Shifting Differentiation and its Implications for the Response Concept[1]

Vicki L. Lee

Monash University

ABSTRACT

This study extended a past experiment on shifting differential reinforcement to a simple game-like task with human participants and followed through the distinction between responses and deeds in analyzing the resulting data. 6 college students worked on a computer-presented task in which button depressions changed a display of icons, provided the cursor was within the target location on the computer screen. The target location changed in a predictable direction but an unpredictable distance after every 15 icon changes. Data analysis related changes in the ratio of icon changes relative to button depressions to (a) cursor relocations after a failure to change the icons, (b) location repetitions after a success, (c) relocations after a failure that moved towards the target location, and (d) distances traversed by relocations. These components of successful performance became increasingly present as ratios of icon changes relative to button depressions increased. The paper concludes by relating an apparently anomalous aspect of the findings to the distinction between responses and deeds.
Key words: response concept, deeds, differential reinforcement, operant conditioning, human operant research.

RESUMEN

Este estudio extiende un experimento anterior sobre reforzamiento diferencial cambiante en una tarea de juego sencilla con sujetos humanos analizando los resultados a través de la distinción entre respuestas y acciones al analizar los resultados. Seis estudiantes universitarios participaron en una tarea presentada por ordenador en la que la presión de un botón cambiaba una serie de iconos, siempre y cuando el cursor estuviera localizado en una zona determinada de la pantalla. La localización de esta zona diana cambiaba en una dirección predecible, pero con una distancia impredecible, cada 15 cambios de iconos. En el análisis de los datos se encontró que la razón entre cambios de iconos y opresión del botón estaba relacionada con: (a) las reubicaciones del cursor tras un fallo en cambiar los iconos, (b) la repetición de la ubicación tras un éxito, (c) la reubicación tras un fallo que se movía hacia la localización objetivo y (d) las distancias recorridas en las reubicaciones. Estos componentes de ejecución exitosa llegaron a incrementarse en la forma de ratios de cambios en el icono relacionados con el aumento de la presión en el botón. El artículo concluye relacionando aparentemente un aspecto anómalo de los resultados a la distinción entre respuestas y acciones.
Palabras clave: Concepto de respuesta, acciones, reforzamiento diferencial, condicionamiento operante, investigación operante en humanos.

As Galbicka (1988) emphasized, differential reinforcement is central to operant conditioning research: «Response differentiation is operant conditioning, and vice versa» (p. 343). In differential reinforcement, responses that meet the experimenter's criterion have the scheduled effect (e.g., access to food) and increase in relative frequency as a result. Other responses, which do not meet the criterion, decrease in frequency relative to the reinforced responses. Differentiation is this change in relative frequencies.

An experimenter can repeat differential reinforcement across two or more related criteria. For instance, in successive differentiation, successive criteria specify increasingly narrow subclasses of the original class of behavior (e.g., Herrick, 1964). As a second example, in shaping, successive criteria specify classes that have some overlap with each other but increasingly less overlap with the original class (and eventually perhaps no overlap). Examples of research on shaping have included training animals to move to a location not previously occupied by them (Pear & Legris, 1987), to deposit ball bearings down a hole (Midgley, Lea, & Kirby, 1989), to jump onto a platform and extend their noses downward over the edge until they lost their balance and fell (Rasey & Iversen, 1993), and to press a bar (Stokes & Balsam, 1991).

An experiment by Eckerman, Heinz, Stern, and Kowlowitz (1980) suggests a third type of repeated differential reinforcement. Pecks by pigeons on the two right-most keys of a row of 20 keys produced food. After many sessions, the experimenter changed the keys at which pecks produced food. The food-effective keys were shifted from the right edge of the row across various intermediate keys to the left edge, then back to the right edge via various intermediate keys, and so on. At each repetition of differential reinforcement, the new food-effective keys were one, two, or three keys from the previous location. This change in the reinforced location followed a varying number of reinforcers for pecks at the currently-reinforced location.

Eckerman et al. (1980) interpreted their experiment as concerning shaping. However, Midgley et al. (1989) doubted this interpretation on three grounds. First, after the first horizontal sweep across the keys, successive locations were increasingly likely to have a previous reinforcement history. In contrast, with shaping, newly shaped acts have no such history. Second, in Eckerman et al. (1980), locations were under continuous reinforcement or extinction. In shaping, intermittent reinforcement commonly precedes extinction. Third, no location was a precursor to, or a component of, an earlier location. In shaping, earlier reinforced acts are precursors to, or constituents of, each newly-reinforced act (e.g., rearing anywhere as a precursor to rearing over the bar).

The experiment by Eckerman et al. (1980) might not demonstrate shaping but nonetheless does demonstrate a type of repeated differential reinforcement. This type of repeated differential reinforcement might be characterized generically as follows. First, there are *n* classes (of locations, forces, latencies, etc.). Second, during any one part of an experiment, only one of these classes is reinforced. Third, the order in which the classes are reinforced follows a pattern specified by the experimenter (e.g., a three-key shift after 100 reinforcers at the currently-reinforced location, then a two-key shift after 100 reinforcers, then a three key shift followed by a two key shift, and so on.). For convenience in the following discussion, differential reinforcement with these

characteristics is referred to as *shifting differential reinforcement*.

The present experiment extended shifting differential reinforcement to a simple game-like task with human participants. The experiment was conducted to describe what participants do under shifting differential reinforcement. Data analysis in Eckerman et al. (1980) focussed on histograms representing changes in the central tendency and dispersion of peck locations. The need for an account *of what participants do* under shifting differential reinforcement is suggested by arguments which question the conventional interpretation of differential reinforcement. These arguments suggest that data from experiments on differential reinforcement warrant conclusions, not about distributions of responses, but instead only about achievements (Gilbert, 1978), instances of the impact of an organism (Jacobs et al., 1988), or the things an individual gets done in the sense of events to which the individual's physical efforts contribute (Lee, 1995, 1996, 1999). The events under discussion here are often referred to in conversational English as deeds, where "deed" means something finished, completed, done, or brought about by someone. Deeds are the critical effects discussed by White (e.g., White, 1980; White & Liberty, 1976) and the functions discussed by Carr (1993). A score increase contingent on a button depression is as much a deed as is the button depression. They are both events (i.e., changes in a state of something) to which the individual's physical efforts (and much else) contribute. This interpretation makes no sense from a response-based perspective, which treats, for instance, (a) a button depression as an event that defines a class of button pressing responses, where the term *response* means an activity of the body and (b) a score increase as an environmental variable that might control responses. In contrast, a deed-based account of the same data (e.g., records of button depressions, score increases, and other events to which a participant contributed) finds only deeds and the relations among them. The question implied by a deed-based account concerns specifically what participants do (i.e., get done, achieve, accomplish) during an experiment rather than what happens to the variables (locations, latencies, durations, etc.) that are conventionally interpreted as the properties of responses.

A way forward is suggested by comparing shifting differential reinforcement with shaping. First, the eventual success of shaping is indicated by reliable occurrences of a target performance. Success during shaping is indicated by maintenance of a reasonably high proportion of reinforced performances, despite repeated changes in the reinforcement criterion. In shifting differential reinforcement, no single performance is established through a series of approximations as in shaping. If the pattern specified by the experimenter is repeated during an experiment, then it seems likely that performance might increasingly correspond to the pattern. If so, then increased correspondence would amount to increasingly efficient performance, such that the success of shifting differential reinforcement might be indicated by a relatively high ratio of success to effort (e.g., high ratio of food access events relative to pecks). Second, in shaping, some reinforced approximations can be conceptualized as components of the target performance. Stokes and Balsam (1991) presented relevant data. Persistence of reinforced approximations, such as rearing, was demonstrated in rats' final bar-pressing performances. Stokes and Balsam interpreted this finding as suggesting that shaping combines existing units (e.g., grasp, rear, body over bar) and directs them towards target locations (e.g.,

bar). It might be asked whether successful performance under shifting differential reinforcement has identifiable components that when combined increase the success-to-effort ratio.

METHOD

Participants

An advertisement placed in the student employment office invited individuals to phone the experimenter if they were interested in participating in an experiment on how people learn. Ten months after five college students had completed the experiment, an advertisement was placed to obtain a sixth participant. In addition to the sixth participant, a previous participant (E) replied. The participant was invited to complete the experiment, without the experimenter (or participant) mentioning past participation. An effort was then made to locate other past participants to take advantage of this unplanned opportunity to repeat the experiment within-participant after a break. Two participants agreed to return for another session. They were given the same information as before (see below). This resulted in an additional set of data for three participants either 10 (E, E2) or 12 (B, B2; D, D2) months later.

Setting and Apparatus

Figure 1. Computer screen showing the icons (specifically a non-matching set), score display, top white area containing the icons and score display, and the grey area that contained the columns (not delineated visually).

The experiment was conducted in an office furnished with a table, chairs, empty

bookshelves, and a Macintosh Centris computer. A mouse and a keyboard were attached to the computer. Software was written in FutureBASIC. Timing was in ticks (1/60th s) which was the unit of timing returned by the operating system when events were detected and recorded.

Five icons (i.e., 32 x 32 pixel bitmaps) extended horizontally from the top left edge of the computer screen (Figure 1). Ten pixels of white space separated each icon. A right-justified score display extended from the top right edge. Beginning at 60 pixels under the top of the window, the window was painted grey. The grey area was partitioned into 16 columns that were not visually differentiated. Participants could see only a continuous grey area with the white area containing the icons and the score above it. Each column was 40 pixels wide and extended vertically the height of the grey area. Columns did not extend into the white area. No instructions were given about the columns.

Procedure

Preliminaries. The participant read and signed a consent form which included an explanatory statement. Table 1 shows the main section of the explanatory statement. Other sections noted that the participant could withdraw at any time. Participants were

Table 1. Main Section of Explanatory Statement

This research is about what people do when asked to complete computer-based tasks. You are asked to interact with tasks presented by the computer until the computer tells you to stop. The session is likely to take as much as an hour or slightly longer. People vary considerably in the time they take to finish these tasks. The tasks are not tests of your intelligence. They are not tests of your personal abilities. We are interested only in finding out what people do when they interact with these tasks. You might or might not be given instructions on how to earn points. If no instructions are given, it will be up to you to figure out what to do.

informed that payment was $2 per quarter hour and that the participant achieving the highest score would earn a $10 bonus.

The following instruction was centered on the computer screen: "You earn points by depressing the up arrow key AS SOON AS the icons match." The experimenter pointed to the instruction and read it. She pointed to the up arrow key on the keyboard as the key was mentioned. The experimenter asked the participant to read aloud a typed instruction taped to the table in front of the computer. It was the same as the instruction on the screen with the following addition: "There are no more instructions. It is up to you to figure out what to do."

The experimenter depressed the button to erase the instruction and used the

mouse to push the cursor over the target column. It was the left-most column for D, D2, E, E2, and F and the right-most column for A, B, B2, and C. The experimenter told the participant to depress the button and to watch the icons. When the icons changed, the experimenter asked the participant if he or she had seen the icons change. The experimenter repeated this procedure for another four icon changes, then left the room after saying, "It's up to you now. The computer will tell you when the session is over."

Contingencies. While the mouse button was up, the cursor shape was an outstretched hand. While the mouse button was down, the cursor was fist-shaped. A button release restored the cursor to its original shape. Changes in cursor shape were independent of cursor location.

If the cursor was at a location that was within the target column, each button depression changed the icon display. If the cursor was located outside the target column when the button was depressed, the icon display did not change. Target columns did not extend into the top area of the screen. Therefore, no mouse button depressions that occurred with the cursor located within the top area could change the icons. Icon changes were contingent on button *depressions*. Button *releases* had no effect.

When the icon display changed, either five matching (i.e., identical) or five non-matching icons appeared. Icons were selected randomly from a pool of 130 black-and-white icons. The icons showed faces, animals, symbols, and other objects. Each set of matching icons differed from the preceding set. Any icon shown in a set of nonmatching icons was not shown in the next set of icons.

After 15 not necessarily consecutive mouse button depressions with the cursor located over the target column (and therefore 15 icon changes), the computer defined a new target column. For each series of 15 changes in the icon display, an array of five ratios from 1:1 through 5:1 determined whether the change was to five matching rather than five non-matching icons. At the start of each series of 15 icon changes, one ratio was selected randomly and without replacement from the array. If the ratio 1:1 was selected, the first icon change was an icon match; if the ratio was 2:1, the second icon change was an icon match; and so on. After an icon match, a second ratio was selected randomly and without replacement from the array, and so on, until 15 icon changes, had occurred, five of them changes to *matching* icons and 10 to *nonmatching* icons.

If the participant depressed the up-arrow key while five matching icons were visible and before depressing the mouse button again, the computer added one point to the score and played a brief tune. If five matching icons were visible and the participant depressed the mouse button before depressing the up-arrow key, a subsequent depression of the up-arrow key deducted one point from the score and caused a squeaking sound (but not if the mouse button depression produced another icon match). If the icons were not matching, each depression of the up-arrow key deducted one point and caused the computer to make a squeaking sound.

Therefore, participants could maximize their score (a) by depressing the up-arrow key each time the icons matched before doing anything else and (b) by desisting from depressing the key otherwise. The contingency relating key depressions to score increases was included to increase the chance that participants would watch for the scheduled visual consequences of depressing the mouse button (i.e., a change in the

icon display). The alternative procedure of increasing the score automatically when the icons matched was not used because it did not *require* participants to watch for their effects on the icons.

Target columns. The initial target column for participants D, D2, E, E2, and F was the left-most column. This column remained the target until the participant had changed the icon display 15 (not necessarily consecutive) times. After the participant met this criterion, the computer defined a new target column. The participant then had to depress the button with the cursor located over this new target column and therefore change the icon display 15 times. Then the computer defined a new target location, and so on, until the target column had been shifted 60 times.

The experiment comprised 61 conditions (i.e., 60 computer-controlled shifts in the target location preceded by the initial condition which did not include a shift). When the 61st condition ended, the computer presented this message: "Thank you for helping with this research. The program has finished now. Please get your supervisor."

Successive target columns were defined as follows ahead of the experiment for all participants. Five shift sizes ranging from one to five columns were used. A shift size was the number of columns added to the current target column to obtain the new target column. A list containing the 120 permutations of the five shift sizes was constructed (2, 1, 3, 4, 5; 3, 1, 2, 4, 5; etc). An array of the five shift sizes was selected randomly and without replacement from this list of permutations. The first value in the array (e.g., 2) was added to the initial target column to define the next target column (e.g., column 3). The next value in the array (e.g., 1) was added to this target column to define the next target column, and so on, until all five shift sizes had been used exactly once.

This series of five shifts moved the target column from the left edge of the screen to the right edge. Excluding the initial target column, these five successive target columns were defined as one sweep across the computer screen.

After the first sweep, a second sweep commenced. The same procedure was used to define the five target columns in this next sweep, except the target columns moved from right to left.

After six sweeps, giving a total of 30 shifts, the same pattern of column shifts was repeated. This gave a total of 12 sweeps across the screen and, given that each sweep comprised five shifts, a total of 60 shifts of the target column.

The same procedure governed shifts in the target column for other participants except as follows. The first target column for A, B, B2, and C was the right-most column. The first sweep proceeded towards the left edge of the screen, the second sweep towards the right edge, and so on.

<center>RESULTS</center>

Horizontal Coordinates

Figure 2 shows the horizontal coordinates of the cursor plotted across successive mouse button depressions for each sweep. These data are in plots of the same size

Figure 2. Horizontal coordinates of the cursor across successive mouse button depressions in each sweep. Numerals at the top of each graph represent the proportion of successful mouse button depressions. Numerals 1 through 12 at the top of the figure represent the 12 sweeps. Values on the vertical axis of each graph represent the number of pixels from the left edge of the screen, so that zero on vertical axes designates the left edge and 640 pixels designates the right edge. Letters A through E2 in the top left corner of each row of graphs are participant identifiers. Each dot shows the cursor coordinate for one mouse button depression. The data are compressed into a rectangle of the same dimensions regardless of the number of mouse button depressions, thus giving a visual impression of the number of mouse button depressions required to achieve the 75 icon changes in each sweep.

across the 12 sweeps for each participant despite variation in the number of mouse button depressions. Scaling the horizontal axes relative to the largest number of mouse button depressions across all participants and all sweeps (4931 for participant A in the

fourth sweep) would have lost the detail visible in most of these graphs.

Numerals above each graph indicate the proportion of all mouse button depressions with the cursor in the target column and that therefore changed the icons (*success proportions*). To conserve space, numbers are not otherwise provided. For all but two sweeps, number of mouse button depressions is estimated by dividing 75 (i.e., 15 icon changes for each of the five phases within the sweep) by the displayed success proportion (e.g., for participant B in the first sweep, 75/.338 (222)). As noted before, in the third and ninth sweeps, the target column shifted in four rather than five conditions. The single condition in these sweeps where the column did not shift is excluded from the calculation of success proportions and of all subsequent variables. Therefore, the number of successes in the third or the ninth sweep for a given participant is estimated by dividing 60 (i.e., 15 icon changes for each of the four conditions in the sweep where the target column *had* shifted) by the corrected success proportion given in the title of the corresponding graph. The density of the data points in each plot is a rough guide to the ratio of unsuccessful to successful mouse button depressions (e.g., compare A's first and last sweeps).

Figure 2 suggests a recurring pattern consisting of the emergence and subsequent persistence of consecutive repetitions of the same horizontal cursor location alternating with series of cursor relocations. A cursor relocation was defined as a difference in the horizontal coordinate of the cursor at the moment of a mouse button depression relative to the preceding mouse button depression. For example, if the horizontal coordinate was 60 pixels from the left edge of the computer screen at button depression t and something other than 60 pixels from the left edge at button depression $t+1$, then a cursor relocation was said to have occurred between the two successive button depressions. A cursor repetition was the absence of a difference in the horizontal coordinate of the cursor location between any two successive mouse button depressions.

The alternation between series of cursor location repetitions and cursor relocations is clear in many graphs (A from the 10th sweep, B from the second sweep, C from the ninth sweep, D from the fifth sweep, E from the third sweep, F from the third sweep, B2 from the fifth sweep, D2 from the sixth sweep, and E2 from the first sweep). Horizontal parts of the data path imply cursor location repetitions. Deviations from a horizontal path imply cursor relocations. In the sweeps listed above, the data path suggests many relocations in an internally-consistent direction. For example, the 12th sweep for A commences with horizontal relocations towards the right edge of the screen followed by a series of repetitions, followed by further horizontal relocations in the same direction, and so on, until cursor locations are at or near the right edge of the screen. Other graphs indicate no alternating pattern (e.g., A across the first seven sweeps, D2 in the first sweep). Yet other graphs suggest approximations to the relocate-repeat pattern (e.g., A for the eighth and ninth sweeps, F for the second sweep).

For E and E2, idiosyncratic patterns occurred in many right-left sweeps (e.g., fourth, sixth, eighth, 10th, and 12th sweeps for E). Target columns started near the right edge of the screen and shifted towards the left edge across successive conditions within the sweep. E (and E2) moved the cursor to the left edge of the screen at the start of these sweeps, then shifted it towards the right edge, until a series of mouse button

depressions occurred with the cursor in the same or similar location. Next, E moved the cursor in small steps until the right edge of the screen was reached. The cursor was then typically shifted back to the left edge before the next mouse button depression. This contrasts with cursor relocations during right-left sweeps for other participants (e.g., B's third, fifth, and seventh sweeps and also the ninth and 11th sweeps for B).

Components of Success Proportions

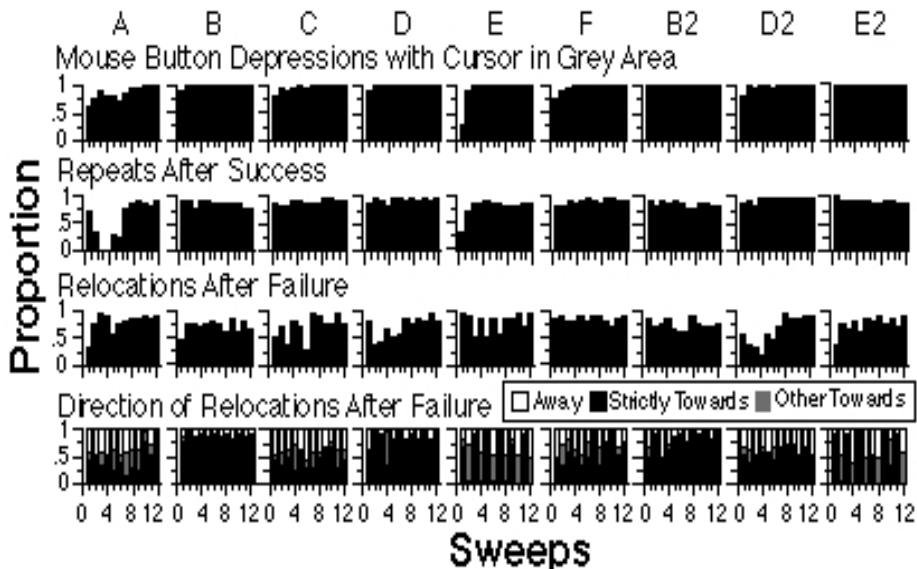Graphs in Figure 3 show five variables (two in the fourth row) which were

Figure 3. Mouse button depressions with cursor in grey area relative to all mouse button depressions (i.e., with cursor located anywhere on screen) across the 12 sweeps for the nine data sets (top row). Repetitions of the same cursor location after a mouse button depression that changed the icons relative to all mouse button depressions that changed the icons (second row). Relocations of the cursor after a mouse button depression that failed to change the icons relative to all mouse button depressions that failed (third row). Relocations of the cursor after a failure that moved away from the target column (white), strictly towards the target column (black), or towards but overshooting the target column or back to the target column after overshooting it (grey) (fourth row). A relocation away from the target column had no chance of reaching the target column regardless of distance. A relocation strictly towards the target column moved in the current sweep direction and ended either short of the target column or in it depending on distance. Relocations strictly towards the target column was a subclass of relocations towards which also included relocations in the current sweep direction that overshot the target column and relocations opposite to the sweep direction after the target column had been overshot.

evaluated as possible components of success proportions. Figure 4 relates the five variables to success proportions to integrate the variables in a way otherwise achieved tediously through detailed inspection of Figure 3.

Data in Figure 4 were extracted as follows. First, the 108 values of each variable from all 12 sweeps and all nine participants were combined. This produced six combined

variables (all success proportions, all grey-area proportions, etc.). Second, combined success-proportions were ranked. Third, an index to the ranked success-proportions was used to reorder each combined variable. For example, the first element in the reordered grey-area proportions was the grey-area proportion that occurred with the lowest success-proportion when success proportions from all conditions and all participants were considered together. Fourth, the quartile bins of the ordered success proportions were obtained. They were as follows: $\geq.02\leq.31$, $>.31\leq.42$, $>.42\leq.56$, $>.56\leq.81$. Fifth, for each quartile bin of success proportions, the distribution of each other variable across five bins of proportions ($\geq0\leq.2$; $>.2\leq.4$; . . ., $>.8\leq1.0$) was determined (Figure 4).

Considered across the experiment, .94 of all grey-area proportions (i.e., 102/108) exceeded .8, and .99 (i.e., 107/108) exceeded .6. Grey-area proportions equal to or less than .8 occurred only with success proportions in the first quartile. For succeed-then-repeat proportions, .87 exceeded .8, and .94 exceeded .6. Succeed-then-repeat proportions no greater than .4 occurred only with success proportions in the first quartile. In sum, for these two variables, high proportions were typical, high proportions occurred with
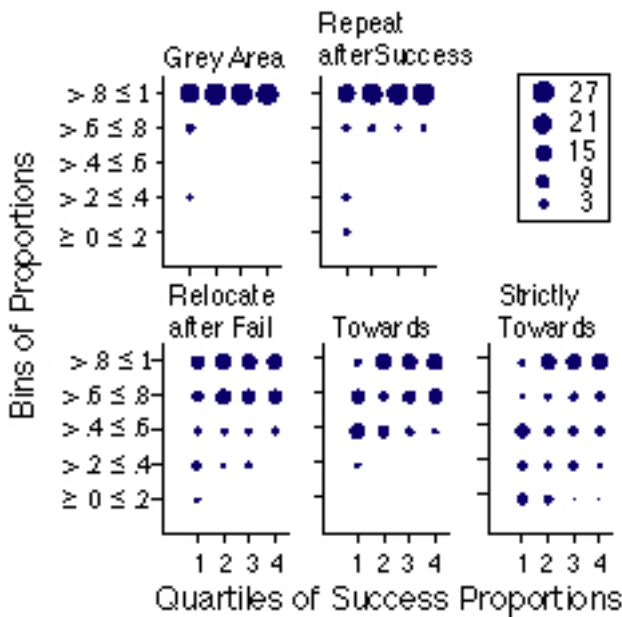

Figure 4. The number of cases of five variables (grey-area selections relative to all mouse button depressions, location repetitions relative to all successful mouse button depressions, relocations relative to all failures, relocations towards the target column relative to all relocations of the cursor after a failure, and relocations strictly towards the target column relative to all relocations after a failure) in each of five bins ($>=0<=.2$, $>.2<=.4$, etc.) for each quartile of the success proportions considered across all conditions and all participants.

low *and* high success proportions, and low proportions occurred only with low success

proportions. This suggests that high grey area proportions and high succeed-then-repeat proportions were necessary but not sufficient for high success proportions.

Figure 3 shows unusually low grey-area proportions persisted only for A beyond the first few sweeps, and unusually low succeed-then-repeat proportions occurred only for A and E early in the experiment. In these sweeps for these participants, success proportions were low (Figure 2).

Considered across the experiment, high proportions of relocating the cursor after a failure were less likely than the previous two variables. Across the experiment, .45 of the 108 cases of this variable were greater than .8, and .81 were greater than .6. Across the four quartiles of success proportions, relative frequencies of relocation proportions greater than .6 increased after the first quartile but subsequently changed only minimally (.63, .89, .85, .85 across the four quartiles). By implication, relocation proportions not greater than .6 were more common when success proportions were in their first quartile than in any other quartile. This result suggests that relocating the cursor after a failure was necessary for high success proportions but not sufficient.

Relocations of the cursor after a failure of a mouse button depression to change the icons could (a) approach the target column or (b) move away from it. Therefore, questions could be asked about relocation direction. Relocations towards the target column were defined as relocations that could end with the cursor within the target column conditional on shift distance.

Six subclasses were defined to support extraction of data. Relocations could move in the same direction as the current sweep from before the target column, *and* the cursor could end either (a) short of the target column, (b) in the target column, or (c) beyond the target column. In addition, relocations could move opposite to the current sweep direction (e.g., left to right in a right-left sweep) from beyond the target column, *and* the cursor could end either (a) short of the target column, (b) in the target column, or (c) on the other side of the target column. (An example is a shift to the left from column 9 when the target column was column 6 during a left-right sweep).

Considered across the experiment, relocations towards the target column with relative frequencies greater than .6 increased across the four quartiles of success proportions (.48, .70, .85, .96). A large increase occurred from the first to second quartile for the smaller subset of proportions greater than .8 with subsequently little change (.11, .56, .52, .56). Low proportions of this variable, and, by implication, higher proportions of relocations that moved away from the target column (i.e., the complement of the variable), were most common when success proportions were in their first quartile. For relocations after a failure in the current sweep direction that either reached or fell short of the target column (i.e., relocations strictly towards), relative frequencies of proportions greater than .8 increased substantially from the first to second quartiles of success proportions but not subsequently (.07, .44, .48, .52 across successive quartiles). Proportions greater than .6 increased in relative frequency across successive quartiles (.11, .52, .63, .78). This result suggests high success proportions depended on a large proportion of relocations after a failure moving towards the target column in the current sweep direction without overshooting it (i.e., relocations strictly towards the target column).

To indicate individual performances, the fourth graph in Figure 3 shows three

subclasses of cursor relocations after a failure. White areas represent relocations that moved away from the target column. Filled areas (i.e., black combined with grey) represent relocations after a failure towards the target column. Black areas represent relocations towards the target column in the current sweep direction that did not overshoot it (i.e., relocations strictly towards the target column). Grey areas represent all other cases of moving towards the target column; specifically, (a) moving in the correct direction but overshooting the target column and (b) moving in the direction opposite to the current direction of the target columns but towards the target column after having overshot it.

Variability persists across sweeps for E and E2, clarifying the idiosyncratic horizontal coordinates in Figure 2. Right-left sweeps for E and E2 (i.e., the even-numbered sweeps for this participant) typically had a low proportion of relocations that moved strictly towards the target column and a high proportion of other relocations including relocations away from the target column. In general across all participants, the higher the proportion of relocations after a failure strictly towards the target column, the clearer the pattern of alternating shifts and repetitions in a consistent direction seen in Figure 2.

. Distances traversed by the cursor when the participant relocated it are considered separately from other variables. Their analysis required some additional steps, as now explained.

The analysis included all relocations after a fail toward, though not necessarily strictly toward, were included. It seemed likely that a predominance of short distances would increase the number of unsuccessful mouse button depressions, therefore reducing success proportions. This relation was suggested empirically by inspecting B and B2's data in Figure 2 where many small relocations occurred with low success proportions. However, as indicated later in the Discussion, on further reflection, it became evident that this relation is logically necessary.

It also seemed likely that a predominance of larger distances (e.g., greater than 40 pixels, being the width of each column) would increase the chance that the new cursor location would bypass the target column, thus increasing the number of unsuccessful mouse button depressions. Distances up to 40 pixels *after a fail* (and therefore after a mouse button depression with the cursor not located in the target column) could not bypass the target column, because even if they began one pixel short of the target column, the new cursor location would be inside the column.

To avoid unnecessarily increasing unsuccessful mouse button depressions per sweep, participants had to desist from large proportions of small or large distances or both. To examine whether and how relocation distances changed during the experiment, three bins of relocation distances were defined (≤ 10 (small), $>10 \leq 40$ (medium), and >40 pixels (large)), and the first and second halves of the experiment were compared with respect to them.

To make this comparison, the proportion of relocations in each sweep with distances in each of the three bins was obtained for each participant. The results are plotted in Figure 5 in a way explained now with reference to participant A (i.e., the first two bars in each graph). White bars show the sums of the proportions of small (top

graph), medium (middle graph), and large (bottom graph) distances across the *first* six sweeps. The three proportions for each sweep necessarily add to 1.0. Therefore the sum of the sums of the three proportions across the six sweeps is 6.0. If distances were


Figure 5. Sum of the proportions of small (<10 pixels) relocation distances (top graph), medium (>10 pixels,<=40 pixels) distances (second graph), and large (>40 pixels) distances (third graph) for the first six sweeps (white bars) and last six sweeps (black bars) for the nine participants (successive pairs of bars in each graph).

evenly distributed across the three bins, the sum of the proportions in each of the three graphs would be 2.0. The black bars are organized comparably for the *last* six sweeps. The next two bars show the same data for participant B, and so on, across participants.

Figure 5 is summarized as follows. First, medium shift distances increased from the first to second half of the experiment for all participants except B. Increases were too small to be considered for B2 and E2. Sum of the medium shift distance proportions was typically greater than the expected 2.0 except for A and C in the first half of the experiment and for B and B2. Second, large relocation distances either rarely occurred (B, D, B2, E2) or decreased from the first to second half (A, C, E, F, D2). Third, small relocation distances increased for B, F, and B2 but otherwise changed minimally (A, E, E2) or decreased (C, D, D2).

Table 2 supplements this interpretation. It shows the number of relocations towards the target column in each of the three bins (i.e., ≤10 (small), >10≤40 (medium),

and >40 pixels (large)) that intersected with each quartile of success proportions considered across all participants and all sweeps (as in Figure 4 for the other variables). The table contains 27 values in total in each of the four quartiles as expected. B and B2's data account for 22 of the 28 cases of medians no more than 10 pixels. Of these 28 cases, 18 occurred with the lower two quartiles of success proportions. Relocation distances greater than 40 pixels were more likely to occur with the lowest success proportions. Relocation distances no more than 10 pixels were unlikely to occur with the highest success proportions.

Table 2 Frequency of Median Relocation Distances in Three Bins Intersecting with Quartiles of Success Proportions

Quartiles	Bins		
	≤10	>10≤40	>40
Q1	2	14	11
Q2	16	10	1
Q3	9	16	2
Q4	1	23	3

DISCUSSION

Considered together, Eckerman et al. (1980) and the present study correspond to two contrasting perspectives on the empirical facts of shifting differential reinforcement. First, shifting differential reinforcement can be conceptualized as a procedure that changes the central tendency and dispersion of distributions of response properties (location, latencies, etc.). Histograms showing these properties of the distributions then seem the obvious choice in analyzing the data, as in Eckerman et al. Second, shifting differential reinforcement can be conceptualized in terms of success-to-effort ratios (e.g., icon changes relative to mouse button depressions) and component-composite relations (e.g., cursor locations always over the grey area and never over the white area as a component of high success-to-effort ratios.

Rather than implying questions about the environmental control of responses, as does the first perspective, this second perspective implies questions only about the things a participant does and the relations among classes of those events. These events

are ordered in how-what or means-end relations, such that one thing is done (e.g., achieving a high success-to-effort ratio) by getting one or more other things done. Success-to-effort ratios and the components of these ratios are not salient when the research question is conceptualized in terms of the central tendency and dispersion of distributions of response properties inferred from the recorded variables (e.g., locations). The second perspective directs attention to aspects of shifting differential reinforcement that analyses of data using histograms do not make apparent.

The present analysis of data combined across all conditions and all participants (Figure 4) indicates differences in what participants did when success-to-effort ratios were relatively high. When achieving a high ratio of icon changes to mouse button depressions, participants kept the cursor in the grey area of the screen and out of the top area, more often repeated (than changed) the cursor location after an icon change, more often shifted (rather than repeated) the cursor location after a failure of a mouse button depression to change the icons, more often shifted the cursor towards the target column (rather than away from it), and more often made the shifts relatively short (rather than long) in distance. A logical analysis of success-to-effort ratios as a composite comprising various components suggested that the obtained differentiations (grey area locations rather than any screen location, shift after a failure to change the icons rather than a repeat, etc.) were *necessary* to achieve an increased ratio of success (icon changes) to effort (mouse button depressions).

The logical analysis is as follows. Mouse button depressions with the cursor over the top area were necessarily unsuccessful because no target column extended into the top area. Relocations after successes #1 through #14 were necessarily successful if they ended with the cursor within the column but were unsuccessful otherwise. Repetitions after successes #1 through #14 were necessarily successful because the locations of the cursor and the target column had not changed since the previous success. Repetitions after successes #15 were necessarily unsuccessful, but so were relocations unless they moved a sufficient distance in the current sweep direction. Repetitions after a failure were necessarily unsuccessful because the cursor had not been relocated since the previous failure. Relocations after a failure away from the target column were necessarily always unsuccessful, by definition. Relocations towards the target column could be successful, depending on (a) relocation distance and (b) distance between the initial cursor location and the new target column. Relocations that moved in the correct direction and did not overshoot the target column had a higher chance of being successful, again depending on relocation distance and initial cursor location. Experimenter-defined columns were 40 pixels wide. Therefore, cursor relocations towards the target column of about 40 pixels had (a) a high chance of reaching the target column with relatively few mouse button depressions and (b) a low chance of missing the column by overshooting it. Smaller distances unnecessarily increased the number of unsuccessful mouse button depressions and therefore decreased success proportions.

The logical analysis implies that participants could increase the ratio of success to effort only by engaging relatively more frequently in the components suggested above (keeping the cursor out of the top area, repeating the location after an icon

change, etc.). Conversely, the logical analysis implies that by engaging increasingly frequently in those components, participants necessarily increased the success-to-effort ratio (which can be thought of as a composite of the components). It appears that shifting differential reinforcement can be thought of in terms of composites and components, both of which can be thought of as things a participant achieves. The experimenter effectively assigns the participant a task to complete (earning a high score by the end of the experiment, consuming a meal by the end of the experiment, etc.). The task can be completed only in a cumulative manner across time (e.g., successive score increases). The empirical phenomena observed during the performance of such a task consist of variation within and across participants both in (a) success-to-effort ratios (e.g., icon changes relative to mouse button depressions), which can be considered composites, and (b) the relative absence vs. presence of the components of these composites (repeat after a success, shift after a fail, etc.). The present emphasis on composites and components is consistent with Stokes and Balsam's (1991) conclusion about shaping as likewise involving the combination of components into a composite.

RESPONSES OR DEEDS

As already indicated, differential reinforcement is conventionally interpreted using the response concept. To maintain consistency with the response concept, a button depression is conventionally interpreted as an environmental event used by the experimenter to define the class of button pressing responses. An event such as an icon change or a score increase is conventionally interpreted as an environmental event that might control the responses. However, if the events are interpreted literally (i.e., as changes in the state of something to which the participant's physical efforts and much else contributed), they do not differ in logical status from button depressions. All these events are deeds in the sense explained earlier.

The conventional interpretation of differential reinforcement presumes it is responses that are differentiated; specifically, by environmental variables. A consequence of this presumption is that the present experiment and its results are likely to seem anomalous. The recorded locations in the present experiment were not of the locations of classes of bodily activities defined by environmental events such as button depressions. Rather, each recorded location was the location of an object (i.e., the cursor), where the word *object* means any relatively enduring item. One of the things done by participants was to relocate the cursor (i.e., change one of the cursor's properties). These events (i.e., the cursor relocations) had various properties (namely, occasion, direction, and distance) that changed and narrowed in variability in this experiment. It was these properties of events which were themselves changes in a property (i.e., location) of an object (i.e., cursor) that were differentiated in this experiment. The fact that the data represented events at objects together with a literal interpretation of what was differentiated is problematic. It might prompt the objection that this experiment is not genuinely a differential reinforcement experiment but is instead only an analogy. However, an alternative view, as argued next, is that the conventional interpretation of past experiments as demonstrating the differentiation *of responses* has been mistaken.

Past experiments on differentiation of locations were interpreted as concerning response (rather than object) locations (Antonitis, 1951; Boren, Moerschbaecher, & Whyte, 1978; Crossman & Nichols, 1981; Eckerman et al., 1980; Eckerman & Lanson, 1969; Eckerman & Vreeland, 1973; Ferraro & Branch, 1968; Moerschbaecher, Thompson, & Thomas, 1979). For example, Boren et al. (1978) interpreted their data in terms of "variability in response location" (p. 66), and Ferraro and Branch (1968) as representing "key-peck response locations" (p. 1025). The underlying assumption is that location, like other variables traditionally classified as response dimensions (duration, interresponse times, etc.), is a property of a response rather than of an object. However, if the procedures of past experiments are examined, it turns out that the recorded locations are accurately conceptualized as the locations of objects. For example, in Eckerman et al. (1988), the recorded locations were the panels affected by the pigeon. The implied locations were the locations of the pigeon's body relative to the panels. The pigeon depressed a panel (i.e., one thing done) *and* changed the location of its body relative to the panels (i.e., another thing done). In Eckerman and Vreeland (1973), the recorded locations were of a written "x" on a page, and the implied locations were of the participant's hand. The changes to which the participant's physical efforts contributed were (a) to the location of his or her hand which grasped the pen and therefore also in the location of the pen relative to the page and (b) to the presence vs. absence of an "x" on the page. In Crossman and Nichols (1981), recorded locations were of keys depressed by the pigeon, and implied locations were locations of the pigeon relative to the keys. Changes to which the physical efforts of the participants' bodies contributed were displacements of the key and relocations of their own bodies relative to the keys. In Pear and Legris (1987), recorded locations were of the pigeon's head, and changes effected by the pigeon's physical efforts were to the location of its head relative to the target sphere. In Antonitis (1951), recorded locations were positions of a rat's nose (and therefore of the rat) along a slot, which were interpreted as locations of "nose-thrusting responses" (p. 275) or of "responses made by the nose" (p. 273).

Data from these past experiments comprised records of properties of objects (e.g., location of head) or of changes in various states of objects (e.g., changes of the location of the nose from outside to inside the slot). In this respect, the present experiment, with its emphasis on the locations of the cursor at the moment of each mouse button depression, is not anomalous. The difference between the present experiment and past experiments on the differential reinforcement of location is not in the type of events recorded. The difference is in how those events have been interpreted. The implied question is why, for instance, locations of the nose should not be interpreted literally as the locations of the nose instead of being interpreted as the locations of "nose thrusting responses" (Antonitis, 1951, p. 274). That is, why were past data, which have *represented* properties of objects and changes in various states of objects, *interpreted* as if they represented responses of a participant's body that must be distinguished from the participant's environment? The reason is because acceptance of the response concept has made it essential to classify all the observed events as belonging either to the organism (as responses) or to the environment (as stimuli). Contrary to that, the present literal interpretation of the data suggests that all these events are at the same time of

the organism *and* of the environment: They are events that have the physical efforts of the participant's body and much else as their constituents. For example, a button depression (i.e., that event at the button) cannot occur without the participant's physical efforts, the button, and much else. Regardless of interpretation, the central problem facing further empirical research is to make available for observation more of the details of what participants do under differential reinforcement. The results of the present study suggest that adopting a literal interpretation of the recorded events might support an increasingly more detailed and integrated account of (a) what participants do (i.e., achieve, get done, accomplish) and (b) how they do it (which is more of what they do).

REFERENCES

Antonitis, J. J. (1951). Response variability in the white rat during conditioning, extinction, and reconditioning. *Journal of Experimental Psychology, 42*, 273-281.

Boren, J. J., Moerschbaecher, J. M., & Whyte, A. A. (1978). Variability of response location on fixed-ratio and fixed-interval schedules of reinforcement. *Journal of the Experimental Analysis of Behavior, 30*, 63-67.

Carr, E. G. (1993). Behavior analysis is not ultimately about behavior. *The Behavior Analyst, 16*, 47-49.

Crossman, E. K., & Nichols, M. B. (1981). Response location as a function of reinforcement location and frequency. *Behaviour Analysis Letters, 1*, 207-215.

Eckerman, D. A., Hienz, R. D., Stern, S., & Kowlowitz, V. (1980). Shaping the location of a pigeon's peck: Effect of rate and size of shaping steps. *Journal of the Experimental Analysis of Behavior, 33*, 299-310.

Eckerman, D. A., & Lanson, R. N. (1969). Variability of response location for pigeons responding under continuous reinforcement, intermittent reinforcement, and extinction. *Journal of the Experimental Analysis of Behavior, 12*, 73-80.

Eckerman, D. A., & Vreeland, R. (1973). Response variability for humans receiving continuous, intermittent, or no positive experimenter feedback. *Bulletin of the Psychonomic Society, 2*, 297-299.

Ferraro, D. P., & Branch, K. H. (1968). Variability of response location during regular and partial reinforcement. *Psychological Reports, 23*, 1023-1031.

Galbicka, G. (1988). Differentiating the behavior of organisms. *Journal of the Experimental Analysis of Behavior, 50*, 343-354.

Gilbert, T. F. (1978). *Human competence: Engineering worthy performance*. New York: McGraw-Hill.

Herrick, R. M. (1964). The successive differentiation of a lever displacement response. *Journal of the Experimental Analysis of Behavior, 7*, 211-215.

Jacobs, W. J., Blackburn, J. R., Buttrick, M., Harpur, T. J., Kennedy, D., Mana, M. J., MacDonald, M. A., McPherson, L. M., Paul, D., & Pfaus, J. G. (1988). Observations. *Psychobiology, 16*, 3-19.

Lee, V. L. (1995). What is a psychological unit? *Behaviour Change, 12*, 98-108.

Lee, V. L. (1996). Things done. *Mexican Journal of Behavior Analysis, 22*, 137-158.

Lee, V. L. (1999). "Behavior" does not mean "behavior of the organism": Why conceptual revision is needed. *Behavior and Social Issues, 9*, 67-80.

Midgley, M., Lea, E. G., & Kirby, R. M. (1989). Algorithmic shaping and misbehavior in the acquisition of token deposit by rats. *Journal of the Experimental Analysis of Behavior, 52*, 27-40.

Moerschbaecher, J. M., Thompson, D. M., & Thomas, J. R. (1979). Effects of methamphetamine and scopolamine on variability of response location. *Journal of the Experimental Analysis of Behavior, 32*, 255-263.

Pear, J. J., & Legris, J. A. (1987). Shaping by automated tracking of an arbitrary operant response. *Journal of the Experimental Analysis of Behavior, 47*, 241-247.

Rasey, H. W., & Iversen, I. H. (1993). An experimental acquisition of maladaptive behavior by shaping. *Journal of Behavior Therapy and Experimental Psychiatry, 24*, 37-43.

Stokes, P. D., & Balsam, P. D. (1991). Effects of reinforcing preselected approximations on the topography of the rat's bar press. *Journal of the Experimental Analysis of Behavior, 55*, 213-231.

White, O. R. (1980). Adaptive performance objectives: Form versus function. In W. Sailor & B. Wilcox (Eds.), *Methods of instruction for severely handicapped students* (pp. 47-69). London: Paul H. Brookes.

White, O. R., & Liberty, K. A. (1976). Behavioral assessment and precise educational measurement. In N. G. Haring & R. L. Schiefelbusch & R. K. Hoyt (Eds.), *Teaching special children* (pp. 31-71). New York: McGraw-Hill.