

Decision-Making and Learning: The Peak Shift Behavioral Response

S. K. Lynn, Boston College, Chestnut Hill, MA, USA

© 2010 Elsevier Ltd. All rights reserved.

Peak Shift Is a Directional Behavioral Bias

In a typical peak shift experiment, control subjects are trained to respond (e.g., button-press) to a positively reinforced stimulus ($S+$, for example a line of a particular orientation or a color of a particular hue). Treatment subjects are trained in a manner identical to control subjects with respect to $S+$ and also to withhold response to an unreinforced or punished stimulus ($S-$, a line or hue similar but not identical to $S+$). Both groups of subjects are then tested without reinforcement on a continuum of stimuli comprising various line orientations or hues, including the training stimuli. **Figure 1** depicts the common finding of such discrimination learning experiments: during the test, control subjects respond most frequently to the $S+$ stimulus. Treatment subjects, however, respond most frequently to a stimulus they have never encountered before.

The treatment subjects' expression of a preference for an unrewarded and novel stimulus over $S+$ is somewhat paradoxical. If all test stimuli are discriminable, why are the treatment subjects ($S+/S-$ trained) not like the control subjects ($S+$ training only), responding most strongly to the stimulus they have learned is rewarding? The phenomenon that treatment subjects display in this type of experiment is known as peak shift. When the subjects' frequency of response is plotted as a function of stimulus value, data show a peaked response gradient. The stimulus receiving the maximum, or 'peak,' response by the treatment subjects is said to be 'shifted' relative to that of the control subjects.

An area shift is often noticeable, even in experiments that do not result in a significant peak shift. An area shift is characterized by an elevation of the rates of response to the novel stimuli on the side of $S+$ away from $S-$. Area shift often co-occurs with peak shift. In addition, a shift of the most strongly avoided stimulus is also produced, off of $S-$ in a direction away from $S+$.

Peak shift is considered to be a general outcome of generalization accompanying discrimination learning. In a typical peak shift experiment, the stimuli are simple sensory perceptions, monotonically increasing in value on the stimulus domain. The experiment consists of a control group and a treatment group each undergoing training (discrimination) and testing (generalization) phases. Stimuli are presented one at a time, for instance, on a

lighted button that the subject presses to indicate its response. The button is lit for a set length of time, and presses to the button while it is lit by a given color are the dependent variable. No reward is given during the testing phase. A variable intermittent reward schedule with all-or-none reward quantities may be given for correct responses during the training phase to minimize extinction during the testing phase. During training, incorrect responses to $S+$ and $S-$ are followed by mild punishment (e.g., lights turn off and a delay is imposed prior to the next trial). Intertrial intervals last 2–3 s, during which the response button is not lit. The training phase lasts until the treatment group reliably responds to the stimuli (e.g., 80% correct). The number of $S+$ training trials may be balanced across both groups, and a variety of reinforcement schedules, stimulus dimensions, and species have been used. As a phenomenon of learning, the magnitude of shift is a function of stimulus presentation parameters such as $S+$ and $S-$ similarity, reward value, and stimulus encounter rates.

Theoretical Accounts of Peak Shift

Peak shift might arise at any of several mechanistic levels of processing, including early peripheral sensory processes, well-learned associative mechanisms at the level of individual stimuli, or via a response to signal-borne risk occurring at the level of stimulus class. Peak shift is thought to arise from relatively uncomplicated mechanisms of learning. However, the 'right' explanation for a particular type of decision will shed light on larger issues of learning, preference establishment, and decision making.

Gradient Interaction Theories

Peak shift has been almost invariably attributed to Kenneth Spence's 1937 theory of overlapping gradients of excitation and inhibition. The summation of Spence's two gradients produces net excitatory and inhibitory behavioral impulses shifted from the training stimuli in the manner described earlier. On this and related accounts, differences in behavioral response strength arise from the level of behavioral excitation associated with each test stimulus. In terms of modern associative learning theory, bell-shaped gradients of positive associative strength (centered on $S+$) and negative associative strength

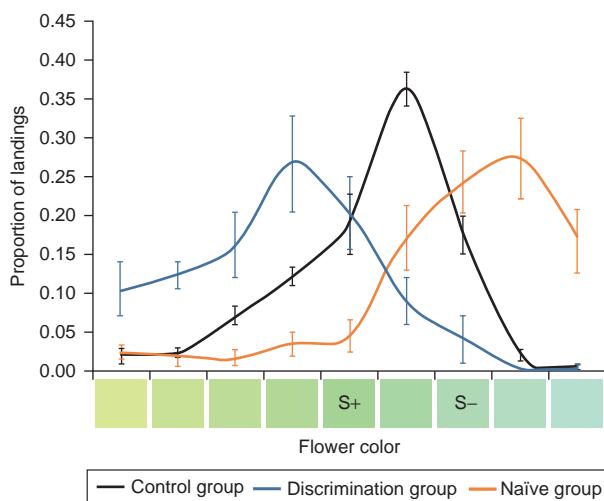


Figure 1 Typical results of a peak shift experiment. A control group was trained to approach the S_+ stimulus. A discrimination group was additionally trained to avoid S_- . AL test, the full range of stimuli was presented. The discrimination group exhibited a shift in their preferred stimulus off of S_+ in a direction away from S_- . These data are from bumble bees trained to discriminate the colors shown while learning to forage on artificial flowers. In addition, a Naïve group received no color training prior to testing, and exhibited an innate preference for bluish flowers. Examination of the control group shows that their peak response was slightly shifted toward the innate preference rather than centered over S_+ itself. Standard error of $n = 10$ bees per group is shown.

(centered on S_-) overlap and interact additively, yielding a maximum net positive associative strength shifted off of S_+ in a direction away from S_- . Differences in behavioral response are due to differences in the association strength between the test stimuli and the reward or punishment associated with responding to those stimuli.

Neural network models of peak shift are a kind of associative account in which nodal weight strength drives behavioral response strength. Nodal weight strength can be mathematically identical to Rescorla–Wagner-based measures of association strength. Peak shift has also been modeled as an additive overlap of sensory receptor excitations, an account at the level of peripheral sensory organs.

All gradient-interaction accounts assume a model of stimulus generalization that involves a spreading of knowledge from training stimuli to similar novel stimuli. The spreading is usually a Gaussian (i.e., bell-shaped) function of the perceptual similarity among stimuli. The S_+ and S_- generalization gradients then interact additively during decision making, producing peak shift.

No satisfactory gradient-interaction explanation of peak shift has arisen that accounts for the various forms of peak shift and the variety of stimulus domains over which it can be induced. The various explanations fail to account for area shift, are not extendable to the full range of stimuli prone to peak shift, or cannot produce the variety of observed gradient shapes.

Adaptation-Level Theory

A peak shift can also be explained by a modification of Harry Helson's theory of sensory adaptation-level, developed in 1947. Over the course of exposure to simple stimuli, such as line orientation, a subject's perceptual system can become 'adapted' to the range of stimulus variation. The adaptation level is centered at the mean of the stimuli encountered. This account of peak shift posits that subjects represent the S_+ and S_- stimuli relative to the adaptation level rather than in more absolute terms. So, during training, subjects learn that S_+ is located at, say, adaptation level +1 unit and that S_- is at, say, adaptation level -1 unit. When, during testing, the adaptation level changes as new stimuli are encountered, responding to adaptation level +1 produces an apparent peak shift. The subjects, however, have not changed their response to the stimuli per se, and this peak shift is not driven by conventional learning parameters. Rather, the baseline against which stimulus differences are evaluated has changed. Because under the adaptation level account stimuli are represented relative to the range of stimulus variation, this kind of peak shift is known as a 'range effect.'

Although a scenario of changing adaptation levels does account for some peak shift results, researchers can control the forces that drive range effects with techniques such as probe-tests administered throughout training. Also, more complex stimuli possessing multiple components or dimensions (e.g., facial expressions, orientation of clock hands) are resistant to range effects. Range effects can, and must, be controlled in studies focusing on peak shift arising from discrimination learning.

The Signals Approach

An account of peak shift can also be derived from signal detection theory (SDT; **Box 1**). This approach postulates that during testing, subjects experience uncertainty about which response is appropriate to give to a particular stimulus. Under SDT, uncertainty of choice-making is due to perceptual similarity of S_+ and S_- stimulus classes and carries a risk of stimulus misclassification. Under a signals approach, peak shift arises from an attempt to optimize stimulus classification rather than as strictly determined by associative strengths. As a signal detection issue, peak shift can be characterized as an aversion to signal-borne risk associated with the uncertainty of stimulus classification.

The three parameters of SDT that govern choice (distribution, relative probability of occurrence, and payoffs – see **Box 1**) correspond to elements of a discrimination learning experiment. (1) The appearance of the S_+ and S_- stimuli constitutes signals. The signal distributions are interpreted as gradients of likelihood that, on the basis of perceptual similarity, a particular stimulus is from the

Box 1: Signal Detection Theory and Extension to Nonthreshold-Based Decision Making

A signals approach to generalization and discrimination takes signal detection theory (SDT) as a descriptive, mechanistic model of decision making, as opposed to the theory's typical use as an analytical tool. SDT is a mathematical description of the trade-offs and risk inherent in the reception of signals (i.e., discerning one signal from another, or signal from noise). Classical SDT provides a functional description of how animals make choices among stimuli under conditions of uncertainty. Typically, that uncertainty is considered to arise from perceived variability in the appearance of stimuli. For example, variability in stimulus appearance could exist in the stimuli themselves or arise from noise in the sensory system. However, the uncertainty modeled by SDT may also arise from stimulus generalization, a process dependant on reinforcement history in addition to perception: inexperienced animals that are able to perceptually distinguish stimuli very well may yet be uncertain as to what response is appropriate to give to a particular stimulus. Like perceptual uncertainty, this response uncertainty can be described by signal detection theory.

Viewing generalization and discrimination as exercises in signal detection is straightforward. Consider a task in which subjects must approach yellowish stimuli and avoid bluish stimuli. In the overlapping greenish region, any given stimulus might be of one class (S_+ , deserving response) or another (S_- , to be ignored). As a model, the signals approach posits that animals estimate information about stimulus encounters to determine to which stimuli it is on average most profitable to respond, such that the number of correct detections of S_+ and correct rejections of S_- are maximized while missed detections of S_+ and false alarm responses to S_- are minimized. Under SDT, response strength is determined by three signal parameters: (1) the stimulus distributions over a perceptual domain, (2) the relative probability of encountering stimuli of one class or another, and (3) the payoff (reward or punishment accrued) for responding to or ignoring S_+ and S_- stimuli.

These signal parameters can be combined in a utility function, the maximum of which locates the optimal placement for a response threshold on the stimulus domain:

$$U(x) = \alpha hP[CD] + \alpha mP[MD] + (1 - \alpha)aP[FA] + (1 - \alpha)jP[CR] \quad [1]$$

where $U(x)$ is the estimated utility over stimulus domain x ; $P[CD]$, the probability of correct detection, measured as the integral of the S_+ distribution from threshold to negative infinity; $P[MD]$, the probability of missed detection, equal to $1 - P[CD]$; $P[FA]$, probability of false alarm, measured as the integral of the S_- distribution from threshold to negative infinity; $P[CR]$, the probability of correct rejection, equal to $1 - P[FA]$; α , the relative probability of encountering an S_+ signal, and $1 - \alpha$ equals the relative probability of encountering a signal from the S_- distribution; h , the benefit of correct detection of S_+ ; m , the cost of missed detection of S_+ ; a , the cost of false alarm response to S_- ; and j , the benefit of correct rejection of S_- . Costs may be negative or simply less positive than benefits, so long as $h > m$ and $j > a$.

In classical SDT, signal distributions over the continuous sensory domain are considered to be probability density functions (PDFs). The probabilities of correct detection, false alarm, missed detection, and correct rejection are calculated by integrating the respective PDFs from each possible threshold location to infinity (or by taking one minus that integral). This integration permits locating the optimal threshold placement (the maximum of eqn [1]). Threshold placement produces a stepped or sigmoid response gradient of dichotomous response strengths: equally strong response on one side of the threshold, and equally weak on the other. To apply SDT as a model of behavior to discrimination tasks in which subjects do not show a threshold-based response, such as peak shift, the assumption of integrated signal distributions can be changed. Substituting the integration of the PDF with the probability density, y_i , of a signal of a given value, x_i , allows signal detection theory to produce continuously variable response strengths.

One way to conceptualize this substitution biologically is to suppose that animals perceive signal variation discretely or construct discrete signal distributions rather than continuous probability density functions. Biologically, one may interpret this as an assumption that although signals may indeed fall along a continuous distribution objectively, animals perceive stimuli in flexible intervals of just-noticeable-difference, the dynamic width and placement of which is determined by contextual factors and the limits of their sensory acuity.

To reflect this substitution, eqn [1] can be modified to yield the expected utility of responding to a signal of $x = x_i$, a particular value, rather than a signal of $x \geq x_i$ as would result from PDF integration:

$$U(x) = [\alpha h f_{S_+}(x_i) + (1 - \alpha) a f_{S_-}(x_i)] - [\alpha m f_{S_+}(x_i) + (1 - \alpha) j f_{S_-}(x_i)] \quad [2]$$

where $U(x)$ is the utility of responding to a stimulus of a given value, x_i (correct detections and false alarms), less the utility of withholding response to that signal (missed detections and correct rejections); $f_{S_+}(x_i)$, the relative frequency of a stimulus of value x_i from the S_+ signal distribution; $f_{S_-}(x_i)$, the relative frequency of a stimulus of value x_i from the S_- signal distribution; and other variables are as for eqn [1].

Equation [2] produces a pulse-shaped gradient exhibiting peak shift. It is positive (utility > 0) for all stimulus values for which responding yields a net benefit. Like eqn [1], it provides a mechanism by which to make choices in the face of uncertain stimulus classification, not by reducing the uncertainty, but by allowing an animal to estimate to which signals it will be on average be most profitable to respond. Equation [2] can still produce threshold-based behavior, by solving for zero, the x -intercept, rather than maximizing the function.

S_+ or S_- stimulus class. (2) The relative probability of encountering an S_+ or S_- signal corresponds to the relative frequency of S_+ and S_- stimulus presentation during training. (3) The payoffs correspond to the value of reinforcement and punishment for responding to and ignoring stimulus presentations during training. The utility function (Box 1, eqn [2]) combines the signal

parameters learned during training to produce a pulse-shaped 'response gradient,' the maximum and minimum of which exhibit peak-shift (Figure 2). Peak shift can thus be interpreted as a signal detection issue. However, rather than using a threshold to dictate choice, as in other applications of signal detection theory, the maxima and minima of the pulse-shaped generalization gradient are used. The

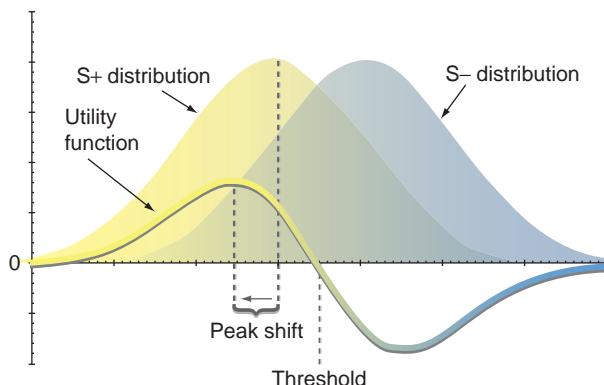


Figure 2 On the signals approach, S_+ and S_- signal distributions (shown here as yellow and blue bell-shaped gradients) represent the subject's estimate of the likelihood that a particular stimulus is from the S_+ or S_- stimulus class. Signal distributions may arise from actual signal variation, noise in the perceptual system, or cognitive generalization. Overlapping distributions produce uncertainty about which response (e.g., approach or avoid) is appropriate to give to any given stimulus. Behavioral response is dictated by a utility function (Box 1, eqn [2]) that integrates the signal distributions, the estimated probability of encountering an exemplar of either stimulus class, and the payoffs expected from correct and incorrect responses. The maximum and minimum of the utility function exhibit peak shift.

magnitude of the peak shift displacement is sensitive to variations in the three signal parameters.

Functionally, the peak shift experiment is a signal discrimination task in which animals are uncertain as to which response, approach or avoidance, is appropriate for any given test stimulus. Though the mathematics of the signals approach provides a functional account of peak shift, the approach can also be interpreted mechanistically, as an alternative to the gradient interaction accounts. According to this interpretation, peak shift is not a failure to perceptually discriminate S_+ from similar signals, nor is it an artifact of overlapping gradients of excitation or associative strength. Rather, peak shift reflects an optimization of response (i.e., committing a number of unavoidable mistakes to achieve correct responses) in cases when subjects experience uncertainty about stimulus classification.

Decision Making at the Intersection of Comparative Psychology and Behavioral Ecology

Peak shift intrigues behavioral researchers for at least two reasons. First, its apparent universality makes peak shift a model system in which to study decision making. Second, when it occurs in situations in which decision-makers evaluate stimuli linked to another organism's reproductive success, then peak shift has the potential to exert selective pressure on the evolution of morphology and communication.

Peak Shift Is a 'Model' Decision

Peak shift is taxonomically widespread: exhibited by birds; mammals, including humans; fish; and at least some arthropods. The phenomenon thus appears to reflect universal attributes of generalization, discrimination learning, and choice-making behavior. As such, peak shift is a 'model' type of decision making, suitable for comparative study at functional and mechanistic levels. Using peak shift as a tractable example of decision making, a variety of organisms can be studied, with strengths differentially well suited to phylogenetic, behavioral, neural, cellular, or molecular investigations.

In addition to being well suited to study at multiple levels, considerations of peak shift go beyond what is typically investigated in research on decision making. Many models of behavioral economics maximize utility: these models consider variability in (1) the costs and benefits of obtaining resources, and how those payoffs change with body state, and (2) the probability of encountering resources of some quality. Game theoretic approaches additionally account for the effect of others' responses on the decision maker's own behavior. However, these models overlook the fact that an animal's estimates of a resource's payoff and probability are based on sensory signals emitted by the resource. Outside of the laboratory, signals, such as color or tail length, vary. This variation may exist independently of any variation in the information encoded by the signals. For example, a signal that indicates a particular food quality (yellow skin on a banana signals ripeness) may vary even if the food quality itself does not (ten bananas of the same ripeness may not share the same yellow color). Typical utility optimization approaches account for variance in resource quality, not variance in the stimuli that signal that quality. Since real world signals are noisy, our understanding of choice behavior will be incomplete without accounting for signal variation and uncertainty. As a signal detection issue, peak shift experiments present an opportunity to investigate the role of this signal-borne risk in decision making and its interactions with those aspects of decision making more commonly investigated.

Significance of Peak Shift for Evolution of Signaling

Peak shift has also been recognized as a possible influence on the evolution of signaling systems. Theoretical development has explored the potential role of peak shift in the evolution of gender or species recognition characters, warning coloration, and sexually dimorphic exaggerated traits. Experiments have shown that peak shift can drive signal evolution in warning coloration, mimicry, and mate-selection systems. Cognitive underpinnings of behavior may thus have a role as a selective mechanism driving evolution, in addition to the better known

interactions between genetic and environmental factors that form the basis of natural selection. Two examples are given in the following section.

Evolution of mimicry and crypsis

Bumblebees (*Bombylius impatiens*) foraging for nectar exhibited peak shift when choosing the flowers to visit (and thus pollinate). In a laboratory study implementing a Bayesian mimicry system, bees were trained to forage on artificial flowers (colored paper disks) under different signal parameter sets. During training, positions of 36 S+ and S- flowers, present simultaneously, were randomized in a 6 × 6 array on the floor of a flight cage. ‘Baseline’ bees received an arbitrary parameter set specifying the color and number of S+ and S- flower types, and the sugar-water reward for visiting the two flower types. Three groups of comparison bees each differed from baseline by manipulating one of the three signal parameters: increased variance of S- appearance (three S- flower colors used, whereas baseline used one), decreased relative abundance of S+ (28% of stimuli were S+ flowers, whereas baseline had 50%), and decreased reward for correct detection of S+ (33% sucrose concentration, whereas baseline used 50%). When tested on a range of nine colors (four exemplars each in random positions in the flight cage), the baseline bees exhibited peak shift relative to a control group that had received no S- training. Furthermore, as predicted by the signals approach, the comparison bees exhibited larger peak shift and area shift over and above that exhibited by the baseline bees, in accordance with the increased signal-borne risk of their training regimes. Simultaneous presentation of all test stimuli was used as a way to mitigate range effects. Also, range effects do not explain the greater shift produced by increased signal variance (which maintained the same adaptation level as the baseline condition). The results indicate that in natural situations of mimicry (two signals resembling one another) or crypsis (a signal being difficult to distinguish from noise), peak shift can influence the evolution of signaling traits.

Sexual selection

Male zebra finches (*Taeniopygia guttata*) exhibited peak shift when deciding which females to court. As nestlings, chicks imprinted on parental beak color, which was manipulated with paint. The nestling period thus corresponded to the training phase of a peak shift experiment. Male chicks learned that beak color could distinguish their father (the S- exemplar since courting another male will not lead to reproductive benefits) from their mother (the S+ exemplar, being a female). As adults, the males were tested by allowing them to court other zebra finches possessing a range of beak colors. The birds exhibited a preference to court birds with beak colors shifted

off that of their mothers in a direction away from that of their fathers. Models for sexual selection of exaggerated phenotypes typically require a genetic association between a sex-linked trait exhibited by one gender and a preference for that trait exhibited by the other gender. In this study, however, beak color carried no inherent fitness advantage, did not communicate the possibility of ‘good genes’ to the courtiers, and did not impart a competitive advantage to future offspring that might possess a particular beak color. The preference for the exaggerated trait was neither genetically predisposed nor based on a sensory bias, but was learned.

Where Does Peak Shift Fit in the Larger Space of Choice-Making?

Peak shift is characterized by uncertainty inherent to perceptual similarity of stimuli that vary on a continuum. A peak shift experiment and the phenomenon of the shift itself thus differ in several ways from topics more commonly treated as decision making, such as choosing among several discrete alternatives (e.g., diet choice), optimizing resource acquisition (e.g., foraging under time, energy, or predation constraints), and investment budgeting (e.g., parental care, life history pattern).

In peak shift, a ‘hidden’ preference for a novel stimulus is established over and above that for a known, rewarded stimulus (S+). From the perspective of models that do not account for generalization, the shift seems somewhat paradoxical. Additionally, preferences are being learned in the absence of the preferred stimuli and are shaped by the presence of unpreferred stimuli (S-). For example, changes in an unpreferred stimulus’ encounter rate can influence choice, contrary to diet-choice models in which the abundance of less-preferred food items has no effect on intake of preferred items.

Many writers have highlighted features shared between peak shift and phenomena such as transitive inference, novelty seeking, extreme seeking, response to supernormal stimuli, artistic caricatures, esthetic preferences, and sensory bias. Though parallels can be seen, peak shift may not be responsible for any of these behavioral phenomena. For example, when driven by an aversion to signal-borne risk, peak shift does not usually produce a preference for stimuli that are extremely different from the training exemplars, or for novelty per se; the peak shifts are enough to only partially mitigate risk of mistakes. A unified mathematical description of choice making under risk and uncertainty, such as the signals approach, could help distinguish among these phenomena or make the parallels more mechanistically concrete.

See also: Decision-Making: Foraging; Rational Choice Behavior: Definitions and Evidence.

Further Reading

- Blough DS (2001) Some contributions of signal detection theory to the analysis of stimulus control in animals. *Behavioural Processes* 54: 127–136.
- Cheng K (2002) Generalization: Mechanistic and functional explanations. *Animal Cognition* 5: 33–40.
- Green DM and Swets JA (1966) *Signal Detection Theory and Psychophysics*. New York: Wiley.
- Lynn SK (2006) Cognition and evolution: Learning and the evolution of sex traits. *Current Biology* 16: R421–R423.
- Lynn SK, Cnaani J, and Papaj D (2005) Peak shift discrimination learning as a mechanism of signal evolution. *Evolution* 59: 1300–1305.
- Purtilo RB (1973) Peak shift: A review. *Psychological Bulletin* 80: 408–421.
- Sperling G (1984) A unified theory of attention and signal detection. In: Parasuraman R and Davies RR (eds.) *Varieties of Attention*, pp. 103–181. Orlando: Academic Press.
- ten Cate C, Verzijden M, and Etman E (2006) Sexual imprinting can induce sexual preferences for exaggerated parental traits. *Current Biology* 16: 1128–1132.
- Thomas DR, Mood K, Morrison S, and Wiertelak E (1991) Peak shift revisited: a test of alternative interpretations. *Journal of Experimental Psychology: Animal Behavior Processes* 17: 130–140.
- Wiley RH (2006) Signal detection and animal communication. *Advances in the Study of Animal Behavior* 36: 217–247.

Relevant Websites

http://eebweb.arizona.edu/Animal_Behavior – Bumblebee Decisions.