Behavioural Processes 85 (2010) 209–214

Contents lists available at ScienceDirect

Behavioural Processes

journal homepage: www.elsevier.com/locate/behavproc

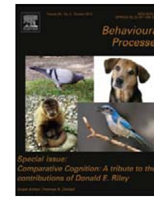

A random-walk interpretation of incentive effects in visual discrimination☆

Donald S. Blough

Department of Psychology, Box 1853, Brown University, Providence, RI 02912, United States

ARTICLE INFO

Article history:
Received 15 April 2010
Received in revised form 29 June 2010
Accepted 29 July 2010

Keywords:
Discrimination
Incentive
Random-walk
Pigeon

ABSTRACT

Previous studies have shown that changes in stimulus discriminability and changes in reward density affect pigeon reaction-time (RT) distributions in different ways. 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.

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 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 s 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.

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

E-mail address: Donald_Blough@brown.edu.

0376-6357/$ – see front matter © 2010 Elsevier B.V. All rights reserved.
doi:10.1016/j.beproc.2010.07.013

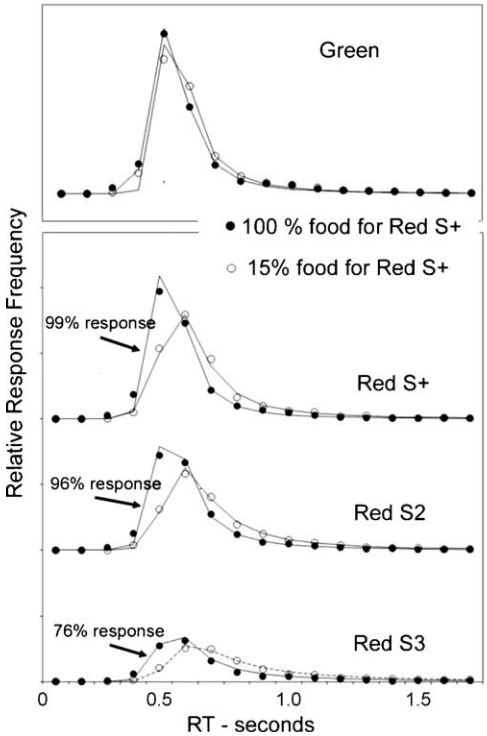

Fig. 1. Mean RT distributions for the stimuli in the reference experiment. At the top are distributions derived from responses to the green spot; responses to this spot always brought reward. Below are distributions from the rewarded stimulus, red S+, and unrewarded red S2 and red S3. Reward for red S+ varied between 100% (black dots) and 15% (white dots). Dots represent data, lines the fits provided by RWP.

Fig. 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 Fig. 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.

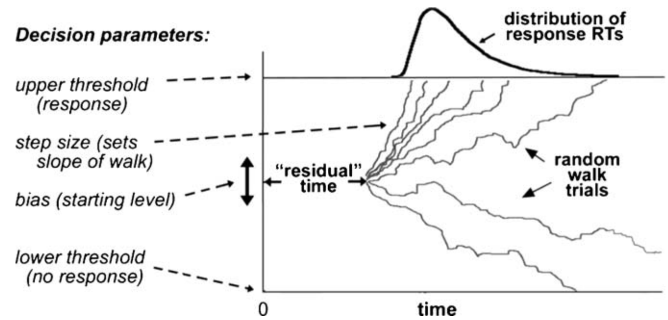

Fig. 2. Sketch of the essential aspects of the RWP model. A stimulus is presented at time = 0. Each irregular line shows the random walk for a single trial. This line represents the accumulation of variable bits of information arriving a successive moments in time; when this accumulation reaches the 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.

Fig. 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 Fig. 2 suggests an RT frequency distribution that might be generated by many trials with a positive stimulus. Additional details about simulation appear in Appendix A; for a full account see Blough (in press).

A computer program implemented the process sketched in Fig. 2 to simulate RTs and find theoretical distributions that best-fit the data distributions. (See below, Appendix A and Blough, in press.) The resulting best-fit distributions appear as lines in Fig. 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. Fig. 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