Learning & Behavior
2008, 36 (4), 352-358
doi: 10.3758/LB.36.4.352

Preference for rewards that follow greater effort and greater delay

Jérôme Alessandri, Jean-Claude Darcheville, and Yvonne Delevoye-Turrell
University of Lille III, Lille, France

and

Thomas R. Zentall
University of Kentucky, Lexington, Kentucky

Humans prefer (conditioned) rewards that follow greater effort (Aronson & Mills, 1959). This phenomenon can be interpreted as evidence for cognitive dissonance (or as justification of effort) but may also result from (1) the contrast between the relatively greater effort and the signal for reinforcement or (2) the delay reduction signaled by the conditioned reinforcer. In the present study, we examined the effect of prior force and prior time to produce stimuli associated with equal reinforcement. As expected, pressing with greater force or for a longer time was less preferred than pressing with less force or for a shorter time. However, participants preferred the conditioned reinforcer that followed greater force and more time. Furthermore, participants preferred a long duration with no force requirement over a shorter duration with a high force requirement and, consistent with the contrast account but not with the delay reduction account, they preferred the conditioned stimulus that followed the less preferred, shorter duration, high-force event. Within-trial contrast provides a more parsimonious account than justification of effort, and a more complete account than delay reduction.

There is evidence from research with pigeons that a positive discriminative stimulus that follows a nonpreferred event, such as a greater number of responses, a longer delay, or a lower probability of reinforcement, is preferred over a different positive discriminative stimulus that follows a preferred event, such as fewer responses, a shorter delay, or a higher probability of reinforcement (see Zentall, Clement, Friedrich, & DiGian, 2006, for a review). This phenomenon has been studied using a procedure in which two kinds of training trial randomly alternate. For example, on one kind of trial, there is an initial preferred event (e.g., a single peck is required) followed by the presentation of a simultaneous discrimination (e.g., red correct, S+; yellow incorrect, S−). On the other kind of trial, there is an initial less preferred event (e.g., 20 pecks are required) followed by the presentation of a different simultaneous discrimination (e.g., green S+; blue S−). After sufficient experience with this procedure, when the pigeons are tested with the two positive discriminative stimuli from training (i.e., red and green) they typically show a preference for the S+ stimulus that followed the less preferred event in training (Clement, Feltus, Kaiser, & Zentall, 2000; Singer & Zentall, 2008).

This observed phenomenon is not explained by traditional theories that postulate that the value of a stimulus should depend only on its consequences (Hull, 1943; Skinner, 1938). Neither can these results be accounted for by recent theories, such as the contextual choice model (Grace, 1994), which predicts that the value of the terminal link stimulus should depend on the delay to primary reinforcement, independent of the context.

The apparent paradox in the inverse functional relationship between preference for the response requirement and preference for the positive discriminative stimulus that follows can be explained by the within-trial contrast model proposed by Clement et al. (2000). This model assumes that the value of a reinforcer (or the discriminative stimulus that predicts it) is judged relative to the value of the event that preceded it. More precisely, the value of a reinforcer depends on the relative improvement in value or the degree of positive contrast between the initial event and the reinforcer (or the stimulus that signals it).

When an analogous experiment has been conducted with humans, the similar paradoxical results have been explained in terms of the justification of effort (Aronson & Mills, 1959), a special case of cognitive dissonance (Festinger, 1957). According to cognitive dissonance theory, when there is a discrepancy between beliefs and behavior or experience, humans should attempt to reduce the discrepancy. The belief is that reinforcement should be commensurate with the effort required to obtain it, but experience does not support this belief. To reduce the dissonance, it is proposed that humans (and other animals; see Lawrence & Festinger, 1962) should give added value to the outcome that required greater effort to obtain it. But even if animals are capable of experiencing a discrepancy

T. R. Zentall, zentall@uky.edu

between expected rewards and the behavior needed to obtain them (dissonance), it is not clear that they have a need to resolve that dissonance. Instead, within-trial contrast offers a more parsimonious account of these effects, especially when the effect is found in animals.

Alternatively, it is possible that results obtained using the procedure developed by Clement et al. (2000)—and similar results—can be explained in terms of the relative reduction in the delay to reinforcement signaled by the discriminative stimulus (Fantino & Abarca, 1985). According to delay reduction theory, a stimulus will be preferred if it signals a greater reduction in the delay of reinforcement (relative to the absence of the stimulus) than does another stimulus. Furthermore, Fantino and Abarca proposed that the reduction in delay should be calculated relative to the total duration of the trial. The longer the trial duration, the closer to reinforcement the discriminative stimulus would be as a proportion of the duration of the trial and, thus, the greater the reduction in delay to reinforcement signaled by the S+. In the procedure used by Clement et al., although the two positive stimuli signaled the same delay to reinforcement, the one that followed the greater response requirement appeared relatively closer to reinforcement than the one that followed the smaller response requirement. Thus, it may not have been the contrast between the more aversive 20-peck requirement and the discriminative stimulus that followed that was responsible for the preference but the relative reduction in delay to reinforcement signaled by discriminative stimulus that followed the 20-response requirement, because the 20 pecks took longer to produce than the single peck did.

Delay reduction theory can also account for results reported by DiGian, Friedrich, and Zentall (2004), in which the introduction of a delay prior to presentation of the discriminative stimuli produced a preference for the S+ stimulus that followed. On the other hand, the results of an experiment by Friedrich, Clement, and Zentall (2005) would seem to pose a problem for delay reduction theory. Consistent with within-trial contrast, they found that an S+ stimulus that followed the absence of reinforcement was preferred over one that followed the presence of reinforcement. However, if one considers the interval between two successive reinforcers as the duration of importance, then the interreinforcement time on a trial with reinforcement would be shorter than the interreinforcement time on a trial without reinforcement, and the S+ stimulus that followed the absence of reinforcement would be a better relative predictor of the next reinforcement.

Recently, we have found that results similar to those reported by Clement et al. (2000) can be obtained with humans (Klein, Bhatt, & Zentall, 2005). When participants worked to obtain each of two positive discriminative stimuli, the one that they had to work harder to obtain was preferred over the one they had to work less hard to obtain.

The purpose of the present study was threefold. First, we sought to examine the preference of humans for low and high effort, or force responses (independent of differential time), and long or short delays (independent of differential force). Second, we sought to determine whether the nonpreferred force required to make a response and the nonpreferred duration of the response would lead to a preference for the discriminative stimuli that follow. Finally, we sought to determine whether delay reduction theory is sufficient to account for the results of this and similar experiments (see Singer, Berry, & Zentall, 2007). If a nonpreferred event occurs on a shorter duration trial, the contrast hypothesis and the delay reduction hypothesis make opposite predictions. If delay reduction is sufficient to account for these results, trial duration should determine the value of the conditioned reinforcer, and the S+ that follows the longer initial event should be preferred. According to the contrast hypothesis, if the force requirement is sufficiently nonpreferred, it should overcome the aversiveness of the delay to presentation of the discriminative stimuli and should result in a preference for the S+ that follows the less preferred, shorter duration, but greater effort event.

METHOD

Participants

The participants were 30 undergraduate students (10 male and 20 female) at the University of Lille III, all volunteers.

Apparatus

A Novatech Mini40 ATi force cell (Tatem Industrial Automation Ltd., Derby, U.K.) served to measure force. All participants were trained and tested with a program created with Labview 8.0 (National Instruments Corporation, Austin, TX), presented on a computer monitor.

Procedure

The experiment was completed in a single session of 40–45 min. The value of high force was defined separately for each participant. We asked each participant to press the force cell with the maximum force possible with his or her dominant hand. For each participant, we defined the high force as at least 50% of this maximum force. The actual high force used varied considerably from a low of 20 N to a high of 150 N. The low force was defined as no less than 2 N and no more than 10 N. The second variable studied was the duration of required force (1 or 5 sec).

Phase 1. The purpose of Phase 1 was to determine the participants' preference for the various combinations of force (high force vs. low force) and time (1 vs. 5 sec). A forced trial consisted of the presentation of a circle at the top or the bottom of the screen whose position depended on the force required (high or low) or on the duration required (short or long). One mouse click (with the nondominant hand) to the circle made it disappear and initiated the force/time response requirement (with the dominant hand). When force was manipulated, at the start of each trial, a message appeared on the screen, indicating that the participant should press the force cell either with low force or with high force (as determined earlier).

If the force criterion was met, the word "correct" appeared on the screen; if not, a message appeared: "You have not pressed with enough force, try again." For the low-force condition, if the participant pressed with a force greater than 10 N, a message appeared: "Be careful, you have pressed with too much force, try again." When this occurred, the trial began again and was repeated until the participant responded correctly. These correction trials occurred very seldom after the first or second trial of each type. In each condition, training consisted of one block of 10 forced trials, 5 of each trial type (randomly presented), followed by 5 free choice trials, then another block of 6 forced trials, 3 of each trial type, followed by 5 additional free choice trials.

When duration was manipulated, the participants were told that they would have to press the force cell for either a short time (1 sec) or a long time (5 sec), and they were given feedback that they had responded long enough or not long enough, similar to the feedback from the high- and low-force conditions. In addition, a temporal constraint was imposed on all trials in all conditions. On all trials, participants had to begin to respond within 1 sec after their response to the circle. Thus, the 1-sec response had to be completed within 2 sec after the response to the circle, and the 5-sec response had to be completed within 6 sec after the response to the circle. The participants were instructed to click on the circle either at the top or at the bottom of the screen. If the circle appeared at the top, they were instructed to immediately press on the cell with high force (for some participants, or low force for others) for a few seconds (for some participants, or briefly for others), until feedback was provided (see Figure 1, left and center columns). They were told that feedback would be provided. If the appropriate response was not made, the trial was repeated.

On choice trials, participants were presented with both circles and were told that they should choose according to their preference and respond with the force and duration that they used in training, but that they would no longer be given feedback following their choice (see right column of Figure 1). The design of the experiment involved various combinations of the two variables (force and duration; see Table 1 for the order of conditions for half of the participants; for the remaining participants, the order of conditions was reversed). The purpose of Condition 1 was to determine the preference for high or low force with duration held constant at the higher level. The purpose of Condition 2 was to determine the preference for long or short duration with force held constant at the higher level. The purpose of Condition 3 was to determine the preference for high or low force with duration held constant at the lower level. The purpose of Condition 4 was to determine the preference for long or short duration with force held constant at the lower level. The purpose of Condition 5 was to determine the preference for the combination of high force and long duration or of low force and short duration. We wanted to know whether combining the two presumably more aversive events would affect the preference. The purpose of Condition 6 was to determine whether pitting the high-force and low-duration condition against the low-force and high-duration condition would reduce the overall preference, or whether the values selected for either of those two dimensions would overshadow the other.

In each condition, training consisted of one block of five forced trials of each kind, intermixed, followed by five free choice trials, then another block of three forced trials of each kind followed by five additional free choice trials. Trials were self-paced. The next trial began immediately after the previous trial ended, with the appearance of the circle at the top or bottom of the screen. There was a 5-min break between Phases 1 and 2.

Phase 2. The purpose of Phase 2 was to determine the relationship between the most and least preferred response from Phase 1 and the positive discriminative stimulus (shape) that consistently followed. A preference score was calculated by summing all choices made of each response over the three conditions in which it appeared. Based on the preference scores for each of the four possible initial events obtained in Phase 1, the most preferred and least preferred responses from Phase 1 served as the two initial events used in Phase 2.

A procedure similar to that used by Clement et al. (2000) was used. Participants were told to click once on an initial red or blue rectangle presented at the middle top of the screen and then to immediately press the force cell with one of the forces they had learned in Phase 1. Participants were informed of the correspondence between the color of the rectangle and the required response. The blue rectangle indicated that the preferred response from Phase 1 was to be made, and the red rectangle indicated that the less preferred response from Phase 1 was to be made. Following performance of the response requirement, a simultaneous discrimination was presented involving a choice between two free-form line-drawn shapes (see Klein et al., 2005). A different pair of line-drawn shapes followed each of the two response requirements. Participants were also told that the choice of one shape would be followed by the feedback "correct" and the choice of the other by the feedback "wrong." The two kinds of trial alternated randomly, with the constraint that each type of trial occurred equally often in a block of four trials. The position of discriminative stimuli (left or right) was randomized, with the constraint that each stimulus occurred equally often on each side. Training continued until there were eight consecutive choices of the correct shapes (four correct choices involving each discrimination). The shape of the correct stimulus was counterbalanced among participants. The design of Phase 2 is presented in Figure 2.

Table 1
Order of Conditions (for Half of the Participants*) and Participants' Mean Preference for the Low Force or Shorter Duration Event With Standard Errors of the Mean

Condition	Force	Duration	Percentage of Choices	
			M	*SEM*
1	high–low	long	73.7	2.37
2	high	long–short	86	1.83
3	high–low	short	75.3	2.74
4	low	long–short	71.7	2.36
5	high–low	long–short	84	2.43
6	high–low	long–short	50	3.44

Note—High force was determined separately for each participant (see the text). Low force was 2–10 N. Long duration was 5 sec; short duration was 1 sec. In Condition 5, the trials involved high force for 5 sec, randomly alternated with low force for 1 sec. In Condition 6, the trials involved high force for 1 sec, randomly alternated with low force for 5 sec (shown is the preference for the high force for 1 sec). *The other half of the participants received the reverse order of conditions.

When the criterion was met, testing followed. Test trials involved a choice between the two positive discriminative stimuli from training. The participants were told that they should now choose the shape that they most preferred but that no feedback would be provided. The testing phase consisted of four test trials, two of which were initiated by the red rectangle and two of which were initiated by the blue rectangle (for each color, the corresponding response was required, as it was in training). The position of the two stimuli was randomized, with the constraint that each stimulus appeared equally often on each side. The trial type (response requirement) of the first test trial was counterbalanced among participants. Following the four test trials, there were an additional eight training trials followed by four more test trials (two with the blue rectangle, two with the red rectangle, as was the case for the initial stimulus). For counterbalancing purposes, the first test trial in the second set of test trials was always the trial type that had not been presented first in the first set of test trials. As in Phase 1, if the appropriate force was not produced, the trial was repeated. There was no break between Phases 2 and 3, other than to verbally tell the participants that there would no longer be a force requirement.

Phase 3. The procedure used in Phase 3 was similar to that of Phase 2, with the following exceptions: In an effort to directly test delay reduction theory, we replaced the response requirement that was preferred by most participants (low force, 1 sec) with a 10-sec delay (the participants were told that there would be no force required) and signaled it with a green rectangle. The participants were trained with these two initial events (the 10-sec delay and their least preferred response determined from Phase 1) as they had been in Phase 1 (see Figure 3). They were then tested for their initial stimulus preference, as they had been in Phase 1 (see Figure 4). They were then trained as they had been in Phase 2, but with four new shapes as discriminative stimuli (see Klein et al., 2005). In training, one of two new shapes was correct following the 10-sec delay and the other new shape was correct following the least preferred response (the high-force 5-sec response). When participants met the discrimination criterion, they were tested for their S+ stimulus preference, as they had been in Phase 2.

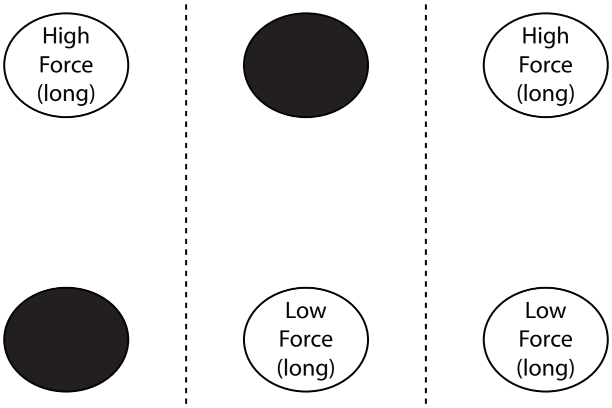

Figure 1. In Phase 1, each training trial started with the illumination of a circle at the top of the screen, indicating that one response would be required (left column, e.g., high force, long), or a circle at the bottom of the screen, indicating the other response would be required (center column, e.g., low force, long). On test trials, circles appeared at both the top and the bottom of the screen and participants could choose the response to be made (i.e., high force, long or low force, long).

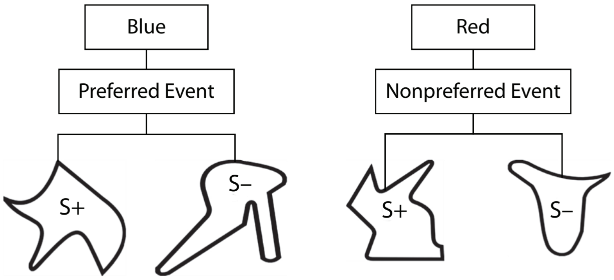

Figure 2. Phase 2 training trials. At the start of each training trial, a red or blue rectangle appeared at the top of the screen. The participants clicked on the rectangle and responded with the required force for the required time to produce a choice between two shapes. Choice of the S+ resulted in the appearance of the word "correct" on the screen. Choice of the S− resulted in the appearance of the word "wrong" on the screen.

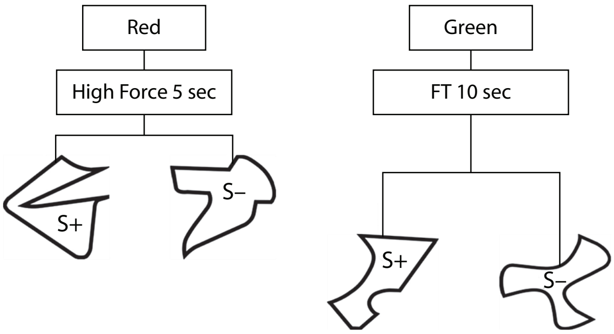

Figure 3. Phase 3 training trials. The preferred response requirement from Phase 2 (typically low force, 1 sec) was replaced by a 10-sec delay with no force response required. When the participants responded with the required force for the required time, they were given a choice between two new shapes. Choice of the S+ resulted in the appearance of the word "correct" on the screen. Choice of the S− resulted in the appearance of the word "wrong" on the screen.

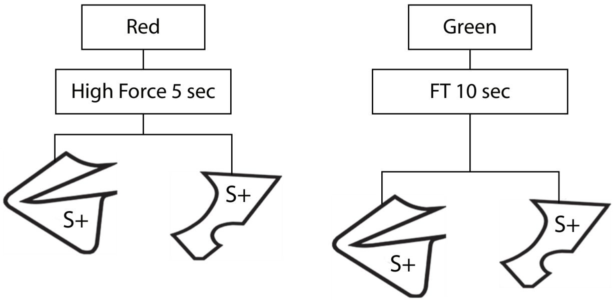

Figure 4. Phase 3 test trials. Participants were given a choice between the two S+ stimuli from Phase 3 training. Half of the test trials were initiated by a red rectangle and half by a blue rectangle.

RESULTS

Phase 1

The results of Phase 1 are represented in Table 1. A participant's preference for the four required responses was determined as follows: As each initial event appeared in three conditions (once with each of the other required responses), the number of choices of each required response was summed and divided by the number of test trials on which the initial event was presented (and multiplied by 100 in order to convert to a percentage preference). Overall, participants preferred the alternative requiring low force and shorter duration (77.0%), followed by the one requiring low force and long duration (53.5%), followed by the one requiring high force and short duration (50.7%), followed by the one requiring high force and long duration (18.8%).

Individual single-sample *t* tests were conducted on the overall choice of the preferred initial stimulus in each condition. In Condition 1, the analysis indicated that the low-force 5-sec response was preferred significantly over the high-force 5-sec response [$t(29) = 5.47$, $p < .01$, $d = 2.03$]. In Condition 2, the high-force 1-sec response was preferred significantly over the high-force 5-sec response [$t(29) = 10.77$, $p < .01$, $d = 4.00$]. In Condition 3, the low-force 1-sec response was preferred significantly over the high-force 1-sec response [$t(29) = 5.07$, $p < .01$, $d = 1.88$]. In Condition 4, the low-force 1-sec response was preferred significantly over the low-force 5-sec response [$t(29) = 5.02$, $p < .01$, $d = 1.86$]. In Condition 5, the low-force 1-sec response was preferred significantly over the high-force 5-sec response [$t(29) = 7.66$, $p < .01$, $d = 2.84$]. However, in Condition 6, the low-force 5-sec response was not preferred significantly over the high-force 1-sec response ($t < 1$). For all but 2 participants, the most preferred was the low-force 1-sec response, and the least preferred was the high-force 5-sec response.

Phase 2

On the first test trial involving choice of the two S+ stimuli, 20 participants chose the S+ stimulus that, in training, followed the nonpreferred response (66.7%). When the data were pooled over all eight of the test trials, a similar preference was found (62.9%, *SEM* = 4.17%). A single-sample *t* test was conducted on the overall S+ stimulus preference on test trials. The analysis indicated that the S+ stimulus preference associated with the less preferred initial response was significantly greater than chance [$t(29) = 3.09$, $p < .01$, $d = 1.15$].

Test trials were divided according to the initial stimulus on test trials to determine whether the required response on test trials affected the S+ preference. When the required response on test trials was the less preferred response, the participants chose the S+ stimulus associated with the less preferred response in training (65.0%, *SEM* = 5.48%), and when the initial event on test trials was the preferred event, they chose the S+ stimulus associated with the less preferred response in training (60.9%, *SEM* = 6.44%). A two-sample *t* test performed on the test data, as a function of the two levels of required response on test trials (nonpreferred vs. preferred response), yielded a nonsignificant effect of the required response on test trials ($t < 1$).

Phase 3

The initial preference assessment indicated that most participants (22/30) preferred the 10-sec initial event over the high-force 5-sec initial event (70.8%, *SEM* = 5.99%). This preference was significantly greater than chance, as indicated by a binomial test ($p < .01$).

Whichever required response was preferred, on the first test trial involving choice of the two S+ stimuli, participants showed a preference for the S+ stimulus that, in training, followed the nonpreferred response (66.7%). When the data were pooled over all eight of the test trials, preference for the S+ stimulus that, in training, followed the nonpreferred response was 64.2% (*SEM* = 5.17%). A single-sample *t* test performed on the overall choice of the S+ stimulus on test trials indicated that preference for the S+ stimulus that followed the less preferred required response in training was significantly greater than chance [$t(29) = 2.71$, $p = .01$, $d = 1.01$].

For the 22 participants who preferred the 10-sec initial event, on test trials involving the choice of the two S+ stimuli, most showed a preference for the S+ stimulus that, in training, followed the less preferred high-force 5-sec event. The mean preference for the S+ stimulus that in training followed the less preferred high-force 5-sec event was 64.2%. A single-sample *t* test was conducted on the overall choice of the S+ stimulus on test trials. The analysis indicated that choice of the S+ stimulus that followed the less preferred required response in training was significantly greater than chance [$t(21) = 2.12$, $p = .05$, $d = 0.93$].

According to the contrast hypothesis, preference for the discriminative stimulus should depend on the degree of aversiveness of the prior event. To test this hypothesis, a Pearson product–moment correlation was performed on the response preference scores (short vs. long duration and low vs. high force) and the positive stimulus preference scores for each participant. The analysis (a one-tailed *t* test) indicated that there was a significant negative correlation between degree of schedule preference and preference for the positive stimulus that followed the preferred response [$r = -.33$; $t(28) = 1.85$, $p < .05$, $d = 0.70$]. Thus, the degree to which a participant showed a preference for a required response predicted the test-trial preference for the S+ stimulus associated with the other required response.

Once again, test trials were divided according to the required response on test trials in order to determine whether the response affected the S+ preference. On test trials, the participants chose the S+ stimulus associated with the less preferred required response in training when the required response on test trials was the less preferred response (74.2%, *SEM* = 2.50%) and when the initial event on test trials was the preferred response (54.2%, *SEM* = 3.65%). A *t* test performed on the test data yielded a significant effect of the initial event on test trials [$t(58) = 2.54$, $p = .014$, $d = 0.67$]. Thus, although there was no need to do so, in Phase 3, participants showed a significant tendency to use the required response on test trials as a conditional stimulus.

DISCUSSION

The results of the present experiment indicate that humans (like other animals) prefer to make responses that require less force and are of shorter durations (Hull, 1943). More important, we found a preference for an S+ stimulus that followed the less preferred response in training over one that followed a more preferred response, using differential force and differential duration of response as the variables. Most of the participants preferred the low-force 1-sec response over the high-force 5-sec response, but consistent with the contrast account, they preferred the S+ stimulus that followed the high-force 5-sec response in training over the S+ that followed the low-force 1-sec response. These results support and extend results reported by Klein et al. (2005).

The results of Phase 2 test trials can be accounted for in terms of either delay reduction or within-trial contrast, because the least preferred response was also the one that took the longest to make. However, the results of Phase 3 support the contrast account, and delay reduction theory by itself, as originally formulated, has difficulty accounting for these results (see also Singer et al., 2007).

According to delay reduction theory, the value of a stimulus as a conditioned reinforcer should depend only on the degree to which it results in a reduction in the delay to reinforcement. Thus, the longer the duration of the required response, the closer the S+ should come to reinforcement, and the greater should be the reduction in delay to reinforcement signaled by the S+. Because delay is generally not preferred, delay reduction theory and the contrast hypothesis often make similar predictions. In Phase 3, however, we pitted delay reduction theory against the contrast account. In Phase 3, the required-response preference assessment indicated that most of the participants preferred to wait 10 sec rather than to respond to the force cell with high force for 5 sec. However, in agreement with the contrast hypothesis, on test trials, participants who initially preferred the 10-sec delay showed a significant preference for the S+ stimulus that followed the high-force 5-sec required response.

Delay reduction theory and the more recent hyperbolic value-added model proposed by Mazur (2001) both use time to reinforcement to predict schedule preference. Although the hyperbolic value-added model describes the schedule preference in terms of value that decreases according to a hyperbolic function, it, like delay reduction theory, bases that value on time—delay to reinforcement. Delay reduction theory is explicit in proposing that it is the relative amount of delay reduction that occurs upon presentation of the discriminative stimuli that predicts preference for the positive discriminative stimulus. The hyperbolic value-added model is not so explicit about the relative reduction in delay as a source of value, but it may be implied from the notion that choice behavior is determined by the amount of value added at the moment of transition from initial to terminal links.

On the other hand, if one were to view value more generally as affected by variables other than delay, in principle, both theories could account for the data from the present experiment. However, if such a theory were to be useful, those variables would have to be more explicitly stated, so that they could be distinguished from contrast and empirically tested.

The present results contribute to findings by Singer et al. (2007). Initially, they trained pigeons on two schedules of reinforcement that involved similar delays to reinforcement: a differential reinforcement of other behavior (DRO) schedule (not pecking for 20 sec was reinforced) and a modified fixed interval (FI) schedule (the first response after 20 sec was reinforced). Because early in training, the pigeons pecked when the DRO schedule was in effect, thus lengthening the interval, the FI schedule was modified on a trial-by-trial basis to match the duration of each preceding DRO trial. When schedule preference was assessed, one pigeon preferred the FI schedule, others preferred the DRO schedule, and still others were indifferent to the schedules. However, when the experimenters followed each of the schedules with a distinctive pair of discriminative stimuli, Singer et al. found, in keeping with the contrast hypothesis, that for every pigeon, the schedule preference predicted the pigeons' preference for the S+ stimulus that followed. Furthermore, there was a strong inverse correlation between the degree of schedule preference and the degree of preference for the positive stimulus that followed the preferred event.

A similar approach was taken by Singer and Zentall (2008). In one experiment, pigeons were trained with a fixed ratio (FR) schedule on some trials and a DRO schedule on others. Again, the trials were matched on an individual basis for duration. Following a test for schedule preference, each of the two schedules was followed by a distinctive pair of discriminative stimuli. Once again, the pigeons' schedule preference significantly predicted their preference for the S+ stimulus that followed.

Finally, Singer and Zentall (2008) trained pigeons with a differential response requirement (FR1 and FR30, where 1 and 30 refer to the number of responses required) similar to that used by Clement et al. (2000). For pigeons in one group, trial duration and number of pecks required were confounded, as they were for Clement et al. For pigeons in a second group, the additional time to complete the higher fixed ratio was added as a delay (in the absence of the stimulus) following the lower fixed ratio requirement. Once again, following a test for schedule preference, each of the two schedules was followed by a distinctive pair of discriminative stimuli, and once again, the pigeons' schedule preference significantly predicted their preference for the S+ stimulus that followed. Thus, it should be clear that this effect that we have referred to as within-trial contrast does not require a differential reduction in delay to reinforcement.

Note that not all studies with pigeons have been able to replicate the contrast effect first reported by Clement et al. (2000; see Arantes & Grace, 2008b; Vasconcelos, Urcuioli, & Lionello-DeNolf, 2007). Why that is so is not clear, but the large number of studies that have reported a significant within-trial contrast effect with relatively aversive prior events consisting of effort (many vs. few responses), delay (long vs. short), the absence of food (when food is presented on other trials), and food deprivation period (long vs. short) suggest that the effect is a real one (Clement et al., 2000; Clement & Zentall, 2002; DiGian et al., 2004; Friedrich et al., 2005; Friedrich & Zentall, 2004; Kacelnik & Marsh, 2002; Klein et al., 2005; Pompilio & Kacelnik, 2005; Pompilio, Kacelnik, & Behmer, 2006; Singer et al., 2007; Vasconcelos & Urcuioli, 2008).

Studies that have failed to replicate the effects of a relatively aversive initial event on the preference for the positive discriminative stimulus that follows are consistent with the contextual choice model of Grace (1994), who proposed that the prior event (or schedule) should not have an effect on the later schedule. The reason for this inconsistency in results is not clear, but it is likely that the values associated with certain variables play an important role in obtaining the effect (see Zentall, 2008). For example, Clement et al. recognized that it was not possible to determine how much training would be needed to develop an association between the initial response requirement and the discriminative stimuli that followed. For this reason, they gave their pigeons 20 sessions of training following the acquisition of the simultaneous discriminations. Friedrich and Zentall (2004) monitored the development of the preference for the S+ that followed the greater response requirement and found that a significant preference did not develop until the pigeons had had over 70 sessions of training. Similarly, Singer et al. (2007) found that a significant preference only developed after 40 sessions of training. Why Clement et al. (2000) were able to find the contrast effect with only 20 sessions of training is not clear, but the later research suggests that it may take considerably more training to establish the association. In this regard, it is interesting to note that in the present study, overtraining was not needed to observe the effect with human participants (see also Alessandri, Darcheville, & Zentall, 2008; Klein et al., 2005). At this point, why overtraining should be necessary with pigeons but not with humans is a matter of speculation. One possibility is that humans may simply acquire the associations between the response requirement and the positive discriminative stimulus that follows considerably faster than pigeons. In the case of humans, faster acquisition may result from the possibility that humans are able to verbalize the associations between the response requirement and the positive discriminative stimulus that follows.

A second study that has failed to report a contrast effect using the Clement et al. (2000) procedure included several pigeons that received extended training (Arantes & Grace, 2008b). The results of this study are a bit more problematic. However, it is possible that the earlier experience that these pigeons had with lean (variable interval) schedules of reinforcement may have inadvertently washed out the contrast effect. According to Zentall and Singer (2007), the basis of the within-trial contrast effect is the contrast between the less preferred schedule and the discriminative stimuli that follow. If, however, the pigeons had been accustomed to lean schedules of reinforcement, the FR20 schedule used in their experiment might not have been sufficiently aversive to produce contrast. On the other hand, according to Arantes and Grace (2008a), even lean variable interval schedules may be less aversive than richer fixed ratio schedules (see also Moore & Fantino, 1975). A test of the hypothesis that prior experience with lean schedules can reduce or eliminate the contrast effect from occurring will have to await further research.

Overall, the present results, together with earlier research with humans (adults and children), pigeons, and starlings, suggest that the within-trial contrast account provides a parsimonious explanation of the preference often found for a positive discriminative stimulus that is preceded by a relatively aversive event over one that is preceded by a less aversive event. Furthermore, although it has been argued that cognitive dissonance theory can be applied to animals in asocial contexts (Lawrence & Festinger, 1962), the present results, as well as the earlier results with pigeons, suggest that within-trial contrast provides a simpler account of the version of cognitive dissonance known as justification of effort (Aronson & Mills, 1959). These results suggest, furthermore, that contrast effects may contribute to other phenomena that have previously been explained in terms of cognitive dissonance.

AUTHOR NOTE

Correspondence concerning this article should be sent to T. R. Zentall, Department of Psychology, University of Kentucky, Lexington, KY 40506 (e-mail: zentall@uky.edu).

REFERENCES

ALESSANDRI, J., DARCHEVILLE, J.-C., & ZENTALL, T. R. (2008). Cognitive dissonance in children: Justification of effort or contrast? *Psychonomic Bulletin & Review*, **15**, 673-677.

ARANTES, J., & GRACE, R. C. (2008a). Contrast and value: Beyond the work ethic effect. A reply to Zentall (2008). *Learning & Behavior*, **36**, 26-28.

ARANTES, J., & GRACE, R. C. (2008b). Failure to obtain value enhancement by within-trial contrast in simultaneous and successive discriminations. *Learning & Behavior*, **36**, 1-11.

ARONSON, E., & MILLS, J. (1959). The effect of severity of initiation on liking for a group. *Journal of Abnormal & Social Psychology*, **59**, 177-181.

CLEMENT, T. S., FELTUS, J. R., KAISER, D. H., & ZENTALL, T. R. (2000). "Work ethic" in pigeons: Reward value is directly related to the effort or time required to obtain the reward. *Psychonomic Bulletin & Review*, **7**, 100-106.

CLEMENT, T. S., & ZENTALL, T. R. (2002). Second-order contrast based on the expectation of effort and reinforcement. *Journal of Experimental Psychology: Animal Behavior Processes*, **28**, 64-74.

DIGIAN, K. A., FRIEDRICH, A. M., & ZENTALL, T. R. (2004). Reinforcers that follow a delay have added value for pigeons. *Psychonomic Bulletin & Review*, **11**, 889-895.

FANTINO, E., & ABARCA, N. (1985). Choice, optimal foraging, and the delay-reduction hypothesis. *Behavioral & Brain Sciences*, **8**, 315-330.

FESTINGER, L. (1957). *A theory of cognitive dissonance*. Stanford, CA: Stanford University Press.

FRIEDRICH, A. M., CLEMENT, T. S., & ZENTALL, T. R. (2005). Discriminative stimuli that follow the absence of reinforcement are preferred by pigeons over those that follow reinforcement. *Learning & Behavior*, **33**, 337-342.

FRIEDRICH, A. M., & ZENTALL, T. R. (2004). Pigeons shift their preference toward locations of food that take more effort to obtain. *Behavioral Processes*, **67**, 405-415.

GRACE, R. C. (1994). A contextual model of concurrent-chains choice. *Journal of Experimental Analysis of Behavior*, **61**, 113-129.

HULL, C. L. (1943). *Principles of behavior.* New York: Appleton-Century-Crofts.

KACELNIK, A., & MARSH, B. (2002). Cost can increase preference in starlings. *Animal Behaviour*, **63**, 245-250.

KLEIN, E. D., BHATT, R. S., & ZENTALL, T. R. (2005). Contrast and the justification of effort. *Psychonomic Bulletin & Review*, **12**, 335-339.

LAWRENCE, D. H., & FESTINGER, L. (1962). *Deterrents and reinforcements: The psychology of insufficient reward*. Stanford, CA: Stanford University Press.

MAZUR, J. E. (2001). Hyperbolic value addition and general models of animal choice. *Psychological Review*, **108**, 96-112.

MOORE, J., & FANTINO, E. (1975). Choice and response contingencies. *Journal of the Experimental Analysis of Behavior*, **23**, 339-347.

POMPILIO, L., & KACELNIK, A. (2005). State-dependent learning and suboptimal choice: When starlings prefer long over short delays to food. *Animal Behaviour*, **70**, 571-578.

POMPILIO, L., KACELNIK, A., & BEHMER, S. (2006, March 17). State-dependent learned valuation drives choice in an invertebrate. *Science*, **311**, 1613-1615.

SINGER, R. A., BERRY, L. M., & ZENTALL, T. R. (2007). Preference for a stimulus that follows a relatively aversive event: Contrast or delay reduction? *Journal of the Experimental Analysis of Behavior*, **87**, 275-285.

SINGER, R. A., & ZENTALL, T. R. (2008). *Effect of prior effort on reward value in pigeons*. Manuscript submitted for publication.

SKINNER, B. F. (1938). *The behavior of organisms: An experimental analysis.* New York: Appleton-Century-Crofts.

VASCONCELOS, M., & URCUIOLI, P. J. (2008). Deprivation level and choice in pigeons: A test of within-trial contrast. *Learning & Behavior*, **36**, 12-18.

VASCONCELOS, M., URCUIOLI, P. J., & LIONELLO-DENOLF, K. M. (2007). Failure to replicate the "work ethic" effect in pigeons. *Journal of the Experimental Analysis of Behavior*, **87**, 383-399.

ZENTALL, T. R. (2008). Within-trial contrast: When you see it and when you don't. *Learning & Behavior*, **36**, 19-22.

ZENTALL, T. R., CLEMENT, T. S., FRIEDRICH, A. M., & DIGIAN, K. A. (2006). Stimuli signaling reward that follow a less-preferred event are themselves preferred: Implications for cognitive dissonance. In E. A. Wasserman & T. R. Zentall (Eds.), *Comparative cognition: Experimental explorations of animal intelligence* (pp. 651-667). Oxford: Oxford University Press.

ZENTALL, T. R., & SINGER, R. A. (2007). Within-trial contrast: Pigeons prefer conditioned reinforcers that follow a relatively more rather than less aversive event. *Journal of the Experimental Analysis of Behavior*, **88**, 131-149.

(Manuscript received October 29, 2007;
revision accepted for publication May 16, 2008.)