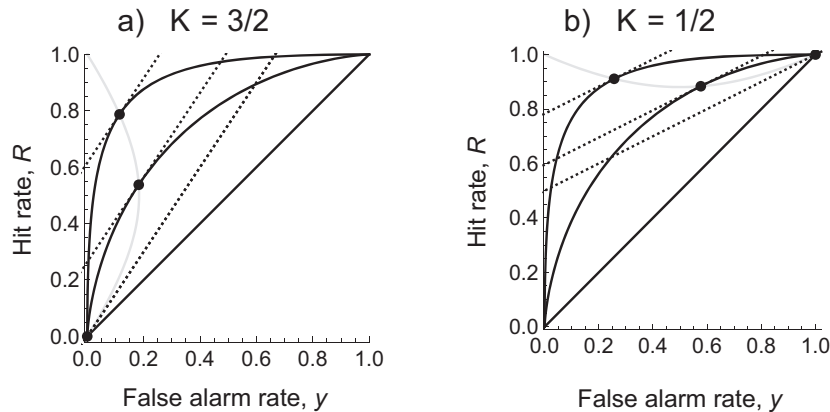

Figure 1: Normal-normal equal variance receiver operating characteristics (ROCs) that differ in discriminability are shown (solid lines; bottom: $d' = 0$; middle: $d' = 1$; top: $d' = 2$). In the single-encounter formulation, the indifference curves (dotted lines) are linear and parallel, with slopes equal to the aggregated parameter K. The optimal receiver strategy (black circle) is the point on the ROC that is intersected by the indifference curve with the highest intercept. The gray curves show how the optimal receiver strategy changes if discriminability d' changes continuously. See the main text for further explanation.

We begin by identifying the expected payoff for the single-encounter formulation, namely,

$$U(y, R) = pRV_{CA} + (1 - p)yV_{FA} + p(1 - R)V_{MI}$$
$$+ (1 - p)(1 - y)V_{CR}. \quad (1)$$

Maximization of this function is subject to the constraints $R = \widetilde{R}(y)$ and $0 \leq y \leq 1$. Letting $U(y, R)$ equal the payoff u and solving for R, we obtain a set of indifference curves that are linear in the (y, R) plane:

$$R = \frac{(1 - p)(V_{CR} - V_{FA})}{p(V_{CA} - V_{MI})}y + \frac{u - (1 - p)V_{CR} - pV_{MI}}{p(V_{CA} - V_{MI})}.$$

These indifference curves all have the positive slope $K = (1 - p)(V_{CR} - V_{FA})/[p(V_{CA} - V_{MI})]$. In SDT terminology, K is also known as the "cutoff" or threshold value of likelihood ratio above which a receiver that maximizes payoff per encounter should accept (e.g., Egan 1975, pp. 18–19; see also appendix B). The higher the payoff u, the higher the intercept of the corresponding indifference curve. The optimal receiver strategy is the point on the ROC that is intersected by the uppermost indifference curve (fig. 1), that is, the one with the highest u. The optimal strategy can be found using this graphical method for any ROC. Let y^* denote an optimal level of false alarm. If $0 < y^* < 1$, then a ROC that is differentiable at y^* will also be tangential to the indifference curve at that point, and we obtain the standard result (e.g., Egan 1975) that the slope of the ROC at the optimal solution equals K.

The classic predictions of the single-encounter model can be inferred from the above expression for K. The receiver will be more likely to accept when the cutoff K is lower (fig. 1), as when the net benefit of accepting in the desirable circumstance A (i.e., $V_{CA} - V_{MI}$) is high relative to the net benefit of rejecting in the undesirable circumstance B (i.e., $V_{CR} - V_{FA}$) or when event A is more common (i.e., p is high). Let n_i represent the (absolute) abundance of a signaler of type i. Assuming equal detection rates, we can reparameterize so that $p = n_A/(n_A + n_B)$, and we see that increasing n_A must change acceptance rates in the same direction as increasing p, while increasing n_B has the opposite effect.

The REP Formulation

Although the single-encounter formulation is a convenient model, it is not always reasonable to focus on one encounter at a time. Models that involve sequential encounters often feature payoffs that cannot be specified in advance but "emerge from the analysis" (Trimmer et al. 2017a). In a search context, for instance, an encountered mate (or prey item) should be accepted only if this leads to higher fitness than would rejection followed by

continued search, but the value of the latter course of action typically depends on the discrimination strategy to be used. We consider sequential models that have payoff functions of the general form

$$U(y, R) = \frac{aR + by + c}{dR + ey + f}, \quad (2)$$

where a, b, c, d, e, and f are constants or composite parameters that are independent of R and y and that have ecological interpretations that vary greatly between the applications we will explore later. As in the single-encounter formulation, the payoff function (2) is maximized with respect to R and y, subject to the constraints $R = \widetilde{R}(y)$ and $0 \leq y \leq 1$. Whenever the receiver encounters both desirable and undesirable events with positive probability (i.e., $0 < p < 1$), we assume that the following two criteria will be satisfied:

$$\begin{aligned} ae &> bd, \\ af &> cd. \end{aligned} \quad (3)$$

These two inequalities are sufficient conditions for the payoff to increase with the hit rate when the false alarm rate is kept constant (i.e., for $\partial U/\partial R > 0$). A model of form (2) with parameters that satisfy (3) is an example of the REP formulation. Setting $U(y, R) = u$ and rearranging, we see that the indifference curves of (2) are linear in the (y, R) plane:

$$R = \frac{eu - b}{a - du}y + \frac{fu - c}{a - du}. \quad (4)$$

If we extend the indifference curves outside the unit square, it is readily shown that they will all meet at the same point (k_1, k_2) in the (y, R) plane (appendix C):

$$(k_1, k_2) = \left(\frac{cd - af}{ae - bd}, \frac{bf - ce}{ae - bd}\right). \quad (5)$$

The first inequality of (3) ensures that the slope of the indifference curves (4) increases with u (appendix C). Together, the two inequalities of (3) ensure that k_1 is negative. REP models therefore have the following graphical solution: the optimal strategy is found at the point on the ROC that is touched by the uppermost straight line that can be drawn from the point (k_1, k_2) to the ROC (fig. 2; see also Holen [2013, fig. S1] and Kloock and Getty [2019, fig. 2]). The optimal strategy can be found using this graphical method for any ROC. If the optimal strategy is in the interior of the unit interval (i.e., $0 < y^* < 1$) and the ROC is differentiable at that point, the line will also be tangential to the ROC (fig. 2).

Note that the single-encounter formulation (1) can be put in the general form $U(y, R) = aR + by + c$, with a, b, and c being composite parameters. We therefore expect the REP formulation to more closely approximate
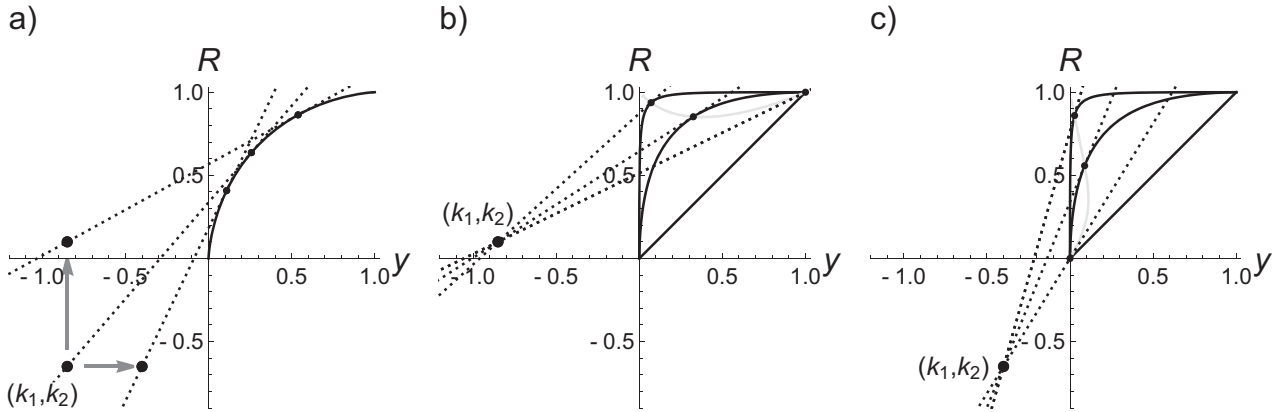
Figure 2: Normal-normal equal variance receiver operating characteristics (ROCs) are shown (solid lines), with R denoting the hit rate and y denoting the false alarm rate. The indifference curves of the repeated-encounter (REP) formulation are linear, and when extended outside the unit square (dotted lines) they meet at the point (k_1, k_2), with k_1 and k_2 being aggregated parameters. The optimal receiver strategy (small circles on ROCs) is the point on the ROC that is intersected by the indifference curve with the highest intercept. *a*, Increasing k_1 moves the point (k_1, k_2) to the right and moves the optimal receiver strategy downward on the ROC, decreasing acceptance rates. Increasing k_2 moves the point (k_1, k_2) upward and the optimal receiver strategy upward on the ROC, increasing acceptance rates. *b, c*, Optimal receiver strategies (small circles on ROCs) are shown for different (k_1, k_2), in each case using three different normal-normal equal variance ROCs (solid lines; bottom: $d' = 0$; middle: $d' = 1.5$; top: $d' = 3$). The gray curves show how the optimal receiver strategy changes if discriminability d' changes continuously. See the main text for further explanation.

the single-encounter formulation when the magnitude of f is large relative to d and e (see eq. [2]).

We will now analyze how the optimal solution to the general REP model changes with parameters before moving on to consider specific applications.

The Effect of Small Changes in k_1 and k_2 on Acceptance Rates. Each model application (to be discussed below) contains a number of ecological parameters that each are part of at least one of the composite parameters a to f. A change in one of these ecological parameters may potentially change both k_1 and k_2 and thus the optimal acceptance rates (hit rate and false alarm rate). We study this using graphical arguments. Suppose the optimal strategy y^* satisfies $0 < y^* < 1$, and consider the straight line drawn through the point (k_1, k_2) and the point $(y^*, \bar{R}(y^*))$ (fig. 2*a*): if a change in an ecological parameter moves the point (k_1, k_2) above the original line, the receiver will as a result become less restrictive and accept more often (it moves toward "adaptive gullibility"; Wiley 1994), but if the parameter change moves (k_1, k_2) below this line, the receiver will become more restrictive and accept less often (it moves toward "adaptive fastidiousness"; Wiley 1994).

The signs of the partial derivatives of k_1 and k_2 with respect to the parameter in question often suffice to decide whether (k_1, k_2) moves above or below the original line. Either sign can be positive, zero, or negative. If the signs differ and either k_1 increases with the parameter change

or k_2 decreases, the receiver will accept less often (fig. 2*a*). If the signs differ and either k_1 decreases with the parameter change or k_2 increases, the receiver will accept more often (fig. 2*a*). If the partial derivatives have the same (nonzero) sign, however, further exploration will be necessary (some examples are given in appendixes D and E).

For ROCs that satisfy our three assumptions, the optimal solution y^* will change gradually with changes in (k_1, k_2). However, for some other types of ROCs, the solution may instead "jump" discontinuously in the predicted direction. A notable example is that of ROCs based on discrete-valued random variables (Egan 1975; Getty 1995, 1996*a*, 1996*b*). Such "discrete" ROCs consist of connected line segments, and the optimal solution of the REP model will be on a vertex where the ROC is not differentiable. A sufficiently small change in (k_1, k_2) will not change the optimal solution, while a larger change may move the optimal solution to another vertex. Another example is that of ROCs that are "improper" sensu Egan (1975), which may have convex sections (see appendix A). The optimal solution of the REP model can never be in a convex section in the interior of the ROC, and a gradual change in (k_1, k_2) may therefore in some cases cause the optimal solution to "jump" discontinuously over an interval that contains a convex section (see some results of Reeve 1989, which implicitly featured improper ROCs). In the remainder of the article, we restrict ourselves to discussing ROCs that are everywhere differentiable and strictly concave.

Finally, if the optimal strategy y^* is found at one of the end points of the ROC, a small change in k_1 and k_2 may not affect the receiver strategy.

The Effect of Changes in Ecological Parameters on Acceptance Rates. Following this elucidation of the effects of changing (k_1, k_2) on acceptance rates, we can explore the effect of small changes in ecological parameters that affect one or more of the composite parameters a to f. It is beyond the scope of our article to make an exhaustive analysis of all ecological parameters in all applications; we focus on those that allow us to compare the REP applications to the single-encounter model. In each of the seven applications we consider, all parameters a to f are positive or zero except b and c, which can be negative.

We first consider the effects of base rate p on acceptance rates, which played such a central role in recent SDT models (Trimmer et al. 2017*a*; McNamara and Trimmer 2019). All seven model applications are formulated either in terms of p or in terms of absolute abundances n_A and n_B, in which case p may enter via the substitutions $n_A = np$ and $n_B = n(1 - p)$ (see below). In most of the applications, p is a factor in a and d, while $(1 - p)$ is a factor in b and e; and in all of the applications, the denominators of k_1 and k_2 are proportional to $(1 - p)$ and p, respectively (after canceling terms). The partial derivatives of k_1 and k_2 with respect to p sometimes have the same sign, and we therefore study the shape of the parametric curve $(k_1(p), k_2(p))$ (appendix D). In every application, the parametric curve is the left branch of a hyperbola with vertical and horizontal asymptotes. Let $Q = [bf - ce]_{p=0} = [bf - ce]_{n_A=0, n_B=n}$. In the limit where p tends to zero from above, k_2 tends to positive infinity if $Q > 0$ and toward negative infinity if $Q < 0$ (we ignore the knife-edge case $Q = 0$). In the first case ($Q > 0$) we say that the less desirable event B is undesirable only in a relative sense (it is always accepted in the absence of A events), and in the second case ($Q < 0$) we say that event B is undesirable in an absolute sense (it is never acceptable). We obtain the following results (appendix D). When event B is undesirable in a relative sense, there is a trivial case in which acceptance rates are constant and equal to 1 on the unit interval $p \in [0, 1]$. There is also a nontrivial case in which acceptance rates as a function of p have a U-shaped dip over a subinterval of the unit interval and are constant and equal to 1 elsewhere. In contrast, when event B is undesirable in an absolute sense ($Q < 0$), acceptance rates will either increase strictly from 0 to 1 over the entire unit interval (e.g., typical for the normal-normal equal variance ROC) or increase strictly from 0 to 1 on a subinterval of the unit interval and remain constant outside this (see appendix D for details).

In addition to p, we have explored the effect on acceptance rates of small changes in ecological parameters that affect one (or two) of the composite parameters a to f (appendix E). The analysis is restricted to strategies satisfying $0 < y^* < 1$. We summarize our findings in table 2; further discussion is given under the various applications. A change in an ecological parameter that is only part of a, is only part of d, or is a factor in both a and d but not part of other composite parameters will have an effect that depends on the sign of k_2 (table 2). In every application except 3, the sign of k_2 is always the same as the sign of Q (appendix D).

We also characterize the effect of changing discriminability on optimal receiver strategies. Since different ROC types have different measures of discriminability, we restrict ourselves to the most common ROC type, the normal-normal equal variance ROC, which has the best-known measure of discriminability, d' (see appendix A and Egan 1975 for details). We ignore the trivial case $k_2 \geq 1$, in which rejection is never optimal. Changes in d' affect the ROC shape but not k_1 or k_2, and so its effects are independent of the specific application. Suppose that $k_1 < k_2$, in which case "always accept" ($y^* = 1$) is the optimal strategy when discriminability is zero: numerical explorations show that as discriminability d' increases from zero toward infinity (perfect discriminability), the false alarm rate decreases from 1 to 0 while the hit rate first decreases and then increases again (see the example in fig. 2*b*). In contrast, suppose $k_1 > k_2$, in which case "always reject" is the optimal strategy when discriminability is zero: numerical explorations show that as discriminability d' increases from zero toward infinity (perfect discrimination), the hit rate increases from 0 to 1 while the false alarm rate first increases and then decreases (see the example in fig. 2*c*).

Table 2: Effect of parameter changes on acceptance rates

Small increase in parameter	Assumptions	Effect on acceptance rates	
		$k_2 > 0$	$k_2 < 0$
a		−	+
b		+	+
c		−	−
d	$c \geq 0$	+	−
e	$c \geq 0$	−	−
f	$c \geq 0$	+	+
δ	$a = \hat{a}\delta, d = \hat{d}\delta$	−	+
δ	$b = \hat{b}\delta, e = \hat{e}\delta$	−	−
δ	$c = \hat{c}\delta, f = \hat{f}\delta$	+	+

Note: The effects of small increases in parameters of the repeated-encounter model on the acceptance rates of the receiver are listed. These results are derived in appendix E and apply when $0 < y^* < 1$. The positive parameter δ is used to represent changes in pairs of parameters. In addition to the assumptions listed in the second column, it is assumed that $a > 0$, $e > 0, f \geq 0$, and $d \geq 0$.

Applications of the REP Formulations

We proceed to review the seven REP applications. To facilitate presentation, we keep the same notation for ecological parameters across these models. Each model has the form (2). We interpret the parameters and explain why the two criteria given by (3) hold (given that desirable and undesirable events both occur with positive probability). We determine the sign of Q and explore model behavior by focusing on how an optimal strategy $0 < y^* < 1$ will be affected by small changes in parameters. Examples of how acceptance rates change with p are shown in figure 3 (all applications); a more systematic exploration is shown for the false alarm rate in figure 4 (all applications except 3).

Application 1: Batesian Mimicry—A Two-Species Model. When modeling prey choice, a drawback with the single encounter model is that it does not take handling time

into account. The disk equation does this (Holling 1959; Stephens and Krebs 1986). A model combining SDT with a two-species disk equation was discussed in Stephens and Krebs (1986, pp. 66–72) in connection to the marginal value theorem, and a very similar model was presented by Reeve (1989) in the context of aiding relatives (the "repeated-search" model).

A predator forages on a Batesian mimic (prey type A) and its model (prey type B). Let s denote area searched per unit time. Using subscripts A and B to refer to prey type, let n denote the number of prey per unit area, h the expected handling time, and E the expected energy intake per attacked prey. All are positive constants except E_B, which we (following Getty 1985; Holen 2013) allow to be positive or negative, the latter reflecting the possibility that net energy is spent repairing damage after handling defended prey. Handling time for model prey may include the time needed to recover from ingesting toxins or from being stung by the prey before search can resume
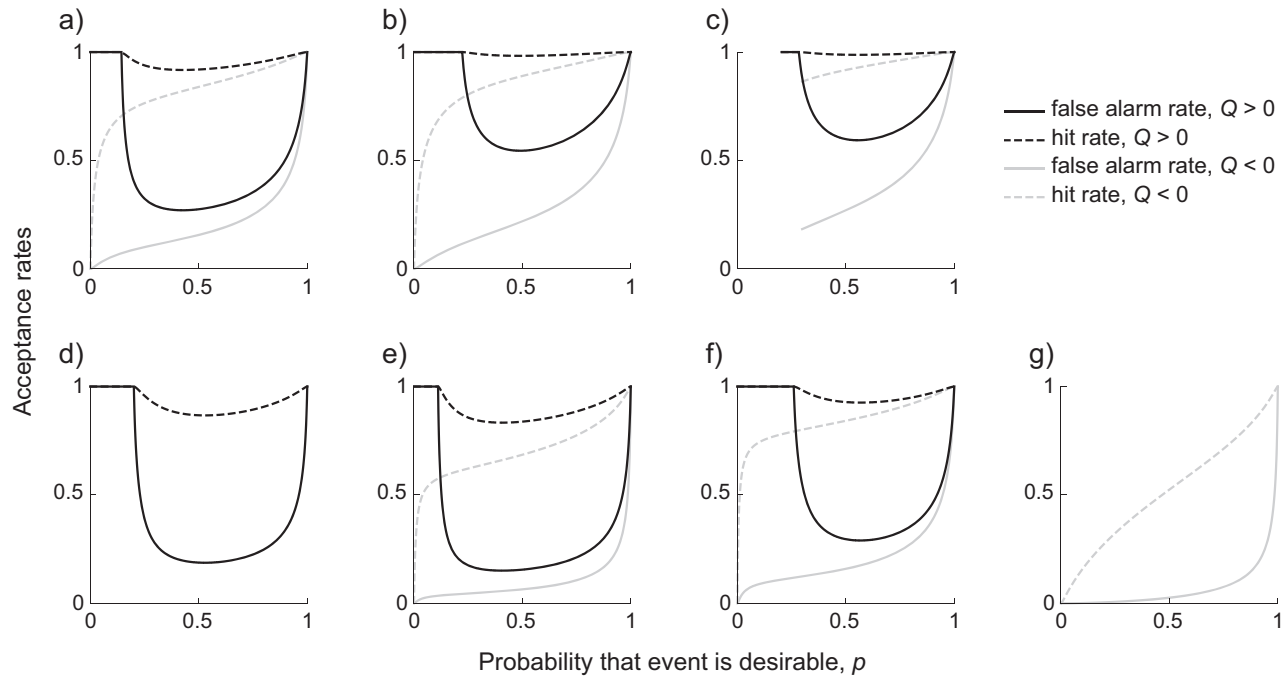

Figure 3: Acceptance rates (hit rate and false alarm rate) are shown as a function of p, the probability that an event is desirable. Each panel shows possible outcomes for one of the seven model applications in the main text (see the main text for the exact parameter values used). When parameters are chosen so that $Q > 0$, the acceptance rates as a function of p are either constant and equal to 1 (trivial case, not shown) or have a U-shaped dip (shown). In contrast, when parameters are chosen so that $Q < 0$, the two acceptance rates increase with p from 0 to 1. Model applications: *a*, two-species disk equation/Reeve (1989) repeated-search model; *b*, Getty (1985) model; *c*, Holen (2013) taste rejection model (the acceptance rates are not plotted in the low range of p, in which the mimicry complex drops out of the optimal diet); *d*, Reeve (1989) search-and-settle model (in this application, Q cannot be negative); *e*, Wiley (1994) search model; *f*, repeated search under mortality risk; *g*, Kloock and Getty (2019) model of aggressive mimicry (in this application, Q cannot be positive). In each panel, we use the normal-normal equal variance receiver operating characteristic (ROC) with discriminability $d' = 2$. See the main text and appendix D for a discussion of other ROC types.
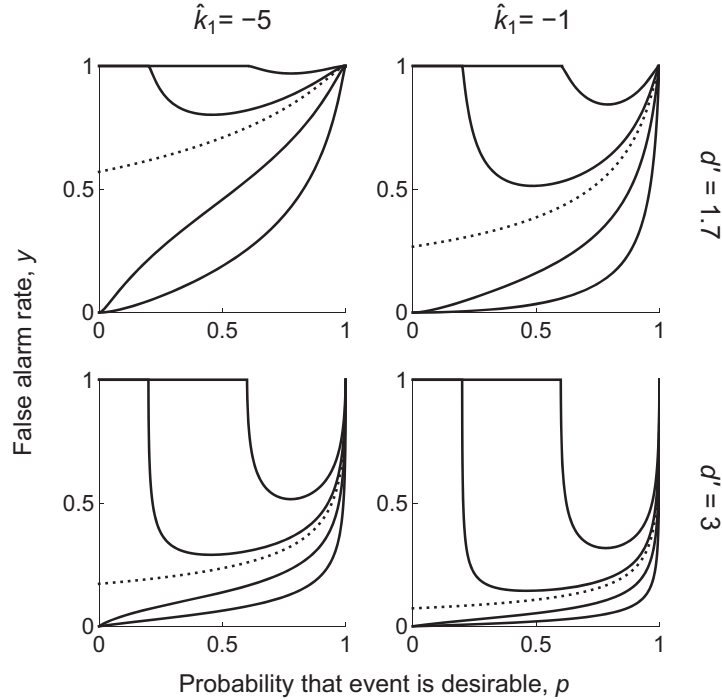
Figure 4: In all model applications except (3), k_1 and k_2 have the forms $k_1 = \hat{k}_1/(1 - p)$ and $k_2 = \hat{k}_2/p$, with $\hat{k}_1$ and $\hat{k}_2$ being (aggregated) parameters that do not depend on p. Each panel shows false alarm as a function of p for (from top to bottom) $\hat{k}_2 = 0.6$ (solid line), $\hat{k}_2 = 0.2$ (solid line), $\hat{k}_2 = 0$ (dotted line), $\hat{k}_2 = -1$ (solid line), and $\hat{k}_2 = -5$ (solid line). The U shape becomes deeper when d' is high, when $\hat{k}_1$ is close to zero, and when $\hat{k}_2$ is close to zero (but positive). For the normal-normal equal variance receiver operating characteristic used here, acceptance rates will be constant and equal to 1 when $\hat{k}_2 \geq 1$.

(Holen 2013). The payoff function (representing energy intake rate) is

$$\frac{s n_A E_A R + s n_B E_B y}{1 + s n_A h_A R + s n_B h_B y}. \quad (6)$$

Note how this function has the general form (2), with (composite) parameters $a = s n_A E_A$, $b = s n_B E_B$, $c = 0$, $d = s n_A h_A$, $e = s n_B h_B$, and $f = 1$. We assume $E_A/h_A > E_B/h_B$ (so that a Batesian mimic provides more energy per unit time than a model). Using this fact, it is straightforward to check that inequalities (3) hold. From (5) we obtain

$$(k_1, k_2) = \left(\frac{-E_A}{s n_B (E_A h_B - E_B h_A)}, \frac{E_B}{s n_A (E_A h_B - E_B h_A)} \right). \quad (7)$$

The signs of k_2 and Q ($= s n E_B$) are here the same as the sign of E_B. Note that if model prey are undesirable in a relative sense ($Q > 0$), decreasing n_A will increase k_2, and there will be a threshold value of n_A below which it will be beneficial to attack both mimics and models indiscriminately.

Table 2 lists the effects of small changes in parameters on acceptance rates, given that it is optimal to discrim-

inate (i.e., $0 < y^* < 1$). As one might expect, an increase in the abundance of model prey (n_B), a decrease in their energy content (E_B), or an increase in their handling time (h_B) will favor reduced acceptance rates in receivers. The results for the mimic parameters depend on the sign of k_2: if $k_2 > 0$, an increase in the abundance of mimics (n_A), an increase in their energy content (E_A), or a reduction in their handling time (h_A) will counterintuitively favor reduced acceptance rates. In contrast, if $k_2 < 0$, these changes to the mimic parameters have the opposite effect and will favor increased acceptance rates.

We can explore the effect of increasing the fraction p of prey in the mimicry complex that are mimics while keeping overall abundance n constant by making the substitutions $n_A = np$ and $n_B = n(1 - p)$ in (6). An example is shown in figure 3*a*. Parameters $sn = 1$, $E_A = 1$, $h_A = 1$, $E_B = 0.25$, and $h_B = 2$ are used to illustrate the case $Q > 0$. Parameters used to illustrate $Q < 0$ are the same except $E_B = -0.25$.

Reeve (1989) focused on recognition within animal societies, and one of the explored scenarios (the repeated-search model) involved a receiver searching for and providing aid to colony members/kin. It was beneficial to

help one type of recipient (e.g., kin) but not so much another type (e.g., more distant kin or intruders). The fitness expression of Reeve (1989) is essentially a rescaled and simplified version of (6) times a constant (the total time devoted to search and aid), obtained by setting $s = 1$, $h_A = h_B = h$, $n_A = np$, $n_B = n(1 - p)$ and then dividing every term in the numerator and denominator by n, so that f is the expected waiting time before encountering a signaler. The payoffs E_A and E_B are interpreted as the benefits of aiding recipients rather than the benefits of consuming them. The general predictions are the same as in the prey model even if interpretations differ. The relationship between false alarm rate and p shown in figure 6b in Reeve (1989) has a U shape similar to the one in figure 3a for $Q > 0$ (this article), even if the choice of parameters and ROC type differs.

Application 2: Batesian Mimicry with Alternative Prey. Getty (1985) analyzed a more complex scenario than the one described above, in which the forager has additional prey types that it may incorporate into its diet. These additional prey can be easily discriminated, and g_i denotes the probability of attacking alternative prey type $i \in (1, ..., k)$ on encounter. The payoff function representing the long-term energy intake rate is (in our notation)

$$\frac{s n_A E_A R + s n_B E_B y + \sum_{i=1}^{k} s n_i E_i g_i}{1 + s n_A h_A R + s n_B h_B y + \sum_{i=1}^{k} s n_i h_i g_i}. \tag{8}$$

There is no coefficient g for the mimicry complex, since decisions about whether to attack mimics and models are based on appearance and thus controlled by the acceptance rates. Using the shorthand notation $E_{alt} = \sum_{i=1}^{k} s n_i E_i g_i$ and $h_{alt} = \sum_{i=1}^{k} s n_i h_i g_i$, this model differs from (6) in that we have $c = E_{alt}$ and $f = 1 + h_{alt}$. As before, all constants are positive except E_B, which may be positive or negative.

We consider a forager that adopts an optimal strategy (i.e., that subject to the constraint $R = \tilde{R}(y)$ maximizes [8] over y and coefficients g_i) that satisfies $0 < y^* < 1$, and with E_{alt}^* and h_{alt}^* referring to the values of E_{alt} and h_{alt} for the optimal set of alternative prey determined by the g_i coefficients. We ignore the case in which parameters balance on a knife edge so that the predator would be indifferent about whether to include a given alternative prey type in the diet. Thus, the g_i coefficients are all 0 or 1 (the zero-one rule; Stephens and Krebs 1986) and associated with strictly higher payoff than any alternative value. For mimics to be attacked, the profitability of a single mimic must be greater than the energy intake rate from the optimal set of alternative prey. Thus, the inequality $E_A/h_A > E_{alt}^*/(1 + h_{alt}^*)$ holds, which together with $E_A/h_A > E_B/h_B$ ensure that inequalities (3) hold.

From (5) we obtain

$$(k_1, k_2) = \left(-\frac{E_A(1 + h_{alt}^*) - E_{alt}^* h_A}{s n_B(E_A h_B - E_B h_A)}, \frac{E_B(1 + h_{alt}^*) - E_{alt}^* h_B}{s n_A(E_A h_B - E_B h_A)} \right). \tag{9}$$

The signs of k_2 and Q are the same as the sign of $E_B/h_B - E_{alt}^*/(1 + h_{alt}^*)$. Note that when model prey are undesirable in a relative sense ($Q > 0$), they have higher profitability than the energy intake rate for the optimal set of alternative prey and are profitable in the absence of mimics.

What are the effects of small changes in parameters, given that it is optimal to discriminate (i.e., $0 < y^* < 1$)? We can consider changes in parameters that are small enough that they change y^* (which is in the interior of the unit interval) but not the optimal g_i coefficients (which remain on the boundary of the unit interval at 0 or 1). Referring to table 2, a small change in n_A, n_B, E_A, h_A, E_B, or h_B changes receiver acceptance rates in the same direction as in the two-species scenario (application 1), with effects of a change in the mimic parameters depending on the sign of k_2 in exactly the same way (although the expression for k_2 is now given by [9]). Predictions obtained for $k_2 < 0$ are consistent with those reported in the numerical explorations of Getty (1985).

As in the two-species model, we may make the substitutions $n_A = np$ and $n_B = n(1 - p)$ and explore the effect of increasing the fraction of prey in the mimicry complex that are mimics. We illustrate changes in p in the example in figure 3b, under the assumption that there is only one type of alternative prey. For the case $Q > 0$ we used parameters $sn = 0.5$, $E_A = 1$, $h_A = 2$, $E_B = 0.4$, $h_B = 2$, $E_{alt}^* = 0.333$, and $h_{alt}^* = 1$. The same parameters were used for $Q < 0$ except $E_B = 0.3$ and $h_B = 3$. In both examples, it is optimal for the forager to include the alternative prey in the diet (not shown). We refer to Getty (1985) for explorations of how different alternative prey types may fall out of the diet.

Application 3: Taste Rejection of Prey. Holen (2013) studied prey choice, assuming that mimics and models were of identical appearance but could be discriminated against on the basis of taste after capture, leading to consumption or rejection. A disk equation model was used with the handling times of mimics and models split into two components: the expected time spent on pursuit and (if the prey is accepted on the basis of its taste) the expected time from start of consumption until search could be resumed (notation below). The payoff function is

$$\frac{(s n_A E_A R + s n_B E_B y) g_{AB} + s n_C E_C g_C}{1 + (s n_A (h_{Ap} + h_A R) + s n_B (h_{Bp} + h_B y)) g_{AB} + s n_C h_C g_C}. \tag{10}$$

Here, there is a single alternative prey with subscript C; coefficients g_{AB} and g_C are the respective probabilities of keeping the mimicry complex and the alternative prey in the diet; parameters h_{Ap} and h_{Bp} are the expected time spent pursuing a mimic and model, respectively; and h_A and h_B are the expected time spent consuming them on capture, including recovery time. We will use the notational shorthand $E_{alt} = s\,n_C E_C g_C$ and $h_{alt} = s\,n_C h_C g_C$. In addition, the model by Holen (2013) had capture rates q specific to mimics and models (which can be absorbed into E_A, h_A, E_B, and h_B) and prey-specific search efficiencies reflecting, for example, prey conspicuousness (which can be absorbed into the prey abundances).

We consider a forager that adopts an optimal strategy satisfying $g_C = g_{AB} = 1$ (meaning that all prey are attacked, although some may be subsequently rejected) and $0 < y^* < 1$ (meaning that a fraction of mimics and models are consumed). The model thus differs from application 1 in that $c = E_{alt}^*$ and $f = 1 + s\,n_A h_{Ap} + s\,n_B h_{Bp} + h_{alt}^*$, with E_{alt}^* and h_{alt}^* referring to the values of E_{alt} and h_{alt} for the optimal strategy. Analogous to application 2, we assume that $E_A/h_A > E_{alt}^*/(1 + h_{alt}^*)$ holds (or it would be unprofitable to consume mimics when having alternative prey in the diet), which implies $E_A/h_A > E_{alt}^*/(1 + s\,n_A h_{Ap} + s\,n_B h_{Bp} + h_{alt}^*)$. This together with $E_A/h_A > E_B/h_B$ ensures that inequalities (3) hold. From (5) we obtain

$$(k_1, k_2) = \left(-\frac{E_A(1 + h_{alt}^* + s\,n_A h_{Ap} + s\,n_B h_{Bp}) - E_{alt}^* h_A}{s\,n_B(E_A h_B - E_B h_A)}, \right.$$
$$\left. \frac{E_B(1 + h_{alt}^* + s\,n_A h_{Ap} + s\,n_B h_{Bp}) - E_{alt}^* h_B}{s\,n_A(E_A h_B - E_B h_A)} \right)$$

(11)

(see Holen 2013, appendix).

From table 2, we see that the effect of changing E_A, h_A, E_B, or h_B on acceptance rates is as in the two previous applications, with the effect of changing E_A and h_A as before depending on the sign of k_2 (appendix E). However, results regarding how acceptance rates are affected by changes in parameters n_A and n_B may differ. This is because parameters reflecting abundance/frequency now also appear in f because of time costs of pursuit. A full analysis of the model is beyond the scope of the current article, but making substitutions $n_A = np$ and $n_B = n(1 - p)$ we have analyzed the application regarding p (appendix D) and have found that it behaves similarly to the others. Note that $Q = ns(E_B(1 + snh_{Bp} + h_{alt}^*) - E_{alt}^* h_B)$. We thus have $Q > 0$ if $E_B/h_B > E_{alt}^*/(1 + snh_{Bp} + h_{alt}^*)$. The latter inequality means that a forager that attacks alternative prey and models but consumes only the former would increase the intake rate if consuming a captured

model. We illustrate changes in p in the example in figure 3c. Parameters corresponding to $Q > 0$ are $sn = 0.5$, $E_A = 1$, $h_A = 2$, $E_B = 0.35$, $h_B = 2$, $E_{alt}^* = 0.333$, $h_{alt}^* = 1$, $h_{Ap} = 0.5$, and $h_{Bp} = 1$. Parameters corresponding to $Q < 0$ are the same except $E_B = 0.1$. Note that at low p, the mimicry complex drops out of the diet since it incurs costs of pursuit for all mimics and models regardless of taste discrimination strategy; acceptance rate curves are not shown in this region (fig. 3c). Explorations of how prey types may fall out of the optimal diet is outside the scope of this article.

Application 4: Search Terminated by Acceptance, with Additive Costs. For other types of behavior that involve sequential encounters, the appropriate fitness measure to be maximized may not be a rate. For instance, sequential search is sometimes terminated as soon as the first acceptance occurs, with the final choice determining the benefit to be obtained. In the "search-and-settle" model (Reeve 1989; Getty 1995, 1996b; Shizuka and Hudson 2020), the receiver seeks to identify a desirable signaler (e.g., a mate, a colony to rob) among many and then settles with that choice. The receiver encounters the candidates sequentially and has two actions at its disposal: it can accept and terminate the search or reject and continue the search. For each candidate explored, an additive cost C (>0) is paid. Settling for a candidate yields the payoff V_{CA} (>0) if the candidate is of the desirable type and V_{FA} $(<V_{CA})$ if it is of the less desirable type. In our notation, the payoff function is (Reeve 1989; Getty 1996b)

$$\frac{pV_{CA}R + (1 - p)V_{FA}y - C}{pR + (1 - p)y}.$$

(12)

We note that the pair of inequalities in (3) hold, given that $V_{CA} > V_{FA}$ and $-Cp < 0$. We proceed to obtain

$$(k_1, k_2) = \left(\frac{-C}{(1 - p)(V_{CA} - V_{FA})}, \frac{C}{p(V_{CA} - V_{FA})} \right).$$

(13)

When search becomes more costly (C increases) or the relative benefit of mating with an optimal candidate (i.e., $V_{CA} - V_{FA}$) decreases, then k_1 decreases while k_2 increases, and the receiver becomes less restrictive and accepts any given mate with a higher probability.

The effect of changing p on acceptance rates (appendix D) is shown in figure 3d. Parameters are $V_{CA} = 1$, $V_{FA} = 0.5$, and $C = 0.1$. Note that this model has the peculiar feature that Q and k_2 are necessarily positive: even if V_{FA} is negative, the sign of the numerator of k_2 (i.e., $bf - ce$) is nevertheless positive, since $c < 0$ and $f = 0$. The less desirable mates cannot be undesirable in an absolute sense, since it is not possible to do worse

than by always rejecting, which leads to an infinite number of additive costs. Because of this, the payoff function (12) is not defined at the origin in the (y, R) plane.

Application 5: Search Terminated by Acceptance, with Multiplicative Costs/Discounting. Wiley (1994) discusses a model of search similar to the search-and-settle model of Reeve (1989) but with multiplicative instead of additive costs. As before, the receiver can accept and terminate the search or reject and continue the search, and settling for a desirable candidate yields the payoff V_{CA} (>0) and an undesirable candidate the payoff V_{FA} ($<V_{CA}$). For each rejection, future payoffs are discounted by a multiplicative factor $0 < \lambda < 1$, which could, for example, represent the probability of finding another candidate if rejecting the current one or the devalued expected payoff for each choice at the next opportunity (since, e.g., mating early often is more valuable than mating late). Wiley's model is derived from the equation $U = pRV_{CA} + (1 - p)yV_{FA} + (p(1 - R) + (1 - p)(1 - y))\lambda U$, which solving for U and rearranging yields

$$U = \frac{pV_{CA}R + (1 - p)V_{FA}y}{1 - \lambda + p\lambda R + (1 - p)\lambda y}. \qquad (14)$$

This is formally similar to the disk equation with two prey (6). Wiley (1994) does not state the full model (14) but states the expression for the slope of the payoff isocontours, which is readily calculated using (4). We note that inequalities (3) hold for $0 < p < 1$, since $V_{CA} > V_{FA}$. We obtain

$$(k_1, k_2) = \left(-\frac{(1 - \lambda)}{\lambda} \frac{V_{CA}}{(1 - p)(V_{CA} - V_{FA})}, \right.$$
$$\left. \frac{(1 - \lambda)}{\lambda} \frac{V_{FA}}{p(V_{CA} - V_{FA})} \right). \qquad (15)$$

Here, k_2 (and Q) has the sign of V_{FA}. If V_{FA} is positive, the less desirable candidate is undesirable only in a relative sense; it is a better choice than not accepting any candidate at all. If V_{FA} is negative, accepting the less desirable candidate is worse than not accepting. Pairing up with, say, a heterospecific mate could easily be worse than not pairing up at all, since it may lead to zero offspring and carry costs in terms of reduced future reproductive value. Regarding λ, from (15) we can infer that $k_1 = (-pV_{CA}/[(1 - p)V_{FA}])k_2$ and thus that k_1 stays proportional to k_2 as λ is changed. The point (k_1, k_2) moves in a straight line from the origin and into the second or third quadrant of the (y, R) plane as λ tends from 1 to 0. The receiver therefore becomes less restrictive as search becomes more costly, that is, as λ is decreased (fig. E1 provides a graphical argument; figs. A1, D1–D3, E1 are available online).

The effect of changing p on acceptance rates (appendix D) depends on the sign of k_2 and is shown in figure 3*e*. Parameters corresponding to $Q > 0$ are $V_{CA} = 1$, $V_{FA} = 0.5$, and $\lambda = 0.9$. Parameters corresponding to $Q < 0$ are the same except $V_{FA} = -0.167$. Other results are given in appendix E and summarized in table 2.

Application 6: Repeated Search under Mortality Risk (Worked Example). In some forms of search, benefits are accrued each time a correct acceptance is made, but a search may be terminated when a false alarm occurs. Using multiplicative costs (i.e., discounting the future), we can construct a simple model of repeated search under mortality risk. The receiver does not settle with its choice but continues searching for new benefits until killed or otherwise incapacitated. Examples might include males searching to obtain (multiple) matings, females searching to find (multiple) resources in which to lay eggs, and so on. Accepting yields payoffs V_A (>0) and V_B in circumstances A and B, respectively, but in circumstance B there is also a probabilistic risk $z > 0$ of dying (e.g., because of predation) before the benefit is obtained. If event B represents an aggressive mimic, it could be reasonable to assume $V_B \leq 0$, with negative values reflecting possible fitness loss even if surviving, but if event B is the same kind of event as A except that it is associated with mortality risk, then it could be reasonable to assume $V_B > 0$. The payoff U satisfies the equation

$$U = pR(V_A + \lambda U) + (1 - p)y(1 - z)(V_B + \lambda U)$$
$$+ p(1 - R)\lambda U + (1 - p)(1 - y)\lambda U,$$

which solved for U yields

$$U = \frac{pV_A R + (1 - p)V_B(1 - z)y}{1 - \lambda + z(1 - p)y\lambda}. \qquad (16)$$

Assuming a nonzero risk of dying ($z > 0$) and noting the fact that $d = 0$, conditions (3) hold for $0 < p < 1$. We obtain

$$(k_1, k_2) = \left(-\frac{1 - \lambda}{\lambda} \frac{1}{z(1 - p)}, \frac{(1 - \lambda)}{\lambda} \frac{(1 - z)V_B}{zpV_A} \right). \qquad (17)$$

The sign of k_2 (and Q) is the same as the sign of V_B. We illustrate changes in p in the example in figure 3*f*. Parameters corresponding to $Q > 0$ are $V_A = V_B = 1$, $\lambda = 0.9$, and $z = 0.30$. Parameters corresponding to $Q < 0$ are the same except $V_B = -0.1$. Other results are given in appendix E and summarized in table 2.

Application 7: A model of Aggressive Mimicry. Kloock and Getty (2019) considered a forager that must discriminate

between food items (models) and aggressive mimics that resemble the food items and can kill the predator. It is assumed that the forager maximizes the ratio of feeding rate to mortality rate. The feeding rate (our notation) is $F = npqV_{CA}R + A$, with n being the combined encounter rate for mimics and models, q the probability of capturing a model given correct classification, and A the feeding rate from alternative prey. Note that here the models are desirable and the mimics harmful; thus, p denotes the fraction of models and V_{CA} the value of ingesting a desirable model. The mortality rate is $\mu = n(1 - p)zy + m$, where z (>0) is the probability of being killed when attacking a mimic and m (>0) is background mortality. Assuming that natural selection maximizes F/μ (Kloock and Getty 2019), this leads to the fitness formulation

$$\frac{F}{\mu} = \frac{npqV_{CA}R + A}{n(1 - p)zy + m}. \tag{18}$$

To see that inequalities (3) hold for $0 < p < 1$, note that $ae > bd$ is equivalent to $n^2p(1 - p)qz > 0$, which holds, and that $cd < af$ is equivalent to $0 < npqm$, which holds. We obtain

$$(k_1, k_2) = \left(-\frac{m}{n(1 - p)z}, -\frac{A}{npqV_{CA}}\right) \tag{19}$$

This result is also given in Kloock and Getty (2019). Note that in this model, k_2 and Q are necessarily negative, reflecting that the aggressive mimic is undesirable in an absolute sense. In this way it resembles our model of repeated search under mortality risk when this was applied to aggressive mimicry. We illustrate changes in p in the example in figure 3g. Parameters are $V_{CA} = 1$, $A = 1$, $z = 0.30$, $m = 0.10$, $q = 0.5$, and $n = 1$. Other results are given in appendix E and summarized in table 2.

Discussion

We have contrasted the predictions of the standard single-encounter formulation with seven models that involve repeated encounters, all of which can be considered variants of the same general SDT model (which we have called the REP formulation). The REP models were originally developed to study prey choice, aggressive mimicry, mate search, and the aiding of kin, and we have included one novel application (repeated search under mortality risk). The REP models feature a richer set of predictions for optimal discrimination than the single-encounter formulation does. We have shown that with one exception (Kloock and Getty 2019), optimal receivers can in the REP models respond to the increased probability of an undesirable or dangerous event not only by being more

cautious but also (counterintuitively) by being less cautious, depending on conditions. Moreover, they can under some conditions respond to an increased payoff from accepting desirable signalers by reducing their acceptance rates. These two counterintuitive findings have been reported previously for two REP models (application 4 and a model formally similar to application 1) featured in Reeve (1989). The first finding was also reported in recent signal detection models on the building up of energy reserves under predation risk (Trimmer et al. 2017a; McNamara and Trimmer 2019). Here, we show that these counterintuitive outcomes not only are very common in models of repeated encounters but are often driven by the same simple and intuitive mechanism, which we describe below. In addition, we have shown that an increase in the absolute abundance of desirable events can either decrease or increase acceptance rates.

The relationship between acceptance rates (hit rate and false alarm rate) and the fraction of signalers that are desirable (p) takes two nontrivial forms in the REP models, depending on whether undesirable events are undesirable in the relative or the absolute sense. First, a signaler or event that is associated with harm will in many situations be undesirable, but if one were to change conditions by reducing the availability or profitability of the better alternatives, it may now become desirable, warranting a different response. Examples may include prey of low profitability or mates of low quality. Such relative undesirability can lead to a U-shaped relationship between the acceptance rates and p: if all events are desirable ($p = 1$), there is no scope for discrimination and the receiver should always accept. However, if some events are less desirable than others ($0 < p < 1$), receivers will often benefit from exercising discrimination. Finally, if all events are undesirable (i.e., $p = 0$), there is again no scope for discrimination but the receiver may accept given that there are no better alternatives. The result is a U-shaped relationship between acceptance rates and p and consequently a range of p in which the receiver responds to the increased probability of undesirable events by being less cautious and accepting more often. The effect is more pronounced for undesirable events (false alarm rate) than for desirable ones (hit rate; fig. 3). Such U-shaped relationships were shown in Reeve (1989, figs. 5c and 6b). They were also featured in the models of Trimmer et al. (2017a) and McNamara and Trimmer (2019), although driven by a slightly different mechanism (see below). All else being equal, the U-shaped dip in the false alarm rate tends to be deeper when discriminability d' is high (fig. 4). The second nontrivial relationship of acceptance rates with p arises when undesirable events are always undesirable, such as encounters with predators or deadly toxic prey that provide no possible benefit. Such absolute undesirability causes acceptance rates to

increase with p in the REP models, in qualitative agreement with the single-encounter formulation (figs. 3, 4; see also fig. D3).

We welcome the novel state-dependent detection theory (SDDT) as a useful addition to the toolbox of evolutionary ecology (Trimmer et al. 2017a, 2017b). It extends SDT to optimality problems that require an explicit consideration of state, as has been showcased for the building up of energy reserves under risk of predation (Trimmer et al. 2017a, 2017b). The potential areas of application of SDDT are many. Yet its heralded signature counterintuitive result that receivers may take more risk when the proportion of dangerous events increases (Trimmer et al. 2017a) is found in simpler, previously published signal detection models that focus on repeated encounters, as also pointed out by McNamara and Trimmer (2019). Note, however, that the conditions under which the counterintuitive prediction applies differ between the REP applications considered here and the energy reserve models: in particular, Trimmer et al. (2017a) and McNamara and Trimmer (2019) showed that the acceptance rates of foragers ("boldness") can increase with the proportion of dangerous foraging opportunities even when the latter is undesirable in the absolute sense (i.e., when false alarm leads to certain death). This is possible because every rejection in these models reduces reserves by one level and thus increases the future number of correct acceptances needed to reach the target reserve level (McNamara and Trimmer 2019). Increasing reserves by two levels requires two consecutive correct acceptances, which becomes exceedingly unlikely as safe foraging opportunities become infrequent, making accepting immediately to try to gain one reserve level a better option. In the REP applications, in contrast, acceptance rates may increase with the proportion of undesirable events only if these are undesirable in the relative sense.

Intriguingly, McNamara and Trimmer (2019) also emphasized that the false alarm rate in some cases could increase monotonically with the proportion of dangerous events, a prediction at odds both with the single-encounter formulation and with each of the seven REP models that we have explored here. This paradoxical pattern implies that receivers must forego some opportunities to accept even when 100% of events are desirable. Indeed, when both background mortality and predator presence are set to zero in their model, individual prey can live forever, and it does not matter how many time steps are spent raising reserves as long as it happens often enough to ultimately raise them. Introducing an infinitesimally small background mortality forces a U-shaped relationship between danger and acceptance rates. It is therefore unclear whether the predicted monotonic relationship can fit biological reality.

In most REP models, the conditions under which increasing the payoff from accepting desirable signalers may reduce receiver acceptance rates (i.e., $k_2 > 0$) coincide fully with relative undesirability (i.e., $Q > 0$). This makes perfect intuitive sense: when the undesirability of event B is dependent on event A being a better option, then increasing the desirability of A will make B even less desirable, while reducing the desirability of event A down to that of event B will make them both acceptable. In application 3 (Holen 2013), however, we are guaranteed only that k_2 and Q have the same sign when p is sufficiently low. Application 3 differs from the others in that measures of abundance (i.e., p or n_A and n_B) occur in the parameter f, and a full in-depth analysis is beyond the scope of the current article.

For the normal-normal equal variance ROC, changes in discriminability d' have similar effects on acceptance rates in the REP formulation and the single-encounter formulation in the following qualitative sense: if the always-accept strategy is optimal when $d' = 0$, the false alarm rate will strictly decrease as d' increases, while the hit rate will first decrease and then increase (figs. 1b, 2b). Likewise, if the always-reject strategy is optimal when $d' = 0$, the hit rate will strictly increase as d' increases, while the false alarm rate will first increase and then decrease (figs. 1a, 2c). This similarity between the REP formulation and the single-encounter formulation has been pointed out by Kloock and Getty (2019) for their aggressive mimicry model, but it holds for every REP model. This is reassuring and suggests a certain robustness to predictions made about mimetic accuracy and the direction of mimicry evolution (Johnstone 2002; Holen and Johnstone 2004, 2006, 2018) that rely on the qualitative relationship between acceptance rates and discriminability.

One of the most common applications of the single-encounter SDT formulation in evolutionary ecology is prey choice, in particular in the context of Batesian mimicry (e.g., Oaten et al. 1975; Johnstone 2002; Sherratt 2002; Holen and Johnstone 2004, 2006, 2018; Reinhold and Engqvist 2004; Lynn 2005; McGuire et al. 2006; Speed and Ruxton 2010; Polnaszek et al. 2017). Classical mimicry theory focuses on visual warning signals and decisions to attack. The single-encounter model predicts that defended model prey and their mimics should be attacked less often on encounter when the benefit of attacking mimics is low, when the cost of attacking models is high, when the fraction of mimics (p) is low, and (if parameterizing the model for absolute abundances) when mimics are rare and models abundant. These results are consistent with classical mimicry theory (e.g., Sheppard 1959a; Turner 1987; Huheey 1988; Ruxton et al. 2018), with experiments designed to test theory (e.g., Brower 1960; O'Donald and Pilecki 1970; Pilecki and O'Donald 1971;

Nonacs 1985; Lindström et al. 1997) and with the predictions of REP application 1 and 2 when model prey are undesirable in an absolute sense (see Getty 1985).

Applications 1 and 2 in addition allow the undesirability of model prey to be relative and dependent on the abundance of more profitable alternatives. Prey may have low but positive profitability if attacks rarely lead to capture (leading to evasive mimicry; Ruxton et al. 2004) or if postcapture handling time is very high. Even toxic prey could be undesirable in the relative sense if the alternative for the predator is starvation (e.g., Sheppard 1959b; Speed 1999; Sherratt 2003). Model prey that are undesirable only in the relative sense reverse classic predictions of mimicry theory: the probability of attack can be a U-shaped function of the proportion of prey that are mimics, as explained above, which means that the probability of attack may sometimes increase with the proportion of prey that are models. Moreover, if the absolute abundance of mimics is below some minimum threshold mimics and models will always be attacked, but above the threshold the probability of attack will decrease with (absolute) mimic abundance, other parameters kept constant. Likewise, the probability of attack will decrease if mimics become more profitable. The disk equation models therefore suggest that a form of Batesian mimicry is possible in which predators are less willing to attack models when mimics are more abundant or profitable. Yet the predicted directions of mimicry evolution are likely to remain as predicted by the single-encounter formulation. To understand the fate of a mutant signaler, we must consider the range of accepted stimuli, and for any ROC generated by the use of a single acceptance threshold (see fig. A1), mimics will be selected to resemble models more closely while models will be selected to become more dissimilar from the mimics.

Two of the REP models consider sequential search in which the receiver makes a final choice and settles with it. Application 4 (Reeve 1989; Getty 1996b; Shizuka and Hudson 2020) assumes additive search costs, while application 5 (Wiley 1994) assumes multiplicative search costs. Both models predict that acceptance rates increase with search costs, consistent with models of mate search (Real 1990; Wiegmann et al. 1996). The additive model has been applied to mate search in female pied flycatchers (*Ficedula hypoleuca*; Getty 1996b). These females choose between unmated males on a primary territory and mated males that have already paired up but defend a second territory; the latter will provide less parental investment if chosen (Alatalo and Lundberg 1990; Slagsvold and Lifjeld 1994). The deception hypothesis holds that females that mate with already-mated males are deceived and would have benefitted by not initiating mating (Alatalo et al. 1981), while the main competing hypothesis holds that such females strategically make the best of a bad situation (Meier

1983; Dale et al. 1992; Slagsvold and Dale 1994). The latter hypothesis was bolstered by the observation that male mating status could be inferred from territorial presence (Dale and Slagsvold 1994). Naturally, there is no logical contradiction between males sometimes being able to deceive females regarding their mating status and females being able to discriminate somewhat between already-mated and unmated males; deception may after all only work sometimes. In the perspective of SDT, being deceived can partly be seen as an outcome of receiver strategy (see "adaptive gullibility"; Wiley 1994). A more useful distinction between the two hypotheses is therefore that already-mated males are undesirable in an absolute sense under the deception hypothesis but undesirable in a relative sense under the best-of-a-bad-situation hypothesis. The multiplicative search cost model (Wiley 1994) is flexible enough to represent both absolute and relative undesirability and thus both hypotheses. In contrast, the additive cost model (Reeve 1989; Getty 1996b; Shizuka and Hudson 2020) cannot represent absolute undesirability. This is an artifact arising from the fact that the strategy of rejecting all males carries infinite cost, which forces females to settle with choices that yield negative payoff.

Some of the applications presented here pertain to aggressive mimics, which take advantage of a receiver by eliciting behavior in the receiver that is usually directed toward a model organism. Many orchid species, for instance, attract pollinators by mimicking the floral signals of food-rewarding plants (i.e., food deception) but provide no benefit to the pollinators (reviewed in Schiestl 2005). The disk equation with alternative prey (application 2) serves as a reasonable model of how receivers should optimally respond to food deception, with the desirable flowers now being the models and the undesirable, nonrewarding flowers the mimics. Nonrewarding flowers are less rewarding than any diet of alternative food sources and are therefore likely to be undesirable in an absolute sense. Kloock and Getty's (2019) model (application 7) focuses on a specific form of aggressive mimicry in which a predator poses as a food item and the undesirable event (the aggressive mimic) provides no benefit and is always undesirable. The model of repeated search under mortality risk (application 6) also serves as a reasonable model of aggressive mimicry when undesirability is absolute (i.e., $V_B < 0$). In applications 2, 6, and 7, the absolute undesirability implies $k_2 < 0$, and the acceptance rates of a predator that discriminates should increase with the frequency of the desirable models, increase with the payoff from accepting desirable models and the payoff from accepting undesirable mimics, and (in application 2) increase with the absolute abundance of the desirable models and decrease with the absolute abundance of undesirable mimics. Thus, the predictions of the REP applications that are relevant for

aggressive mimicry are consistent with those of the single-encounter formulation when applied to aggressive mimicry (e.g., Holen 2004, 2006).

In conclusion, we have shown that a number of superficially different signal detection models involving sequential choice can be analyzed and understood as specific cases of the same general framework (the REP formulation). The seven applications we have investigated are not exhaustive, and we hope that the solution techniques presented here will make it easier to analyze and compare future signal detection models involving repeated encounters that fall into the same general framework. We have shown how (and why) these models exhibit many properties that single-encounter models do not, properties that were not always recognized at the time of publication. The counterintuitive nature of the predictions of these more complex models makes them excellent candidates for experimental testing. We agree with recent critics (Trimmer et al. 2017a; McNamara and Trimmer 2019) that it often is necessary to explicitly incorporate sequential encounters into models, but we point out that there are already a number of published signal detection models of this kind. SDDT allows SDT to be applied to problems that require a state-dependent approach but does not invalidate the use of SDT in other settings. Our comparison of REP models to the single-encounter formulation revealed well-defined circumstances under which they make similar qualitative predictions. This strongly suggests that the single-encounter formulation remains a reasonable approach for many scenarios, including classic Batesian mimicry and aggressive mimicry. That different signal detection models make different predictions depending on their economic assumptions is not a reason for concern but a testament to the versatility of SDT.

Acknowledgments

Statement of Authorship

Ø.H.H. and T.N.S. conceived and planned the article together. Ø.H.H. did the model analysis. Both authors contributed to verifying and interpreting results and to writing and critically revising the manuscript. Both authors gave their final approval for publication.

Literature Cited

Abbott, K. R. 2010. Background evolution in camouflage systems: a predator-prey/pollinator-flower game. Journal of Theoretical Biology 262:662–678.

Abbott, K. R., and T. N. Sherratt. 2013. Optimal sampling and signal detection: unifying models of attention and speed-accuracy trade-offs. Behavioral Ecology 24:605–616.

Alatalo, R. V., A. Carlson, A. Lundberg, and S. Ulfstrand. 1981. The conflict between male polygamy and female monogamy: the case of the pied flycatcher *Ficedula hypoleuca*. American Naturalist 117:738–753.

Alatalo, R. V., and A. Lundberg. 1990. Polyterritorial polygyny in the pied flycatcher. Pages 1–27 *in* P. J. B. Slater, J. S. Rosenblatt, and C. Beer, eds. Advances in the study of behavior. Academic Press, San Diego, CA.

Blough, D. S. 2001. Some contributions of signal detection theory to the analysis of stimulus control in animals. Behavioural Processes 54:127–136.

Bogaardt, L., and R. A. Johnstone. 2016. Amplifiers and the origin of animal signals. Proceedings of the Royal Society B 283:20160324.

Brower, J. V. Z. 1960. Experimental studies of mimicry. IV. The reactions of starlings to different proportions of models and mimics. American Naturalist 94:271–282.

Dale, S., H. Rinden, and T. Slagsvold. 1992. Competition for a mate restricts mate search of female pied flycatchers. Behavioral Ecology and Sociobiology 30:165–176.

Dale, S., and T. Slagsvold. 1994. Polygyny and deception in the pied flycatcher: can females determine male mating status? Animal Behaviour 48:1207–1217.

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 B 263:925–931.

Dor, R., H. Kedar, D. W. Winkler, and A. Lotem. 2007. Begging in the absence of parents: a "quick on the trigger" strategy to minimize costly misses. Behavioral Ecology 18:97–102.

Egan, J. P. 1975. Signal detection theory and ROC analysis. Academic Press Series in Cognition and Perception. Academic Press, New York.

Ehlman, S. M., P. C. Trimmer, and A. Sih. 2019. Prey responses to exotic predators: effects of old risks and new cues. American Naturalist 193:575–587.

Getty, T. 1985. Discriminability and the sigmoid functional response: how optimal foragers could stabilize model-mimic complexes. American Naturalist 125:239–256.

———. 1987. Crypsis, mimicry, and switching: the basic similarity of superficially different analyses. American Naturalist 130:793–797.

———. 1995. Search, discrimination, and selection: mate choice by pied flycatchers. American Naturalist 145:146–154.

———. 1996a. The maintenance of phenotypic plasticity as a signal detection problem. American Naturalist 148:378–385.

———. 1996b. Mate selection by repeated inspection: more on pied flycatchers. Animal Behaviour 51:739–745.

———. 2002. The discriminating babbler meets the optimal diet hawk. Animal Behaviour 63:397–402.

Green, D. M., and J. A. Swets. 1988. Signal detection theory and psychophysics. Peninsula, Los Altos, CA.

Holen, Ø. H. 2013. Disentangling taste and toxicity in aposematic prey. Proceedings of the Royal Society B 280:20122588.

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.

———. 2018. Reciprocal mimicry: kin selection can drive defended prey to resemble their Batesian mimics. Proceedings of the Royal Society B 285:20181149.

Holen, Ø. H., and T. O. Svennungsen. 2012. Aposematism and the handicap principle. American Naturalist 180:629–641.

Holling, C. S. 1959. Some characteristics of simple types of predation and parasitism. Canadian Entomologist 91:385–398.

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

Johnstone, R. A. 1998a. Conspiratorial whispers and conspicuous displays: games of signal detection. Evolution 52:1554–1563.

———. 1998b. Efficacy and honesty in communication between relatives. American Naturalist 152:45–58.

———. 2002. The evolution of inaccurate mimics. Nature 418: 524–526.

Kloock, C. T., and T. Getty. 2019. A mathematical model of aggressive mimicry. Behavioral Ecology 30:134–141.

Laubach, Z. M., W. Perng, D. C. Dolinoy, C. D. Faulk, K. E. Holekamp, and T. Getty. 2018. Epigenetics and the maintenance of developmental plasticity: extending the signalling theory framework. Biological Reviews 93:1323–1338.

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 264: 149–153.

Lynn, S. K. 2005. Learning to avoid aposematic prey. Animal Behaviour 70:1221–1226.

Lynn, S. K., and L. F. Barrett. 2014. "Utilizing" signal detection theory. Psychological Science 25:1663–1673.

Macmillan, N. A., and C. D. Creelman. 2005. Detection theory: a user's guide. Lawrence Erlbaum, New York.

McGuire, L., H. Van Gossum, K. Beirinckx, and T. N. Sherratt. 2006. An empirical test of signal detection theory as it applies to Batesian mimicry. Behavioural Processes 73:299–307.

McNamara, J. M., and P. C. Trimmer. 2019. Sequential choices using signal detection theory can reverse classical predictions. Behavioral Ecology 30:16–19.

Meier, P. T. 1983. Polyterritorial polygyny in the pied flycatcher: male deception or female choice? American Naturalist 121: 145–147.

Metcalf, C. J. E., A. T. Tate, and A. L. Graham. 2017. Demographically framing trade-offs between sensitivity and specificity illuminates selection on immunity. Nature Ecology and Evolution 1:1766–1772.

Nonacs, P. 1985. Foraging in a dynamic mimicry complex. American Naturalist 126:165–180.

O'Donald, P., and C. Pilecki. 1970. Polymorphic mimicry and natural selection. Evolution 24:395–401.

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.

Pilecki, C., and P. O'Donald. 1971. The effects of predation on artificial mimetic polymorphisms with perfect and imperfect mimics at varying frequencies. Evolution 25:365–370.

Polnaszek, T. J., T. L. Rubi, and D. W. Stephens. 2017. When it's good to signal badness: using objective measures of discriminability to test the value of being distinctive. Animal Behaviour 129:113–125.

Real, L. 1990. Search theory and mate choice. I. Models of single-sex discrimination. American Naturalist 136:376–405.

Reeve, H. K. 1989. The evolution of conspecific acceptance thresholds. American Naturalist 133:407–435.

Reinhold, K., and L. Engqvist. 2004. Inaccurate mimicry and predator ecology. Journal of Theoretical Biology 229:433–434.

Ruxton, G. D., W. L. Allen, T. N. Sherratt, and M. P. Speed. 2018. Avoiding attack: the evolutionary ecology of crypsis, aposematism, and mimicry. Oxford University Press, Oxford.

Ruxton, G. D., M. Speed, and T. N. Sherratt. 2004. Evasive mimicry: when (if ever) could mimicry based on difficulty of capture evolve? Proceedings of the Royal Society B 271:2135–2142.

Schiestl, F. P. 2005. On the success of a swindle: pollination by deception in orchids. Naturwissenschaften 92:255–264.

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

———. 1959b. Natural selection and heredity. Hutchinson University Library, London.

Sherratt, T. N. 2001. The evolution of female-limited polymorphisms in damselflies: a signal detection model. Ecology Letters 4:22–29.

———. 2002. The evolution of imperfect mimicry. Behavioral Ecology 13:821–826.

———. 2003. State-dependent risk-taking by predators in systems with defended prey. Oikos 103:93–100.

Sherratt, T. N., and Ø. H. Holen. 2018. When should receivers follow multiple signal components? a closer look at the "flag" model. Behavioral Ecology 29:e6–e8.

Shizuka, D., and E. J. Hudson. 2020. To accept or reject heterospecific mates: behavioural decisions underlying premating isolation. Philosophical Transactions of the Royal Society B 375:20190484.

Slagsvold, T., and S. Dale. 1994. Why do female pied flycatchers mate with already mated males: deception or restricted mate sampling? Behavioral Ecology and Sociobiology 34:239–250.

Slagsvold, T., and J. T. Lifjeld. 1994. Polygyny in birds: the role of competition between females for male parental care. American Naturalist 143:59–94.

Speed, M. P. 1999. Batesian, quasi-Batesian or Müllerian mimicry? theory and data in mimicry research. Evolutionary Ecology 13:755–776.

Speed, M. P., and G. D. Ruxton. 2010. Imperfect Batesian mimicry and the conspicuousness costs of mimetic resemblance. American Naturalist 176:E1–E14.

Sperling, G. 1984. A unified theory of attention and signal detection. Pages 103–181 in R. Parasuraman and D. R. Davies, eds. Varieties of attention. Academic Press, New York.

Staddon, J. E. R., and R. P. Gendron. 1983. Optimal detection of cryptic prey may lead to predator switching. American Naturalist 122:843–848.

Stephens, D. W., and J. R. Krebs. 1986. Foraging theory. Monographs in Behavior and Ecology. Princeton University Press, Princeton, NJ.

Sumner, C. J., and S. Sumner. 2020. Signal detection: applying analysis methods from psychology to animal behaviour. Philosophical Transactions of the Royal Society B 375:20190480.

Treisman, M. 1975. Predation and the evolution of gregariousness. II. An economic model for predator-prey interaction. Animal Behaviour 23:801–825.

Trimmer, P. C., S. M. Ehlman, J. M. McNamara, and A. Sih. 2017a. The erroneous signals of detection theory. Proceedings of the Royal Society B 284:20171852.

Trimmer, P. C., S. M. Ehlman, and A. Sih. 2017b. Predicting behavioural responses to novel organisms: state-dependent detection theory. Proceedings of the Royal Society B 284:20162108.

Trimmer, P. C., A. I. Houston, J. A. R. Marshall, R. Bogacz, E. S. Paul, M. T. Mendl, and J. M. McNamara. 2008. Mammalian choices: combining fast-but-inaccurate and slow-but-accurate decision-making systems. Proceedings of the Royal Society B 275:2353–2361.

Turner, J. R. G. 1987. The evolutionary dynamics of Batesian and Muellerian mimicry: similarities and differences. Ecological Entomology 12:81–95.

Wiegmann, D. D., L. A. Real, T. A. Capone, and S. Ellner. 1996. Some distinguishing features of models of search behavior and mate choice. American Naturalist 147:188–204.

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.

Associate Editor: Jürgen Groeneveld
Editor: Daniel I. Bolnick

The Eurasian bee beetle, *Trichius fasciatus*, a Batesian mimic of bumble bees. Photo © Øistein Haugsten Holen.