

PEAK SHIFT DISCRIMINATION LEARNING AS A MECHANISM OF SIGNAL EVOLUTION

SPENCER K. LYNN,^{1,2} JONATHAN CNAANI,^{1,3} AND DANIEL R. PAPAJ^{1,4}
¹*Ecology and Evolutionary Biology, University of Arizona, Tucson, Arizona 85721*
⁴*E-mail: papaj@email.arizona.edu*

Abstract.—“Peak shift” is a behavioral response bias arising from discrimination learning in which animals display a directional, but limited, preference for or avoidance of unusual stimuli. Its hypothesized evolutionary relevance has been primarily in the realm of aposematic coloration and limited sexual dimorphism. Here, we develop a novel functional approach to peak shift, based on signal detection theory, which characterizes the response bias as arising from uncertainty about stimulus appearance, frequency, and quality. This approach allows the influence of peak shift to be generalized to the evolution of signals in a variety of domains and sensory modalities. The approach is illustrated with a bumblebee (*Bombus impatiens*) discrimination learning experiment. Bees exhibited peak shift while foraging in an artificial Batesian mimicry system. Changes in flower abundance, color distribution, and visitation reward induced bees to preferentially visit novel flower colors that reduced the risk of flower-type misidentification. Under conditions of signal uncertainty, peak shift results in visitation to rarer, but more easily distinguished, morphological variants of rewarding species in preference to their average morphology. Peak shift is a common and taxonomically widespread phenomenon. This example of the possible role of peak shift in signal evolution can be generalized to other systems in which a signal receiver learns to make choices in situations in which signal variation is linked to the sender’s reproductive success.

Key words.—Batesian mimicry, decision-making, learning, peak shift, sexual selection, signal detection theory, signal evolution.

Received May 2, 2004. Accepted March 3, 2005.

An appreciation for the role of cognitive phenomena as part of the mechanism of natural selection has been building for several years (e.g., receiver psychology [Guilford and Dawkins 1991], cognitive ecology [Dukas 1998]). One topic that has surfaced repeatedly at this interface between comparative psychology and behavioral ecology is the discrimination learning phenomenon “peak shift.” Peak shift is a common finding in discrimination learning experiments (Rilling 1977): Control subjects are trained to respond (e.g., button press) to a positively reinforced stimulus (S+, e.g., a line of a particular orientation or a light of a particular wavelength). Treatment subjects are trained identically to control subjects with respect to S+ and are also trained to withhold response to an unreinforced stimulus (S–, e.g., a line orientation or hue similar but not identical to S+). Both groups of subjects are then tested without reinforcement on a continuum of similar stimuli. During the test, control subjects respond most strongly to the S+ stimulus. However, treatment subjects respond most strongly to a novel stimulus in preference to the one they learned was rewarding during training.

The Peak Shift Discrimination Phenomenon

The phenomenon that treatment subjects display in this type of experiment is known as “peak shift” (Hanson 1959, and reviewed in animal learning textbooks, e.g., Domjan 1998; Shettleworth 1998). A plot of the subjects’ strength of response as a function of stimulus value shows a bell-shaped

response gradient. The stimulus receiving the maximum (“peak”) response by the treatment subjects is said to be “shifted” relative to that of the control subjects. The shift is in a direction away from S–, but is limited in its extent; extreme stimuli do not receive increased response. There is also a shift of the most strongly avoided stimulus off of S–, away from S+ (Guttman 1965). The seemingly paradoxical preference for an unrewarded and novel stimulus over the stimulus that subjects have learned is rewarding during training can be accounted for mechanistically by additive interaction of bell-shaped S+ and S– generalization gradients (gradients of the strength of the learned association between stimulus morphology and response reinforcement) that overlap each other across the range of test stimuli (reviewed by Cheng 2002).

The response bias inherent to peak shift has lead several researchers to hypothesize that it could be involved in the evolution of signaling systems. Much thought has been applied to the potential role of peak shift in the sexual selection of sex or species recognition characters (Guilford and Dawkins 1991; Weary et al. 1993; Enquist and Arak 1998) and in the evolution of warning coloration (Leimar et al. 1986; Gamberale and Tullberg 1996; Enquist and Arak 1998; Yachi and Higashi 1998; Lindström et al. 1999). In the context of the evolution of warning coloration Leimar et al. (1986) and Yachi and Higashi (1998) developed mathematical models of learning. Their results indicated that for the evolution of warning coloration, peak shift could mitigate hypothesized requirements for kin selection, nonlethal attacks, prey aggregations, and large mutational changes imposed by theories that did not take features of discrimination learning into account. Leimar and Tuomi (1998) applied peak shift toward the understanding of plants’ anti-herbivore defenses. Weary et al. (1993) recognized that peak shift might

² Present address: Harvard Medical School, Department of Psychiatry, Center for Mental Health and Media, 271 Waverley Oaks Road, Waltham, Massachusetts 02452-8403; E-mail: spencer@nmr.mgh.harvard.edu.

³ Present address: Arava Institute for Environmental Studies, Kibbutz Ketura, D.N. Eilat, 88840, Israel. E-mail: jcnaani@arava.co.il.

have effects on within-species discriminations other than sexual selection, and suggested discriminations of adult versus adolescent and low versus high quality individuals as possible domains in which to search for an influence.

Here, we present an experiment within a novel functional account of peak shift based on signal detection theory. Modeling learned behavior at a functional level in this manner highlights the commonalities across the diverse sensory and behavioral domains within which generalization and discrimination have fundamental roles (Ghirlanda and Enquist 2003). As a model of cognition, signal detection theory provides a powerful comparative framework within which to understand cognitive behavior as both agent and object of natural selection.

The Signals Approach to Discrimination Learning

Signal detection theory (SDT, Green and Swets 1966; Wiley 1994) is a mathematical description of the trade-offs and uncertainty inherent in the reception of signals (i.e., discerning one signal from another or signal from noise). The theory specifies three parameters that every signal has and that, optimally, any receiver should take into account when responding to a signal: (1) a signal's frequency distribution over the sensory domain, (2) a signal's abundance relative to alternative signals, and (3) the payoffs for responding to or ignoring a signal. From the perspective of behavioral ecology, resources (e.g., food, mates) can be considered as sensory stimuli that emit signals (e.g., their appearance). As a functional model of decision-making, signal detection theory thus links choice behavior to estimates of ecologically valid variables upon which to base those choices: the distribution, abundance, and quality of resources.

In behavioral ecology, the uncertainty modeled by SDT is typically considered to arise from perceived similarity in the appearance of stimuli (Wiley 1994). Perceptual variability in prey appearance might mirror phenotypic variability in the prey or arise from sensory noise (Boneau and Cole 1967). Additionally, the uncertainty modeled by SDT may arise from stimulus generalization (Blough 1967, 1969; Lynn 2005), a process dependant on reinforcement history in addition to perception. Reciprocally, the account of peak shift developed here suggests that the phenomenon, while traditionally assumed to arise from learning, might also arise from conventional perceptual uncertainty.

SDT uses information about stimulus encounters to place a behavioral response criterion (a threshold) on the stimulus domain; stimuli on one side of the threshold receive a response, stimuli on the other side are ignored. Optimal criterion placement maximizes the number of correct detections of S+ (signals to which response is relatively advantageous) and correct rejections of S- (signals to which response is relatively disadvantageous) while minimizing missed detections of S+ and false-alarm responses to S-. The influence of the three signal parameters on optimal behavioral response is described by a utility function (Sperling 1984; Wiley 1994):

$$U(x) = \alpha h P[CD] + \alpha m P[MD] + (1 - \alpha) a P[FA] + (1 - \alpha) j P[CR] \quad (1)$$

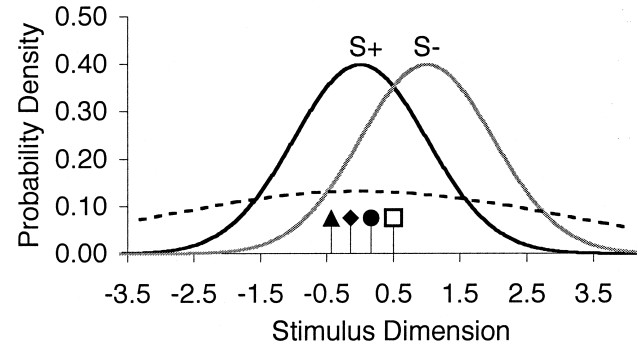


FIG. 1. The influence of signal parameters on the behavioral response criterion. Bell-shaped signal distribution parameters are attributed to the S+ and S- training stimuli, such that novel stimuli are assigned to the S+ or S- signal class with a likelihood based on perceptual similarity to the respective training stimuli. Particular combinations of the three signal parameters (Eq. 1) locate different optimal decision criteria on the stimulus domain. Four examples are shown: (□) arbitrary "baseline" signal parameters, signal distributions modeled as Gaussian probability density functions with variance = 1, means = 0 (S+, solid dark curve) and 1 (S-, solid light curve), $\alpha = 0.50$, $h = 1.00$, $m = -0.20$, $a = -1.00$, $j = 0.20$; (●) a parameter set that differs from baseline in the payoff for correct detection, $h = 0.66$; (◆) a set differing from baseline in the variance of the S+ signal distribution, $\sigma^2 = 3$ (dashed curve, tails truncated for display purposes); and (▲) a set differing from baseline in the relative abundance of the signal types, $\alpha = 0.28$. Relative to the criterion location for the baseline parameters, the other parameter sets are shifted away from S- and result in a lower probability of false alarm responses (integral of the S- distribution from criterion to $-\infty$).

where $U(x)$ = estimated utility over stimulus domain x ; $P[CD]$ = probability of correct detection (measured as the integral of the S+ distribution from threshold to $-\infty$); $P[MD]$ = probability of missed detection (equal to $1 - P[CD]$); $P[FA]$ = probability of false alarm (integral of the S- distribution from threshold to $-\infty$); $P[CR]$ = probability of correct rejection ($= 1 - P[FA]$), $\alpha = \text{alpha}$, the relative probability of encountering an S+ signal ($1 - \alpha$ = the relative probability of encountering a signal from the S- distribution); h = payoff of correct detection; m = payoff of missed detection; a = payoff of false alarm; and j = payoff of correct rejection (where $h > m$ and $j > a$).

The point of maximum utility (the optimal criterion location) shifts as signal parameters (or an animal's functional representations of them) change. For example, relative to arbitrary "baseline" parameters, increasing the relative abundance of S- stimuli, increasing the variance of S+ stimuli, or decreasing the benefit for correct detection of S+ stimuli each increase the risk of mistaking S+ for S- and cause the point of maximum utility to shift away from S-, in the direction of decreasing probability of false alarm (Fig. 1).

With this background, the three signal parameters can each be seen to correspond to elements of a discrimination learning experiment. (1) The signals are the appearance of the S+ and S- stimuli. The signal distributions correspond to gradients of relative likelihood, based on perceptual similarity, that a particular stimulus is from the S+ or S- stimulus class. (2) The relative abundance of S+ and S- signals corresponds to the relative frequencies of stimulus presentation during

training. (3) The payoffs correspond to the reinforcement for responding to or ignoring S+ and S- stimulus presentations.

Under the signals approach, the peak shift experiment is a signal discrimination task in which animals are uncertain as to which response (approach or avoid) is appropriate for any given test stimulus. This perspective implies that the response "peak," although not a threshold, is a type of decision criterion. The magnitude of the peak shift displacement should thus be sensitive to variations in the three signal parameters. Manipulation of training contingencies similar to the different parameter sets of Figure 1 should produce different magnitudes of peak shift. To investigate the possible influence of peak shift on signal evolution in a more naturalistic setting than that of typical peak shift experiments, we trained bumblebees to forage for sugar water among a patch of artificial flowers. Flowers within the patch were of two types, rewarding (S+) and unrewarding (S-), which were similar in color, modeling a Batesian mimicry system.

METHODS

The experiment took place in a flight cage, 60 cm on each side. During experiments, the lab fluorescent lights were off and the flight cage was illuminated by a 150-watt GE Plant Gro-n-Show lamp, model 150r40/PL. The floor of the cage held a 6 × 6 array of artificial flowers, 5-cm high on a 10-cm grid. Flowers were made from paper disks, 4-cm diameter, colored by an inkjet printer. In HSB color-space, the hue of nine flower colors ranged from 80° (yellow-green) to 160° (green-blue) in 10° increments (saturation = 50%, brightness = 100%). Hue and brightness were not otherwise controlled among the stimuli; however, bees will ignore brightness if hue differences provide enough discrimination (Backhaus 1991; Chittka et al. 1992). Bees were trained on a subset of the colors and tested on all nine colors. Flower positions on the array were randomized. Reinforcer was pipetted onto a thumbtack at the center of each flower prior to a training or testing session and not replenished during the session. During tests, each color was present on the array four times and deionized (DI) water was used as the reinforcer. The number of landings on each of the nine colors was counted for each bee. Counts were converted to relative proportion of total landings per bee and arc-sin transformed. In SPSS (r11; SPSS, Inc.; Chicago, IL) a split-plot ANOVA design (across bee group) was then applied to the transformed data, utilizing landings on the nine colors as within-subjects measures. Group and color were considered fixed effects. The factor of interest was the landings-by-group interaction. Because Mauchly's Test indicated that sphericity assumptions were not met, both the conservative Greenhouse-Geisser and more liberal Huynh-Feldt corrections were computed. Following Stevens (1992) the mean degrees of freedom of both corrections was used to determine the *P*-value.

The subjects were bumblebees, *Bombus impatiens*, obtained from Koppert Biological Systems (Romulus, MI). The experiments involved six groups of bees, 10 inexperienced bees per group. To preclude an effect of hive or hive age on the results, group assignment was proportionally balanced across three hives. One group received training under arbitrary "baseline" signal parameter values. Three comparison

groups received training that differed from baseline by the manipulation of one signal parameter. Additionally, two control groups were utilized. Signal parameter values used during training were: Baseline Group—relative frequency S+ = 0.50 (18 S+ flowers, 18 S- flowers present on training array); S+ stimulus = 120° hue; S+ reinforcement = 4 μL 50% sucrose solution, half of S+ flowers reinforced with sugar water, half with DI water; S- stimulus = 140° hue; S- reinforcement = 4 μL 3% NaCl solution, half of S- flowers reinforced with saltwater, half with DI water. Relative Frequency Comparison Group—relative frequency S+ = 0.28 (10 S+ flowers, 26 S- flowers present on training array), other conditions as for baseline group. Payoff Comparison Group—S+ reinforcement = 4 μL 33% sucrose solution, other conditions as for baseline group. Distribution Comparison Group—S+ stimuli = 110°, 120°, and 130° hue stimuli present simultaneously, each color present six times, other conditions as for baseline group. Naive Control Group—neither S+ nor S- training. Discrimination Control Group—S+ training only, as per baseline group. All bees received pretraining on unprinted (white) flowers prior to training on colored flowers in order to learn to feed from the flowers (4 μL 50% sucrose solution on 18 randomly placed flowers). Additionally, all but the Naive Control Group received further pretraining on colored flowers (4 μL of S+ reinforcement on 18 randomly placed S+ flowers). Without this additional pretraining, many bees would never visit S+ during the first experimental training session, instead landing only on S- (reflected in bias of Naive bees; see Results). The groups that received S- training were tested after reaching a criterion of 80% of landings on S+ flowers during one training session. The control groups were tested after receiving the number of training sessions averaged by the other groups. Training and testing sessions ended when a bee returned to the hive or had not landed on a flower for 3 min.

RESULTS

Bumblebees exhibited peak shift. The flower choices of naive, discrimination control, and baseline bees differed from one another (Table 1). Baseline bees preferred a novel, yellow stimulus over S+ (Fig. 2A). Additionally, as predicted, bees showed sensitivity to signal-borne risk. The payoff and distribution parameter manipulations produced greater shifts in flower choice off of the S+ training stimuli, in a direction away from S-, than that produced by the baseline group's parameters (Fig. 2B; Table 1). The relative abundance manipulation produced a skew in the predicted direction but not a significant shift in flower choice relative to the baseline group's shift (Table 1). This skew is known as "area shift" in the comparative psychology literature (Rilling 1977; Cheng 2002) where it is considered a weak form of peak shift.

Responses of naive bees toward the bluer colors could have been caused by the pretraining that all bees received but are in agreement with prior research on bee innate color preferences (reviewed by Gumbert 2000). The unimodal nature of both control groups' response gradients (vs., e.g., bimodal) indicates that the bees subjectively perceived the stimuli in the rank order of their objective hue. The peak shift exhibited

TABLE 1. ANOVA results of bee flower landings.

Group comparison	df ¹	F	P (alpha = 0.05)
Baseline vs. Control	4.65, 83.72	9.94	$P < 0.0005$
Baseline vs. Naive	5.98, 107.62	13.14	$P < 0.0005$
Control vs. Naive	4.31, 77.60	13.82	$P < 0.0005$
Payoff vs. Baseline	4.36, 78.50	2.56	$0.050 > P > 0.025$
Distribution vs. Baseline	4.70, 79.82	2.53	$0.050 > P > 0.025$
Relative abundance vs. Baseline	4.45, 80.15	0.41	$P > 0.25$

¹ Mean of Greenhouse-Geisser and Huynh-Feldt corrections for lack of sphericity.

by other groups was therefore not an artifact of perceptual idiosyncrasy that might, for example, judge a stimulus of 110° hue as more similar to 130° than to 120°.

DISCUSSION

Peak shift is taxonomically widespread. In addition to the bees of the present study, it is exhibited by rats; birds; fish;

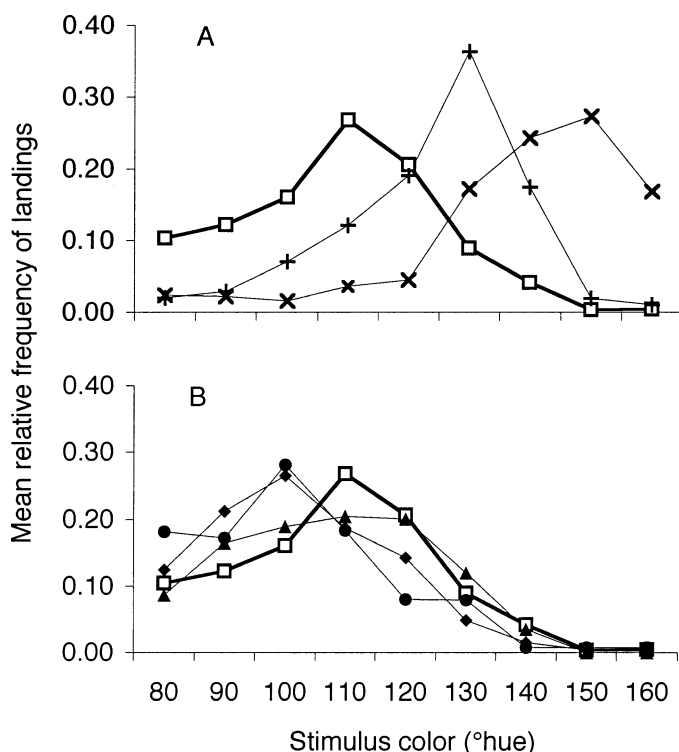


FIG. 2. Bumblebees exhibit peak shift and aversion to signal-borne risk. (A) Responses of naive bees (x) were toward bluer colors (higher stimulus values). Control bees (+), in the absence of discrimination training with S⁻, landed at the 130° stimulus more frequently than at S⁺ (120°) reflecting the underlying predisposition for bluer stimuli revealed by the naive group. The baseline group (□), exposed to S⁻ (at 140°) exhibited a shift in flower choice relative to both the naive and the control groups ($P < 0.0005$), and landed more frequently on a novel and unrewarded stimulus than on S⁺ itself. (B) In accordance with the increased signal-borne risk of their training regimes, payoff (●) and distribution groups (◆) showed a shift in flower preference over and above that of the baseline group (□; $0.050 > P > 0.025$). The relative frequency group (▲) showed a response skewed in the predicted direction, but not significantly different choices from the baseline group. Data are mean proportion of responses per stimulus, $n = 10$ bees per group.

primates, including humans (Purton 1973); and possibly moths (Daly et al. 2001). The signals approach to discrimination learning developed here casts the peak shift phenomenon as a decision-making strategy. This characterization contrasts with those of either the evolutionary or psychological literatures, in which peak shift has been viewed as a heritable trait (e.g., O'Donald 1980; Weary et al. 1993) or as strictly a learning phenomenon (e.g., Domjan 1998; Shettleworth 1998), respectively. The perspective developed here highlights the functional relevance of the cognitive processes of discrimination and generalization as agents of evolutionary selection in a broader and more generalizable way than prior accounts. Though classical signal detection theory (Eq. 1), with its threshold-based response, does not adequately model the biological mechanism producing a bell-shaped "peaked" response gradient, SDT does capture the *functional* ecological parameters of the task, and characterizes the "shift" as due to an aversion to signal-borne risk. As a *mechanistic* model for understanding bell-shaped response gradients, associative learning theory may be used or the "integrated signal distributions" assumption of classical SDT can be modified (Lynn 2003).

Some authors have hypothesized that there might be heritable variation within a population in the magnitude of a peak shift (O'Donald 1980, p. 165; Weary et al. 1993). For example, Weary et al. suggested that peak shift could, in circumstances in which it leads females to select low quality or rare males, be a costly trait and selected against. Alternatively, they posited, in cases of little variation in male character or little difference in benefit among males, peak shift could be neutral with respect to a female's fitness. Although mechanistic aspects of signal parameter estimation are undoubtedly heritable, the present experiments show that the magnitude of the shift is related to the values of each of the three signal parameters under which subjects learn to distinguish S⁺ from S⁻. The amount of shift is thus already more or less optimized to reflect stimulus characteristics such as relative abundance of males of differing quality and the difference in benefit among males. For experienced decision-makers it is conceivable that the costs of searching for rare resources is incorporated into the payoff parameter. As long as the stimulus parameters that subjects have learned (or are provided with as innate predispositions) are accurate at the time a decision is made, a peak shift should always be more beneficial than no shift at all when S⁺ and S⁻ may be mistaken for one another. Except for conditions in which parameters are learned only once (e.g., imprinting [Weary et al. 1993]), payoffs cannot be evaluated, or the decision is only made once, the shift will be changing more quickly as

a plastic learning process than as a fixed behavioral character under selection across generations.

The bee experiments suggest that peak shift could drive the evolution of flower characteristics for rewarding flowers when the rewarding species can be confused with a non-rewarding species. Floral mimicry by orchids is such a system. Many species of orchids offer no nectar reward to visiting pollinators. The orchids are Batesian mimics. Orchid floral mimics rely on the similarity of their appearance to that of nearby rewarding species to produce false alarm visitation mistakes by pollinators. For example, *Orchis boryi* is a Batesian mimic of a purple variety of *Vicia villosa* (Gumbert and Kunze 2001). In peak shift terms, the orchid is an S⁻ stimulus and the *Vicia* S⁺. Gumbert and Kunze have established the occurrence of false alarm visitations by pollinators, so it is known that the signal distributions overlap sufficiently. If the other signal parameters have appropriate values (e.g., false alarms are sufficiently costly in terms of time or energy, or the orchid is sufficiently common) and there is variation in the morphology of the *Vicia* upon which selection can act, then pollinators are predicted to show a peak shift toward a color variant of *V. villosa* that is less similar to the orchid and, by definition, rarer than the habitat's average *villosa*.

Under conditions of uncertainty, flowers with which pollinators have less experience, but that are less likely to be confused with costly flowers, are predicted to be at a pollination advantage. The change in pollination frequencies will have an effect on the host plant evolution (selection for the variant coloration), feeding back to drive the evolution of the signals upon which the decisions are based. In such a circumstance, changes to the signal parameters of unpreferred stimuli induce receivers to change their response to preferred stimuli without any change having occurred in the preferred stimuli themselves. The orchid mimic is affecting the model's evolution via the effect of signal-borne risk on the flower visitation decisions.

Several authors have made arguments for the role of peak shift in the sexual selection of secondary sexual traits. For a female choosing a mate, S⁺ has been considered to be conspecific males, but S⁻ has been assumed to be either other females (Guilford and Dawkins 1991; Weary et al. 1993) or males of different species but similar appearance (Guilford and Dawkins 1991). Peak shift might thus drive evolution of sexual dimorphism and species recognition signals as far as the initial peak shift, but not beyond (Weary et al. 1993). This limitation arises because the S⁻ morphology is not co-evolving with the S⁺ morphology, and as the S⁺ and S⁻ stimuli become more distinguishable the magnitude of the shift diminishes (Purtle 1973; Rilling 1977).

The signals approach, however, indicates how both S⁺ and S⁻ distributions may originate from the same population of potential mates, opening up the diversity of sexually selected, elaborated characters to the influence of generalization and discrimination learning. In a typical peak shift experiment, payoffs for S⁺ or S⁻ are not considered to change as the stimulus value changes, but to be equal for all stimuli within each class. More naturalistically, however, traits evaluated during mate choice may honestly signal better quality mates with increasing stimulus value (Andersson and Iwasa 1996). Consider a hypothetical signal identification issue in

which a female songbird chooses a mate by evaluating the males' rate of singing (e.g., Bradbury and Vehrencamp 1998). Males emit signals (songs) at some rate with some relative frequency and information about health is honestly encoded in song rate. Healthy males tend to sing at a faster rate than unhealthy males. Objectively, of course, there are not two categories of males, healthy and sick. Rather, there is a single distribution of males arrayed on a continuum of health, which may be unimodal. However, if the female's experience is limited or if she has an imprinted or innate expectation that remains separate from a distribution compiled from experience, then she will represent males with two distributions. Because she carries two signal distributions that have different average payoffs she may exhibit a peak shift in the direction of faster singing males.

Peak shift could thus lead females to choose males of increasingly exaggerated traits in a runaway evolutionary process. As a product of generalization and discrimination, this runaway selection occurs even if females have no predisposition or sensory bias for the signal (Ryan 1998) and without a genetic correlation between male trait and female preference (Fisher 1930). In the case of peak shift, females have no "knowledge" that "faster is better;" they will not choose the fastest singing males they encounter. Rather, females will avoid false alarms: males that they perceive to have a high probability of being from the less profitable type. Here, type (sick-and-slow or healthy-and-fast) is an artifact of the female's stimulus generalization, imposed on the actual male distribution.

Signaling and communication provide good systems in which to study the evolutionary feedback between selection by the environment on cognitive characteristics, and the action of cognition as an agent of selection on the environment. Generalization and discrimination learning can be used to test ideas about the perception of signals, the payoffs of decisions based on those signals, and how decisions may come to influence the morphology of signals. The applicability of signal detection theory to phenomena in generalization and discrimination learning, such as peak shift, has interesting implications regarding a role for cognitive mechanisms in the evolution of signaling systems as diverse as mate choice (Weary et al. 1993), warning coloration (Gamberale and Tullberg 1996), anti-herbivore defense (Leimar and Tuomi 1998), diet choice (Getty 1985), and Batesian mimicry and extravagant advertisement (present study). In neither behavioral ecology or comparative psychology is it altogether intuitive that animals should learn a preference for stimuli they have never experienced or that preferences may be altered without the preferred stimulus itself undergoing some change. However, when costly and beneficial signals are perceptually similar, for example, for deception by mimicry, or even when generalizing from past experience over honest signals, there are cognitive mechanisms widespread in the animal kingdom that yield such nonintuitive results. These examples of the possible role of peak shift in signal evolution can be generalized to other domains in which a signal receiver learns to make choices in situations in which signal variation is linked to the sender's reproductive success.

ACKNOWLEDGMENTS

The theoretical outline of peak shift presented here was developed while SKL was supported as a visiting student at The Media Laboratory, Massachusetts Institute of Technology. He thanks I. Pepperberg and N. Gershenfeld for that opportunity. The bee experiments were supported in part by a National Science Foundation grant to DRP. A. J. Figueredo provided statistical advice. Portions of this material were submitted by SKL in partial fulfillment of the requirements for the degree of Doctor of Philosophy at The University of Arizona. We thank I. Pepperberg, A. J. Figueredo, J. King, and W. Maddison, as well as A. Cutter and K. Prudic for insightful discussion.

LITERATURE CITED

- Andersson, M., and Y. Iwasa. 1996. Sexual selection. *Trends Ecol. Evol.* 11:53–58.
- Backhaus, W. 1991. Color opponent coding in the visual system of the honeybee. *Vision Res.* 31:1381–1397.
- Blough, D. S. 1967. Stimulus generalization as signal detection in pigeons. *Science* 158:940–941.
- . 1969. Generalization gradient shape and summation in steady-state tests. *J. Exp. Anal. Behav.* 12:91–104.
- Boneau, C. A., and J. L. Cole. 1967. Decision theory, the pigeon, and the psychophysical function. *Psych. Rev.* 74:123–135.
- Bradbury, J. W., and S. L. Vehrencamp. 1998. Principles of animal communication. Sinauer, Sunderland, MA.
- Cheng, K. 2002. Generalization: mechanistic and functional explanations. *Anim. Cognit.* 5:33–40.
- Chittka, L., W. Beier, H. Hertel, E. Steinmann, and R. Menzel. 1992. Opponent colour coding is a universal strategy to evaluate the photoreceptor inputs in Hymenoptera. *J. Comp. Physiol. A* 170:545–636.
- Daly, K. C., S. Chandra, M. L. Durtschi, and B. H. Smith. 2001. The generalization of an olfactory-based conditioned response reveals unique but overlapping odour representations in the moth *Manduca sexta*. *J. Exp. Biol.* 204:3085–3095.
- Domjan, M. 1998. The principles of learning and behavior, 4th ed. Brooks/Cole, Pacific Grove, CA.
- Dukas, R. 1998. Cognitive ecology: the evolutionary ecology of information processing and decision-making. Univ. of Chicago Press, Chicago, IL.
- Enquist, M., and A. Arak. 1998. Neural representation and the evolution of signal form. Pp. 21–87 in R. Dukas, ed. *Cognitive ecology: the evolutionary ecology of information processing and decision-making*. Univ. of Chicago Press, Chicago, IL.
- Fisher, R. 1930. The genetical theory of natural selection. Clarendon, Oxford, U.K.
- Gamberale, G., and B. S. Tullberg. 1996. Evidence for a peak-shift in predator generalization among aposematic prey. *Proc. R. Soc. Lond. B* 263:1329–1334.
- Getty, T. 1985. Discriminability and the sigmoid functional response: How optimal foragers could stabilize model-mimic complexes. *Am. Nat.* 125:239–256.
- Ghirlanda, S., and M. Enquist. 2003. A century of generalization. *Anim. Behav.* 66:15–36.
- Green, D. M., and J. A. Swets. 1966. Signal detection theory and psychophysics. Wiley, New York.
- Guilford, T., and M. S. Dawkins. 1991. Receiver psychology and the evolution of animal signals. *Anim. Behav.* 42:1–14.
- Gumbert, A. 2000. Color choices by bumble bees (*Bombus terrestris*): innate preferences and generalization after learning. *Behav. Ecol. Sociobiol.* 48:36–43.
- Gumbert, A., and J. Kunze. 2001. Colour similarity to rewarding model plants affects pollination in a food deceptive orchid, *Orchis boryi*. *Biol. J. Linn. Soc.* 72:419–433.
- Guttman, N. 1965. Effects of discrimination formation on generalization measured from a positive-rate baseline. Pp. 210–217 in D. I. Mostofsky, ed. *Stimulus generalization*. Stanford Univ. Press, Palo Alto, CA.
- Hanson, H. M. 1959. Effects of discrimination training on stimulus generalization. *J. Exp. Psychol.* 58:321–334.
- Leimar, O., and J. Tuomi. 1998. Synergistic selection and graded traits. *Evol. Ecol.* 12:59–71.
- Leimar, O., M. Enquist, and B. Sillén-Tullberg. 1986. Evolutionary stability of aposematic coloration and prey unprofitability: a theoretical analysis. *Am. Nat.* 128:469–490.
- Lindström, L. R., V. Alatalo, J. Mappes, M. Riipi, and L. Vertainen. 1999. Can aposematic signals evolve by gradual change? *Nature* 397:249–251.
- Lynn, S. K. 2003. The signals approach to decision-making in behavioral ecology. Ph.D. dissertation, University of Arizona, Tucson, AZ.
- . 2005. Learning to avoid aposematic prey. *Anim. Behav.* *In press*.
- O'Donald, P. 1980. Genetic models of sexual selection. Cambridge Univ. Press, Cambridge, U.K.
- Purtilo, R. B. 1973. Peak shift: a review. *Psych. Bull.* 80:408–421.
- Rilling, M. 1977. Stimulus control and inhibitory processes. Pp. 432–480 in W. K. Honig and J. E. R. Staddon, eds. *Handbook of operant behavior*. Prentice-Hall, Englewood Cliffs, NJ.
- Ryan, M. J. 1998. Sexual selection, receiver biases, and the evolution of sex differences. *Science* 281:1999–2003.
- Shettleworth, S. J. 1998. Cognition, evolution, and behavior. Oxford Univ. Press, New York.
- Sperling, G. 1984. A unified theory of attention and signal detection. Pages 103–181 in R. Parasuraman and R. R. Davies, eds. *Varieties of attention*. Academic Press, Orlando, FL.
- Stevens, J. 1992. Applied multivariate statistics for the social sciences, 2d ed. Lawrence Erlbaum Associates, Hillsdale, NJ.
- Weary, D. M., T. C. Guilford, and R. G. Weisman. 1993. A product of discriminative learning may lead to female preferences for elaborate males. *Evolution* 47:333–336.
- Wiley, R. H. 1994. Errors, exaggeration, and deception in animal communication. Pp. 157–189 in L. A. Real, ed. *Behavioral mechanisms in evolutionary ecology*. Univ. of Chicago Press, Chicago, IL.
- Yachi, S., and M. Higashi. 1998. The evolution of warning signals. *Nature* 394:882–884.

Corresponding Editor: M. Björklund