Published in final edited form as:

Behav Processes. 2010 October ; 85(3): 209–214. doi:10.1016/j.beproc.2010.07.013.

A random-walk interpretation of incentive effects in visual discrimination

Donald S. Blough

Brown University, Providence, RI 02912

Abstract

Previous studies have shown that changes in stimulus discriminability and changes in reward density affect pigeon reaction-time (RT) distributions in different ways (Blough, 2004). A random-walk model (“RWP”) accounts for these differences and assigns a single parameter to each of the independent variables. This paper briefly reviews the model and illustrates its findings with hue discrimination data. A new analysis then presents fits to data showing that increased reward for stimulus “A” lengthens RT of pecks to an alternative stimulus “B”, and that this effect on RT distributions is much the same as the effect caused by reduction of reward to B. RWP account for both effects by changes in its “bias” parameter. The remainder of the paper comments on the relations between reward, RT, incentive and bias.

Keywords

discrimination; incentive; random-walk; pigeon

1. Introduction

This paper reviews some effects of reward manipulations on discriminative reaction-time (RT) distributions and describes a random-walk model that accounts for these effects. The review integrates material on this topic drawn from other sources (Blough, 1989, 2000, 2004, 2007, 2009, in press) and also presents a new analysis of RT contrast effects. It is a pleasure for me to discuss these matters in an issue honoring Al Riley, for over the years Al and his students have contributed a great deal to our understanding of discrimination and variables that affect it.

Discriminative RTs must somehow reflect the time taken by underlying processes, and, as a behavioral measure, RT has several unique and informative characteristics. Unlike a single “yes/no” or choice response, each RT offers quantitative information. Many RTs together yield a distribution with a characteristic shape, and a pattern of results may emerge that strictly limits the set of applicable models, especially when such distributions are used in conjunction with other dependent variables. In particular, RT data coming from studies of human memory and discrimination, and now from animal research, strongly point to a model that generates each

Address correspondence to Donald S. Blough, Department of Psychology, Box 1853, Brown University, Providence RI 02912.
Donald_Blough@brown.edu.

Publisher's Disclaimer: This is a PDF file of an unedited manuscript that has been accepted for publication. As a service to our customers we are providing this early version of the manuscript. The manuscript will undergo copyediting, typesetting, and review of the resulting proof before it is published in its final citable form. Please note that during the production process errors may be discovered which could affect the content, and all legal disclaimers that apply to the journal pertain.

RT from the accumulation of discriminative information through time, with the accumulation taking the form of a random-walk or diffusion process (e.g. Luce, 1986).

The model I outline here is an application of the random walk idea to data from pigeon visual discriminations. The model is called “RWP”, for “random walk, pigeon”; it is essentially a variant of a model that Ratcliff has applied with impressive success to a range of human cognitive data (e.g. 1978, 2002). The Ratcliff model and others developed around human data are generally more detailed and theoretically sophisticated, but RWP seems to isolate key elements of discriminative decision making. The model is especially valuable here in accounting for reward effects that have been relatively neglected in the work with humans. A full account of the model appears elsewhere (Blough, in press). In the following pages I briefly describe the model and summarize sample data to illustrate how it works and to introduce some key results. Then I apply the model to a set of RT distributions from a visual search study involving reward variation. The remainder of the paper discusses theoretical issues raised by these results.

2. A reference experiment: hue discrimination

The following brief account of a reference experiment will illustrate key findings and provide a context for a description of the RWP model. The experiment used a discrete-trial go/no-go procedure (for details see Blough, in press, Experiment 2). On each trial a stimulus spot appeared on a screen in a standard operant chamber; this spot stayed on until pecked or for 3 sec. at most. An RT was defined as the time between stimulus onset and a peck at the spot. On some trials the spot was a bright red “S+”; pecks to this hue brought food 100% of the time during Phase 1, which lasted for a block of experimental sessions, then 15% in Phase 2, then 100% again in Phase 3. Less saturated red spots appeared on most trials and pecks to these always ended the trial without food. Thus there were two independent variables, percent reward to the S+, which varied between phases (blocks of sessions), and similarity of stimuli to the S+, which varied within sessions. A green spot also appeared on some trials, and pecks to it always brought food.

Figure 1 illustrates the effects of the reward and similarity variables on RT distributions to the red S+, two unrewarded S- reds, and the green stimulus that was rewarded 100% throughout. Shown are mean frequency distributions of RTs under each of the reward conditions, with data from the two 100% phases averaged together. Data appear as points and model fits, described below, appear as lines. A reward effect appears clearly in each pair of “red” distributions, where the 15% reward function is shifted to the right of the 100% function. The distribution for the green stimulus also shifted slightly when reward for red S+ dropped to 15%. The difference between mean RTs for the two reward conditions was significant for the data from each of the four stimuli. (See Blough, in press, for all statistics.)

A hue effect appears across the three sets of “red” distributions. The change in hue from S+ to S2 and S3 caused the distributions shrink in size; that is, fewer responses were emitted as stimuli became less similar to S+. The S2 and S3 distributions also moved slightly to the right of that for S+. The percentage of responses emitted did not change significantly between the reward conditions for any stimulus, for the lower peak value of the 15% curves was compensated by an increased number of long RTs..

3. A random-walk model of discrimination (“RWP”)

Like the Ratcliff diffusion model on which it is based, RWP models the processes involved in simple two-choice decisions, such as the peck/no-peck decisions used in the reference experiment just described. RWP is used to fit RT distributions, as in Figure 1, and it does this by generating simulated distributions from a large number of simulated trials. Each of the

simulated RTs is the sum of two intervals. One of these is a “residual time” taken by events, such as head motions, that are unrelated to the response decision. The other interval is the “decision time,” which is the time taken for the accumulation of choice-related information to reach a threshold. The information arrives over time in variable amounts, and the running sum of these amounts constitutes a random walk.

Figure 2 sketches the process that generates RTs. Parameters of the equations that control the process are identified at the left. The RT interval starts at “0” on the horizontal axis, when the stimulus appears. The “residual time” is inserted here; it represents the sum of all times taken by non-decision events, although these might actually occur at any time. Next, random walks for several trials are sketched as wavy lines. “Bias” specifies the level at which each walk starts; for example, a large positive bias would mean that the walk starts near the upper response threshold. “Step-size” specifies the average speed of accumulation, represented by the overall slope of the random walk. The random walk goes up when incoming information favors a response, or down if it favors non-response. The steps up or down are highly variable; for a positive stimulus (S+) the walk tends to rise steeply, and for stimuli dissimilar to S+ it rises more slowly, or falls. The random walk continues until the total accumulation reaches either the upper threshold, which triggers a response, or the lower threshold, which yields no response. The RT of a response is the time taken for the walk to reach the response threshold, plus the residual time. The function at the top of Figure 2 suggests an RT frequency distribution that might be generated by many trials with a positive stimulus. Additional details about simulation appear in the Appendix; for a full account see Blough (in press).

A computer program implemented the process sketched in Figure 2 to simulate RTs and find theoretical distributions that best fit the data distributions. (See below, Appendix and Blough, in press.) The resulting best-fit distributions appear as lines in Figure 1; these are means across the simulated data of 6 birds. The fitting lines track the data points rather well, including the shifts to the right induced by reward reduction, the reduction in response percentage induced by decreased similarity, and the small RT shifts induced by changes in hue similarity.

The step-size and bias parameters are key to fitting these data. Recall that step-size controls the speed of the random walk’s overall upward or downward trend, whereas bias controls the starting level of the walk. Thus, step-size and bias both affect the time of arrival at a threshold and they might seem to affect RT in similar ways. But, as it turns out, the pattern of their effects is quite different. Figure 3 underlines the difference by displaying the best-fitting values of these parameters for the various conditions. Values for the various stimuli are separated along the abscissa. Black dots show means for the 100% reward phases and white dots show those from the 15% phase. The top panel shows that step-size varied strongly with stimulus hue but not with reward. The bottom panel shows that bias varied strongly with reward but not with stimulus hue. These effects appear to be independent of each other; a second experiment confirmed this separation (Blough, in press). The remainder of this paper further explores the nature of these relationships, particularly the effect of reward and its relation to bias.

4. RT contrast in visual search

Previous research has used RWP to fit data from variations on the hue discrimination just described. Another opportunity to explore the reward-bias relationship comes from RT distributions generated in an old experiment on visual search (Blough, 1989). That experiment demonstrated an RT “contrast” effect; that is, searches for one target took longer when the reward percentage for responses to a second target was increased. Here I show fits of RWP to this search data set and discuss their implications. Unfortunately, only mean data from the experiment are still available, so the fits are less certain than those based on fits to individual data, but they appear to capture consistent main effects.

Briefly, the search experiment went as follows (for details see Blough, 1989). Three birds learned to peck at the target letters “A” and “E.” Each experimental session consisted of many brief trials; on half of these “A” appeared on a screen, on the other half “E” appeared. Pecks at these letters brought food on some trials; RT was the time from display onset until the target letter was pecked. Each target letter appeared either alone or together with distractor letters scattered randomly over the screen. The number of distractors varied (0, 2, or 18), as did their similarity to the target. However, the key variable for our purposes was a change in reward density from one block of experimental sessions to the next: During the “baseline” phase, pecks to both the “A” target and the “E” target brought food on 10% of trials. During “contrast” phase, pecks to “E” brought food on 30% of trials, whereas “A” remained at 10%. Each block of sessions contained 540 trials of each display type, and responses to these yielded an RT distribution for each bird and display type.

Figure 4 shows the mean RT distributions for pecks to the “A” target; pecks to “E” are not shown (see below). The five pairs of curves come from five display conditions. Black dots show data from the baseline phase, open circles show data from the contrast phase, and RWP fits appear as lines drawn near the data points from each phase. Fits were produced as outlined above and in the Appendix, with the important exception that variation in the “residual time” parameter was necessary to achieve fits of data from the various display conditions, as described below. Comparing the curves in the different panels of Figure 4, we see that the RT distributions shifted to the right when distractors were added, and that this effect was a bit larger for the “P” distractor than for the “Z”. However, our main interest is in the difference between baseline and contrast conditions. Contrast RTs were longer than baseline RTs. This difference seems to be about the same size under the different distractor conditions, and it looks much like the reward-induced difference in the reference experiment (Figure 1). As the fitting lines show, RWP gave a good account of all the distributions.

The “best-fits” shown in Figure 4 required the bias parameter of RWP to differ substantially for baseline and contrast phases, but, as in the reference experiment above, this change was remarkably constant across different stimulus conditions. For four of the pairs shown in Figure 4 the change in “best-fitting” bias between baseline and contrast was the same (9 steps); only the “A & 18P” condition called for a somewhat higher bias change (12 steps). By comparison, the step-size parameter showed only small changes between the phases ($-.03$ to $+.08$ steps), and these were not consistently related to the display variables. The other parameters of the model were held constant across the baseline-contrast fits, but varied otherwise as will be described below.

In summary, the indirect effect of increasing reward to an alternative target in the search experiment was empirically similar to the direct effect of decreasing reward in the reference experiment. RWP assigned both of these effects to a change in bias, and the amount of this change was about the same across display conditions. It seems appropriate to hypothesize that both reward effects stem from a common source, about which more is said below.

The relative constancy of the contrast effect across display conditions suggests that the effect of reward occurs independently of the display changes. However, in this search experiment, the display conditions clearly affected the RTs, for displays that contained distractors yielded distributions shifted to the right of those with “A alone” (Figure 4). Although RWP provided good fits for these various conditions, the meaning of the adjustments required by the different displays is problematic. The best fits for the different displays called for substantial adjustments in bias, but they also required changes in the mean and variance of the residual distribution, which theoretically is not involved in decision making.

That fitting search data required a departure from the model assumptions is not surprising, for visual search differs a great deal from hue discrimination, both in procedure and also in the behavior required for correct response. In fact, existing data support the idea that a major component of these RTs arises from a serial search of the display during which a bird's chance of finding a target is constant from moment to moment (Blough, 1988). This in turn generates an exponential component of the search RT distributions that is not found in RWP. Evidently the RWP fitting program compensated for this lack by varying the bias and residual parameters. Despite this departure from the model, it is important to note that within each pair of baseline-contrast distributions the display conditions were identical and the residual parameters were held constant. Thus it seems reasonable to assume that the difference between baseline and contrast results can be attributed to a change in bias as described above.

Figure 4 shows no data for responses to the "E" target, despite the fact that reward to "E" changed from 10% to 30% between phases. These data are omitted because "E" distributions for the baseline and contrast phases fell almost exactly on top of each other and the mean RTs to "E" for the two phases differed unsystematically by an average of less than .007 sec. This identity in the results from the two reward conditions is quite interesting in itself; it might have a peripheral source, possibly some sort of ceiling on speed, but more likely it has to do with the dynamics of the process by which incentive value is determined. This mystery receives further mention below. The effect of the display variables on searches for the E target were very similar to those for A that appear in Figure 4 and they will not be discussed further.

The analysis just described supports the previous conclusion that reward manipulations shift RT distributions along the time axis and that these shifts correspond to a bias change in the RWP model. The new conclusion is that these consequences follow a relative downward shift in reward to the target stimulus that happens when reward to an alternative stimulus increases. In the comments that follow, I call the RT effects that are associated with reward changes simply "the RT-reward" effect. This effect, and its theoretical interpretation, dominate the following discussion.

5. RT, incentive and bias

The experiments above, together with other results (Blough, 2000, 2004, in press) showed that, in the discrete-trial discriminations tested, a relative decrease in reward shifted RT distributions to the right in a characteristic way. The rest of this paper discusses this RT-reward effect: I suggest that RT is largely determined by the incentive value of the stimulus responded to, that this incentive value flows from a Pavlovian association of the stimulus and reward, and that the bias parameter of RWP represents this incentive value.

5.1 Incentive from Pavlovian association

Because RT varies with reward density, one might think that it would correlate highly with other measures of response or associative strength that also change with reward. However, the reference experiment above showed that percent response and RT responded differently to changes in reward density. RWP gave one interpretation of these results, but they might be understood in a different way via a familiar associative formulation, the "three-term contingency" of Skinner, sometimes known as the "stimulus-response-outcome" ("S-R-O") triad. The consequence of this contingency is that a stimulus comes to control a response when, in its presence, a response yields a motivationally significant outcome. We may assume that this triad controls the emission of discriminative responses in the above experiments. However, the S-R-O triad also contains a Pavlovian "S-O" contingency, and this contingency may confer additional properties on the stimulus that are unrelated to the "R" that yields the "O". For example, if the "O" is reward, the stimulus may arouse or activate behavior, it may be preferred

to other stimuli, and it may elicit preparatory responses, including approach, that are not explicitly reinforced (e.g. Dickinson & Balleine, 2002).

Two experiments suggest that RT may act like a Pavlovian “approach” response in the preparations reviewed above (Blough, 2009). These studies separated the stimulus-reward correlation from the response-reward correlation: the first correlation was varied and the second correlation was not. In both experiments, pigeons chose left for blue and right for green on discrete trials. The birds got food more often for one of the correct choices, “blue-left” for example, than for the other. However, this imbalance of reward for choice was achieved in two different ways. In some blocks of sessions reward probability was the same for each hue but the relative numbers of “blue” and “green” trials were unequal. In other blocks of sessions there were equal numbers of “blue” and “green” trials but reward probabilities differed. The two treatments were the same with respect to the relation of response to reward, because they provided the same disparity between percent reward for “left” and percent reward for “right.” However, the two treatments differed with respect to the percent reward assigned to each stimulus.

As in the experiments outlined earlier, this experiment produced different results for percent response than for RT. For both experimental treatments the birds correctly chose the more rewarded side more often than they correctly chose the less rewarded side, and this choice bias was about equal whether it had been produced by changing the relative reward for the two stimuli or changing the relative frequency of the two stimuli. However, RTs showed a different pattern. When hue frequencies were equated but the reward probability differed for the two hues, the bird pecked quickly when the hue signaled frequent reward and slowly when the hue signaled infrequent reward. Note that here the “S-O” (Pavlovian) association is imbalanced: more rewards per unit time follow one hue than the other. In contrast, RTs did not differ when the “S-O” association remained in balance; that is, when each hue was rewarded at an equal rate. In sum, RTs were not low for the more rewarded *response*; rather, they were low for the more rewarded *stimulus*. This stimulus-reward connection was especially explicit in the second of the two experiments. There, the stimulus appeared in one location and the choice responses were made in other locations. RTs of “observing” pecks to the stimulus varied in the manner just stated; RTs of the choice responses did not vary. Both results suggest that RT in this sort of experiment reflects the incentive value of the stimulus, dictated by the Pavlovian correlation of reward and stimulus. (see Blough, 2009 for further discussion).

In addition to this evidence that RT reflects a stimulus-specific incentive effect, there are indications that the RT-reward effect may include a general arousal component. An example of this appears in the small but significant shift in the “green” RT distributions depicted at the top of Figure 1. This shift followed the change in reward to the red stimulus which is mirrored in the small decrease in the bias parameter used to fit these distributions (Figure 3). This effect might be attributed either to stimulus generalization from the red to the green stimulus (both were small spots on a dark background) or to a more general effect of the change in overall reward frequency. Evidence of a general arousal effect of reward on RT also surfaced in another discrete-trial discrimination experiment, where it was found that RT could be predicted by summing a large stimulus-specific effect with a smaller general effect (Blough, 2007).

5.2. RT, incentive and contrast

In the preceding section I suggested that the RT-reward effect may represent a Pavlovian stimulus-reward association. The efficacy of this association has often been noted before, and it is commonly said to confer an “incentive value” on the stimulus (e.g. Dickinson & Balleine, 2002). Because I have suggested that the RT-reward effect includes contrast results, it is appropriate to ask whether incentive might play a role in other observations of behavioral contrast.

In a comprehensive review of the contrast literature Williams (2002) touched on key data and discussed the major alternative models that have been used to account for them. The data he addressed consisted mostly of response rates that came from operant procedures, but a few other measures were mentioned, including pecking elicited by signaled reward, consummatory licking, and choice percentages. It may come as a surprise, however, that Williams dismissed discrete-trial RT as an effective indicator of contrast. This conclusion arose from a finding by Mackintosh (1974) that discrete-trial responses to a positive stimulus were not speeded when trials with nonreinforced S- stimuli were included in the testing sessions.

RT contrast in discrete trial situations is something of a puzzle. The Mackintosh outcome just mentioned seems to disagree with the evidence for RT contrast presented here (from Blough 1989), and also with observations of strong RT contrast that appeared in another pigeon discrimination experiment (Blough, 1992). Note, however, that both of those experiments showed an *increase* in RT in the contrast condition, consequent on increased reward to an alternative stimulus. Thus the results might be reconciled if one is willing to speculate that reward manipulation can *increase* RT, either directly (as the reference experiment above) or indirectly, as in the visual-search experiment, but that it does not *decrease* RT below a pre-established baseline (as in Macintosh 1974). This idea is further supported by two observations noted above. First, we saw in the search experiment that RT to the “E” target did not change when reward to it was increased. Second, we noted in the reference experiment a slight shift in the RT distribution for the green stimulus (Figure 1) toward *longer* RTs, rather than shorter, when reinforcement dropped to the alternative red stimulus. I noted that this shift might represent a general arousal from an overall increase in reward, but in any case it is contrary to the contrast expectation. This interesting asymmetry calls for verification and exploration; it could pose a problem for any complete model that addresses RT.

Whatever the role of RT in contrast, it can still be said that incentive shift is a strong candidate among theoretical accounts of behavioral contrast. Though Williams (2002) was unable to endorse any theory as a full explanation of contrast, he reviews evidence based on stimulus preferences that positive contrast with respect to response rate does reflect an increase in stimulus value. Other evidence is not so supportive, but the literature also reflects some difficulties in determining the roles of Pavlovian contingencies in the various experimental paradigms that he considered.

5.3. Random-walk bias and SDT bias

The above account suggests that the RT-reward effect can be attributed to incentive changes, and these changes are reflected in the bias parameter of RWP. Many readers will surely be reminded that reward (“payoff”) manipulations also affect “bias” in signal detection theory (SDT). This familiar information-processing model distinguishes the effects of reward from the effects of stimulus discriminability (“signal”), and it shows how these may combine to control discriminative responses (e.g. Blough, 2001; Macmillan & Creelman, 2005). In the standard version of SDT stimulus information is represented by values distributed on a decision axis, and reward (among other variables) affects a criterion value on this axis. The proportions of the stimulus distribution that fall on either side of the criterion determine which response will occur.

Both SDT and RWP distinguish reward from stimulus effects, but they do this in different ways. Note particularly that SDT separates the influences of signal and bias as they combine to control the *choice response*. However, as we have seen, response percentages and RTs respond differently to reward changes. Indeed, SDT says nothing about RT, nor does it address incentive, arousal, or approach. Thus the present interpretation of the bias parameter in RWP, which is reflected primarily in RT rather than choice and perhaps derives from a Pavlovian

process, is quite different than the bias of SDT, which participates in controlling choice and has not been linked to RT¹.

Despite the differences just described, SDT bias and RWP bias are not necessarily incompatible. They can be reconciled by letting an SDT-like process modulate the input (“step-size”) to the random walk, leaving RWP’s starting-point bias to act as outlined above. Essentially this combination of processes has already been used in Ratcliff’s diffusion model, of which RWP is an offshoot (see Ratcliff, 1978, 2002).

5.4. Other theoretical relationships

I am not aware of other models in the animal-behavior literature that attempt to deal with the sort of RT data presented here, but related models might help to expand our understanding of the present approach. One particularly exciting example is the drift-diffusion model (“DDM”) model recently proposed by Simen and his associates (Simen, Cohen & Holmes, 2006; Simen et al, 2009) which they have used to generate classic reward-matching phenomena, among other things (Simen & Cohen, 2009). This mathematical model incorporates both a random walk process and a reward-density adjustment, and the authors also suggest a neural-network implementation, including some possible neural sites for the necessary processing. Though DDM differs markedly in detail from RWP, the basic structures are quite similar. Within that structure, an important difference may be that in DDM reward density adjusts the response threshold, whereas in RWP reward density adjusts bias. However, threshold and bias are closely related (see Appendix), so a reconciliation of this difference might not be too difficult.

DDM provides a way for rewards to influence thresholds and possibly bias, as just discussed. Other models may suggest ways that stimulus similarity could influence step-size (accumulation speed). For example, similarity might affect the speed of a memory search through a store of samples or “exemplars” of previous encounters with the various stimuli. Several models in the animal literature relate response accuracy to the contents of such a memory, and it would not seem difficult to derive the speed of successful search from this mechanism (e.g. Blough, 1998; Chase & Heinemann, 1991; Heinemann, 1983). An extension of RWP in this direction might yield a more complete and more broadly testable theory of discriminative processing.

6. Concluding Comments

This paper reviewed some RT data from discrete-trial discriminations and showed how they could be fit with a random-walk model, RWP. This analysis led to a somewhat speculative line of argument according to which (1) RT in discrete-trial situations reflects the Pavlovian-induced incentive value of discriminative stimuli, (2) this incentive value is influenced by the relative density of reward assigned to the discriminative stimuli, and (3) incentive value is represented by the bias parameter of RWP. Whatever the validity of this argument, it seems fair to conclude that RT supplies important information not carried by other behavioral measures, that stimulus incentive probably affects RT, and that random-walk models are worth serious consideration in connection with animal stimulus discrimination. These findings add to the growing body of research on animal cognition that, in many instances, helps to bring basic human and animal discriminative processes under a common conceptual umbrella.

¹Although the reward shift did not reduce response percentages in Experiment 1, an extreme bias shift toward the negative threshold can reduce percentages substantially. One can see this by imagining that the random-walk starts so close to the lower threshold that many walks end there. Thus, by varying bias enough, one can get large enough percentage changes to derive theoretical receiver-operating-characteristic (“ROC”) curves. These curves look superficially like those commonly displayed in SDT analyses, but the extreme bias shifts needed to produce them also yield RT changes far larger than anything seen in experimental data.

Acknowledgments

This research was supported in part by National Institutes of Health Grant MH61782.

References

- Blough DS. Quantitative relations between visual search speed and target-distractor similarity. *Perception & Psychophysics* 1988;43:57–71. [PubMed: 3340500]
- Blough DS. Contrast as seen in visual search reaction times. *Journal of the Experimental Analysis of Behavior* 1989;52:199–211. [PubMed: 2584913]
- Blough DS. Effects of stimulus frequency and reinforcement variables on reaction time. *Journal of the Experimental Analysis of Behavior* 1992;57:47–50. [PubMed: 16812648]
- Blough DS. Context reinforcement degrades discriminative control: A memory approach. *Journal of Experimental Psychology: Animal Behavior Processes* 1998;24:1–15.
- Blough DS. Effects of priming, discriminability and reinforcement on reaction-time components in pigeon visual search. *Journal of Experimental Psychology: Animal Behavior Processes* 2000;26:50–63. [PubMed: 10650543]
- Blough DS. Some contributions of signal detection theory to the analysis of stimulus control in animals. *Behavioural Processes* 2001;54:127–136. [PubMed: 11369465]
- Blough DS. Reaction time signatures of discriminative processes: Differential effects of stimulus similarity and incentive. *Learning & Behavior* 2004;32:157–172. [PubMed: 15281388]
- Blough DS. Reaction times in discriminations of varying difficulty: Decision modulated by arousal. *Behavioural Processes* 2007;74:234–243. [PubMed: 17101238]
- Blough DS. Reaction times identify a Pavlovian component in a two-choice discrimination. *Behavioural Processes* 2009;81:195–204. [PubMed: 19429212]
- Blough DS. A random-walk model of accuracy and reaction time applied to three experiments on pigeon visual discrimination. *Journal of Experimental Psychology: Animal Behavior Processes*. (in press).
- Chase, S.; Heinemann. Memory limitations in human and animal signal detection. In: Commons, ML.; Nevin, JA.; Davison, MC., editors. *Signal detection: Mechanisms, models, and applications*. Hillsdale, NJ: Erlbaum; 1991. p. 11–138.
- Dickinson, A.; Balleine, B. The role of learning in the operation of motivational systems. In: Pashler, H.; Gallistel, R., editors. *Stevens' Handbook of Experimental Psychology: Vol. 3. Learning, Motivation and Emotion*. 3. New York: Wiley; 2002. p. 497–533.
- Heinemann, EG. A memory model for decision processes in pigeons. In: Commons, ML.; Herrnstein, RJ.; Wagner, AR., editors. *Quantitative analyses of behavior, Vol. 4: Discriminative processes*. Cambridge, MA: Ballinger; 1983. p. 3–19.
- Luce, RD. *Response times*. New York: Oxford University Press; 1986.
- Macmillan, NA.; Creelman, CD. *Detection Theory: A User's Guide*. Cambridge University; New York: 2005.
- Mackintosh NJ. A search for contrast effects in discrete-trial discrimination learning by pigeons. *Learning and Motivation* 1974;5:311–327.
- MATLAB. Computer Program. Mathworks, Inc; Natick, MA: 2007.
- Ratcliff R. A theory of memory retrieval. *Psychological Review* 1978;85:59–108.
- Ratcliff R. A diffusion model account of reaction time and accuracy in a brightness discrimination task: Fitting real data and failing to fit fake but plausible data. *Psychonomic Bulletin & Review* 2002;9:278–291. [PubMed: 12120790]
- Simen P, Cohen JD. Explicit melioration by a neural diffusion model. *Brain Research* 2009;1299:95–117. [PubMed: 19646968]
- Simen P, Cohen JD, Holmes P. Rapid decision threshold modulation by reward rate in a neural network. *Neural Networks* 2006;19:1013–1026. [PubMed: 16987636]
- Simen P, Conteras D, Buck C, Hu P, Holmes P, Cohen JD. Reward rate optimization in two-alternative decision making: empirical tests of theoretical predictions. *Journal of Experimental Psychology: Human Perception and Performance* 2009;35:1865–1897. [PubMed: 19968441]

Williams BA. Behavioral contrast redux. *Animal Learning & Behavior* 2002;30:1–20. [PubMed: 12017964]

Appendix: Comments on simulation using RWP

Parameter variance

Doubtless all the parameters listed in the text and in Figure 2 have some variability, but for simplicity only two were allowed random variance. One of these is step-size, with a mean that set the average slope of the random walk, and a Gaussian variance that added a large variability to the walk. This arrangement simulates “noisy” input of information to the decision process. (The account of step-size given here is essentially correct, but its computation was actually more complex; see Blough, in press.) The residual time also had Gaussian variability, with a mean and variance constant for each bird and condition. (An exception to this was made for the search data, as explained in the text.) Bias had no random variability, though it changed to fit different stimuli and experimental conditions. The thresholds were held constant.

Treatment of thresholds

The upper and lower thresholds (Figure 2) were set symmetrically above and below the zero-bias point. Their values were based on preliminary fits and were the same for all birds, making them the most arbitrary feature of the simulations. Constancy was necessary to avoid overwhelming the simulation procedure by fitting too many parameters. Also, threshold and bias are such closely related parameters that the available pigeon data did not permit a clear choice between them. Both threshold and bias control the amount of accumulation needed for the random walk to trigger a decision; either could be used to simulate good fits, and allowing both to vary produced a confusing mix of values. I chose to hold the thresholds constant and allow bias to vary. This follows a common assumption that response thresholds are stable within sessions. Bias provided a parameter that could vary from trial to trial, as required by these fits. This arrangement also allows thresholds to accommodate variables that were not tested here, such as motivation and arousal.

The simulation process

Equations written in MATLAB (2007) simulated the RWP processes described in the text and sketched in Figure 2. Simulation was done separately for each bird (except for the search experiment) and for every pairing of the independent variables. Simulation started with parameter values “guessed” from exploratory work. For each parameter, some values larger and smaller than the “guess” were added, resulting in a list that spanned a short range. The lists from the various parameters were embedded in large multidimensional matrix, with one dimension for each parameter. The simulation program then ran 25000 simulated trials with each combination of values in the parameter matrix. It collected the RTs from each run into a distribution, and each distribution was compared with the corresponding empirical distribution by computing variance accounted for (ω^2). The parameter combination giving the maximum ω^2 was selected, with the restriction that no value for a given parameter could be the largest or smallest of the tested values for that parameter. If any parameter failed this test, a new list was generated, a new matrix was filled, and a new run commenced. The parameter values that emerged from this process controlled a final run that produced fitting distributions such as those shown by lines in Figures 1 and 4.

Figure 1.

Mean RT distributions for the stimuli in the reference experiment. At the top are distributions derived from responses to the green spot; responses to this spot always brought reward. Below are distributions from the rewarded stimulus, Red S+, and unrewarded RedS2 and RedS3. Reward for RedS+ varied between 100% (black dots) and 15%, (white dots). Dots represent data, lines the fits provided by RWP.

Figure 2.

Sketch of the essential aspects of the RWP model. A stimulus is presented at time = 0. Each irregular line shows the random walk for a single trial. This line represents the accumulation of variable bits of information arriving at successive moments in time; when this accumulation reaches the upper threshold, a response occurs. If the accumulation reaches the lower threshold, no response occurs (or, in two-choice situation, an alternative response occurs.) Total RT also includes a “residual” time from events irrelevant to decision. Decision-related parameters are named at left.

Figure 3.

Mean across birds of the step-size and bias parameters that gave the data fits shown in Figure 2. Top panel: Quite different step values were needed to fit data from the different red stimuli, but these values were almost the same for the two reward percentages. Bottom panel: Quite different bias values were needed to fit the data from the different reward conditions. These hardly differed among the different stimuli.

Figure 4.

Mean RT distributions from the search experiment. The search target was an “A” on some trials and an “E” on others, but data appear for “A” trials only. The dots are data, the lines are fits by RWP. At the top are curves resulting when the target “A” appeared alone on the screen. The other panels show the results when distractor letters were on the screen with the “A.” During the “baseline” condition (black dots) reward followed pecks to both “A” and “E” on 10% of trials. During the “contrast” condition reward for “A” stayed at 10% and reward for “E” jumped to 30%. Note especially that the baseline-contrast difference was about the same, regardless of the distractor condition (see text).