



The midsession reversal task: A theoretical analysis

Thomas R. Zentall¹

Published online: 27 April 2020
© The Psychonomic Society, Inc. 2020

Abstract

In the *midsession reversal task*, choice of one stimulus (S1) is correct for the first half of each session and choice of the other stimulus (S2) is correct for the last half of each session. Although humans and rats develop very close to what has been called a win-stay/lose-shift response strategy, pigeons do not. Pigeons start choosing S2 before the reversal, making anticipatory errors, and they keep choosing S1 after the reversal, making perseverative errors. Research suggests that the pigeons are timing the reversal from the start of the session. However, making the reversal unpredictable does not discourage the pigeons from timing. Curiously, pigeons' accuracy improves if one decreases the value of the S2 stimulus relative to the S1 stimulus. Another form of asymmetry between S1 and S2 can be found by varying, over trials, the number of S1 or S2 stimuli. Counterintuitively, if the number of S2 stimuli varies, it results in a large increase in anticipatory errors but little increase in perseverative errors. However, if the number of S1 stimuli varies over trials, it results in a large increase in perseverative errors but no increase in anticipatory errors. These last two effects suggest that in the original midsession reversal task, the pigeons are learning to reject S2 during the first half of each session and learning to reject S1 during the last half of each session. These results suggest that reject learning may also play an important role in the learning of simple simultaneous discriminations.

Keywords Midsession reversal · Timing · Attention · Learning to reject · Pigeons

Introduction

Some might say that the comparison of cognitive abilities among animals is a questionable enterprise because an animal's cognitive abilities should evolve to the extent that they are needed for their survival or, if they do not, the species will go extinct. That is, animals without what we call cognitive abilities, such as our ability to develop a complex language, to form abstract concepts, and creatively manipulate our environment to inhabit much of the world, typically have not needed them. If that view is correct, we humans have needed them and presumably we have evolved them to survive. In our case, as measured by our numbers and our ability to survive in a large number of ecological niches, those cognitive abilities have allowed our species to be quite successful; however, in general, our closest relatives (e.g., the other great apes, other primates, and perhaps some of the sea mammals such as dolphins and porpoises) have not fared as well. Nevertheless, we

tend to view our cognitive abilities as positive attributes and we have been interested in comparing our abilities to those of other animals (Romanes, 1882).

The problem with that comparison is that differences in sensory abilities, methods of responding, and motivational systems can easily prevent us from appropriately assessing the abilities of other species. Thus, a rat will appear unable to learn to discriminate between a red light and a green light, if only because it is color blind, yet the ability of rats to discriminate between two different odors exceeds our own.

Bitterman (1965) described several ways to attempt to circumvent the problems associated with the assessment of cognitive abilities. He proposed that we should not try to assess, for example, the rate at which an animal can learn, because rate of learning will be affected by the sensory, response, and motivational differences among species that may have little to do with cognition. Instead, Bitterman suggested that one find a discrimination that the species in question can readily learn and consider the animal's flexibility in learning the reverse of the original discrimination. For example, if it takes a rat 50 trials to learn to choose consistently a particular arm of a T maze, one can ask how many trials it takes the rat to learn to reverse consistently that choice, and now choose the other arm of the maze. If one considers the number of trials it takes the

✉ Thomas R. Zentall
zentall@uky.edu

¹ Department of Psychology, University of Kentucky,
Lexington, KY 40506-0044, USA

animal to learn the reversal, relative to the number of trials to learn the original discrimination, one has a relative measure of the animal's cognitive ability.

Although Bitterman's (1965) logic seems reasonable, in practice even the relative ability to acquire the reversal of a discrimination appears to depend, at least in part, on the conditions under which the learning and reversal are assessed. For example, in studying reversal learning by fish, not only does the way in which one provides reinforcement affect the rate of original learning, but it also affects the rate of reversal learning, relative to original learning (Mackintosh & Cauty, 1971).

The problems with the comparison of reversal learning among species notwithstanding, reversal learning has been studied as an important kind of learning. In fact, Bitterman (1965) has attempted to order several species according to how much they can benefit from a serial reversal task, one in which there is a series of reversals, the measure being the asymptotic proportion of errors made with a series of reversals, relative to original learning.

The midsession reversal task

A version of the serial reversal task is one that was described by Rayburn-Reeves, Molet, and Zentall (2011; see also Cook & Rosen, 2010) in which, in each session, animals are trained on a two-alternative simultaneous discrimination to choose S1 rather than S2 for the first number of trials (e.g., 40) and then choose S2 rather than S1 for the remaining trials. If this midsession reversal task is repeated every session with the same two stimuli, one can ask if animals develop an efficient response "strategy": Choose S1 until it stops being correct and then choose S2. Not surprisingly, this so-called win-stay/lose-shift strategy (Levine, 1975; Restle, 1962) is what humans (Rayburn-Reeves, Molet, & Zentall, 2011) typically develop with experience. Even with considerable experience with this task, however, pigeons do something quite different. Pigeons begin to choose S2 well before the reversal, making what are called anticipatory errors, and once the reversal has occurred, they continue to make what are called perseverative errors (Rayburn-Reeves, et al., 2011; see Fig. 1, which depicts choice of the first correct stimulus over the trials in a session). Rats, however, when tested under conditions similar to that of pigeons, show accuracy closer to that of humans than of pigeons (Rayburn-Reeves, Stagner, Kirk, & Zentall, 2013; see Fig. 2).

Although the function presented in Fig. 2 looks quite symmetrical (anticipatory and perseverative errors look to be relatively similar in number) the two kinds of errors provide very different kinds of feedback. Anticipatory errors let the animal know that the reversal has not yet occurred, but they provide little information about the consequence of choosing S1 or S2 on the next trial. Perseverative errors, however, let the animal know that the reversal has already occurred, and they indicate

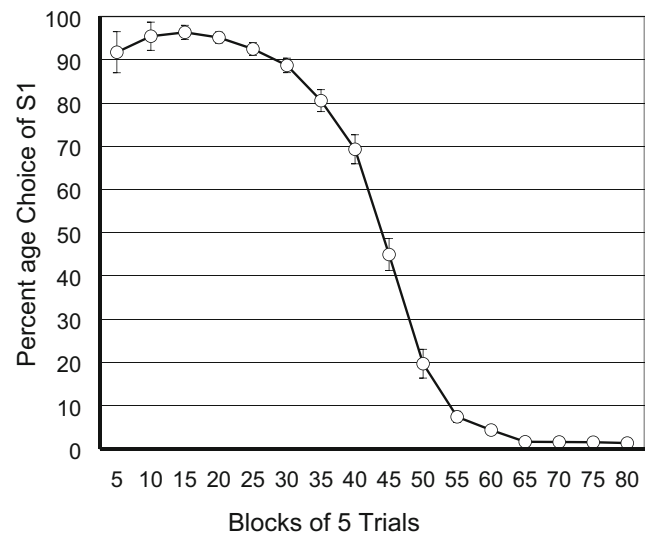


Fig. 1 Pigeons' performance on the midsession reversal task as a function of trial number. S1 was correct for the first 40 trials, S2 was correct for the last 40 trials (after Rayburn-Reeves et al., 2011)

that choice of S1 will no longer be reinforced. Thus, although successive anticipatory errors may be justified, successive perseverative errors are not.

Timing the occurrence of the reversal

The shape of the function presented in Fig. 2 appears very similar to the shape of a psychophysical timing function found after training animals on a conditional temporal discrimination. For example, after training pigeons to choose a red comparison stimulus when presented with a short (e.g., 2 s) sample stimulus and a green comparison stimulus when presented with a longer (e.g., 8 s) sample stimulus, one can test the pigeons with sample stimuli of intermediate durations. A plot of the choice of the red stimulus (correct after the short sample) as a function of the sample duration looks very similar to




Fig. 2 Rats' performance on the midsession reversal task as a function of trial number. S1 was correct for the first 40 trials, S2 was correct for the last 40 trials (after Rayburn-Reeves, Stagner et al., 2013)

the plot of the choice of the S1 stimulus as a function of trial number in Fig. 2. That is, to determine when to reverse the discrimination, the pigeons appear to be timing from the start of the session to the reversal.

Given the repetitive nature of this task, various cues are available to the pigeon to signal the point in the session at which the reversal will occur. The pigeons could be sensitive to internal cues such as the amount of food they have eaten, the number of trials that have passed, or the amount of time that has passed. Alternatively, they could be sensitive to external cues, such as the feedback received from choice of S1 or S2 on the preceding trial(s). Although feedback from the preceding trial(s) would appear to be the most reliable cue, that does not appear to be the source of primary stimulus control for pigeons.

Cook and Rosen (2010) showed that amount of food eaten does not control the transition between choice of S1 and choice of S2 because prefeeding the pigeons had little effect on when the pigeons began choosing S2. The estimation of trials into the session and time into the session both also provide cues for the occurrence of the reversal. Given the fact that time from the start of the session to the reversal is likely to be a bit more variable because of the potential for variability in the latency of choice, estimation of the number of trials into the session would appear to be the more reliable cue. However, rather than using the estimation of trials to the reversal or what one might consider the more reliable feedback from the preceding trial as the basis for switching from choice of S1 to choice of S2, pigeons appear to be timing from the start of the session to the reversal.

The pigeons' attempt to estimate the time from the start of the session to the reversal can be seen, after training with a particular intertrial interval, by testing pigeons with longer or shorter intertrial intervals (McMillan & Roberts, 2012; Smith, Beckmann, & Zentall, 2017). If during training the intertrial interval is 5 s and it is lengthened to 10 s, pigeons begin to reverse much earlier, as measured by trials into the session, because the reversal is now delayed. If, however, the intertrial interval is shortened to 2.5 s, the pigeons now begin to reverse well after the reversal because the reversal now comes earlier, as measured by trials into the session (see Fig. 3).

Another way to demonstrate the pigeons' strong tendency to estimate the time from the start of the session to the reversal is what happens when the trial in the session at which the reversal occurs is made unpredictable (Rayburn-Reeves et al., 2011). Surprisingly, when the reversal happens to occur early in the session, pigeons make a large number of perseverative errors, but they make few anticipatory errors. However, when the reversal happens to occur late in the session, the pigeons make a large number of anticipatory errors and few perseverative errors (see Fig. 4). In fact, although the reversal is no more likely to occur at the midpoint of the session than anywhere else, the pigeons tend to be most accurate, making the fewest total

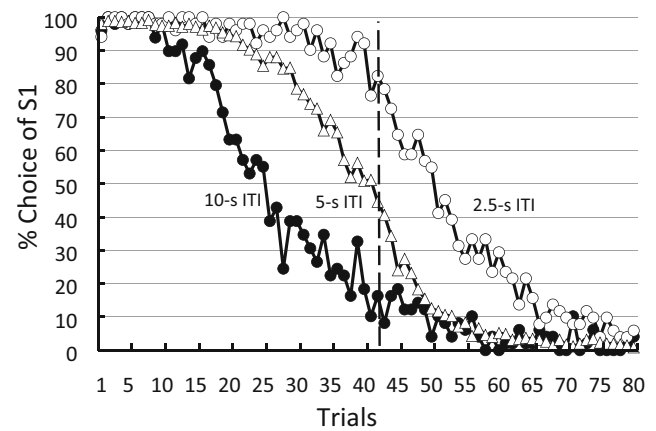


Fig. 3 Pigeons were trained on the midsession reversal task with 5-s intertrial intervals and were transferred to 10-s or 2.5-s intertrial intervals (after Smith et al., 2017)

errors, when the reversal occurs at the middle of the session. This result was found whether pigeons were originally trained with the reversal after trial 40 of an 80-trial session (Rayburn-Reeves et al., Experiment 1) or pigeons were trained with the reversal at a variable trial in the session from the start of training (Rayburn-Reeves et al., Experiment 2).

The use of time from the start of the session as a cue to reverse would appear to be a poor choice of cue by pigeons. Although one could argue that many environmental events have temporal periodicities, such as the natural day/night cycle or seasons of the year, those events are also accompanied by external stimuli: sunrise and sunset in the case of day/night, temperature change and other external cues in the case of




Fig. 4 Pigeons were trained on the midsession reversal task with S1 correct at the start of the session and S2 correct at the end of the session with a single reversal that could occur unpredictably after trial 10, 25, 40, 55, or 70 (after Rayburn-Reeves et al., 2011). Note the large number of perseverative errors when the reversal occurred early in the session and the large number of anticipatory errors when the reversal occurred late in the session

changes in season. Although the timing of short intervals may be somewhat unnatural, animals do reasonably well at it.

One of the most studied procedures used to study animal timing is the peak procedure in which animals are trained on a fixed-interval schedule, in which the first response after a fixed time provides reinforcement (S. Roberts, 1981). After such training, on so-called empty trials, reinforcement is not provided but responses continue to be recorded. A plot of the rate of responding over time on those nonreinforced empty trials has been taken as evidence of animals' timing ability. It is known as the peak procedure because when the response rate is averaged over those nonreinforced trials, the resulting function has a peak value at roughly the point in the trial at which reinforcement would have occurred on a fixed-interval training trial. However, if one examines responding on individual nonreinforced trials, it shows very little in common with the average rate of responding, over those trials, as a function of time into the trial (Cheng, Westwood, & Crystal, 1993). What generally occurs is what is known as break-run-break. The animal starts responding before the timed interval at a relatively constant rate and continues responding sometime after the timed interval. What appears to vary is not so much the rate of responding but the time at which responding starts and stops.

Another procedure that has been used to study animal timing is the temporal discrimination task, already mentioned, in which animals learn to discriminate between a relatively short and a longer interval and test trials with intermediate durations are used to determine the underlying psychophysical function. The psychophysical function can provide evidence for the nature of the underlying temporal scale. That is, at what sample duration the animal chooses the comparison stimulus associated in training with the short duration as often as the comparison stimulus associated in training with the long duration. In keeping with Weber's Law, that point is typically close to the geometric mean of the two training intervals.

The similarity of the psychophysical function obtained following training on the temporal discrimination task and the similar function obtained following training on the midsession reversal task suggests that the underlying processes are similar. It should be noted, however, that an important difference between the two tasks is, in the temporal discrimination, on intermediate duration test trials, there is no correct response, whereas in the midsession reversal task, there is always a correct response and a valid external cue for reinforcement is typically the feedback from the preceding response. Why pigeons do not make better use of those feedback cues is not clear but research by McMillan, Sturdy, and Spetch (2015) provides some clues.

Competition between S1 and S2

McMillan et al. (2015) trained pigeons on a midsession reversal task that involved successive presentation of the S1 and S2

stimuli. With this go/no-go version of the task, the response measure was whether the stimulus was pecked within 3 s. They found that the pigeons almost always pecked the S1 stimulus before the reversal and they almost always pecked the S2 stimulus after the reversal. However, increasingly, the pigeons began pecking the S2 stimulus as the reversal approached, and they continued to peck the S1 stimulus well after the reversal had occurred (see Fig. 5). But of course, an important difference between the simultaneous two-alternative discrimination and the successive discrimination is the consequence of making an "error." In the case of the successive discrimination, there is virtually no cost to pecking a stimulus in error, whereas failing to peck in error means the loss of reinforcement. Thus, pigeons may adopt the general rule, when in doubt, peck.

The successive discrimination midsession reversal task offers some suggestion as to the mechanism responsible for anticipatory and perseverative errors. On initial trials in a session, when S1 is correct but not S2, there is little competition between the two stimuli, but as the reversal approaches, inhibition to S2 decreases, resulting in increased competition between S1 and S2. That increased competition likely results in anticipatory errors. Following the reversal, competition between S1 and S2 is still present, resulting in perseverative errors, but as inhibition to S1 begins to build up, responding to S1 declines.

Smith, Zentall, and Kacelnik (2018) confirmed this hypothesis using a hybrid, simultaneous-successive discrimination version of the midsession reversal task. They presented trials in blocks of five, with the first four trials involving randomly presented successive presentations of S1 (two trials) and S2




Fig. 5 Pigeons were trained on a hybrid midsession reversal task in which in each block of five trials there were four go/no-go trials (two S1 and two S2 trials) and one choice trial. For the first 40 trials S1 was correct. For the last 40 trials S2 was correct. The graph shows the latency of the response to S1 and S2 as a function of trial number (after Smith, Zentall, & Kacelnik, 2018). Note that the latency to the correct stimulus does not change as a function of the trial number

(two trials), with the fifth trial a simultaneous discrimination trial. Smith et al. found that the latency difference between pecks to the S1 stimulus and pecks to the S2 stimulus predicted accuracy on the simultaneous discrimination trial that followed. As the latency difference approached zero, anticipatory and perseverative errors increased.

The results of the Smith et al. (2018) study support the sequential choice (or horse-race) model of discrimination learning, which proposes that the tendency to choose S1 and S2 increases independently and the one that reaches threshold first wins (Kacelnik, Vasconcelos, & Monteiro, 2011). The sequential choice model is in contrast to the, perhaps more intuitive, tug-of-war model in which as the reversal approaches, the competition between S1 and S2 should result in an increase in the latency to make a choice.

Both the McMillan et al. (2015) and the Smith et al. (2018) studies identify competition between S1 and S2 and the cause of anticipatory and perseverative errors. And that competition comes from the reduction in inhibition to S2 prior to the reversal and the relative lack of inhibition to S1 immediately after the reversal.

Persistence in the use of time from the start of the session to the reversal as a cue to reverse is not to suggest that the pigeons totally avoid using the feedback from recent trials as a cue. As already noted, when the point of the reversal in the session is made unpredictable, although pigeons show more perseverative errors when the reversal occurs early in the session, they do begin to reverse sooner than they would if the reversal had not occurred. Nevertheless, even when the reversal is unpredictable, the pigeons make the fewest errors when the reversal occurs in the middle of the session. That is, it appears that the pigeons use as a cue the average time of the occurrence of the reversal from the start of the session. Furthermore, when the reversal is unpredictable and it happens to occur at the midpoint of the session, the pigeons do no better than when the reversal occurs at the midpoint of the session on all sessions. However, although pigeons continue to use temporal cues, even when the reversal occurs at an unpredictable time in the session, they do show some sensitivity to the outcomes of the choices made on the most recent trials.

The role of spatial cues

Research on the source of stimulus control in the midsession reversal task with pigeons has led to some interesting findings. For example, McMillan and Roberts (2012) trained pigeons on the midsession reversal task with redundant visual and spatial cues and found improved accuracy relative to visual cues alone (in which the visual cues changed location randomly from trial to trial). Not only did accuracy improve, but when they manipulated the intertrial interval, the resulting small shift in the point of the reversal indicated that these pigeons were not using the

passage of time as much as the control group that had visual cues that varied in location from trial to trial. McMillan and Roberts suggested that the addition of spatially relevant cues allowed the pigeons to anticipate where to make the next response. When spatial cues alone are used with pigeons, however, the result is not very different from the use of visual cues alone, at least with the typically used 5-s intertrial interval (Laude, Stagner, & Zentall, 2014; see Fig. 6).

Rayburn-Reeves, Stagner, Kirk, and Zentall (2013) found support for the hypothesis that pigeons sometimes may use relevant spatial cues to anticipate where to respond next. Using the midsession reversal of a spatial discrimination, they manipulated, between groups, the intertrial interval. When the intertrial interval was either 5 s or 10 s, a typical number of anticipatory and perseverative errors were found. When the intertrial interval was only 1.5 s, however, the pigeons showed accuracy that approached a win-stay/lose-shift strategy (see Fig. 7). Their explanation for this finding was, with a very short intertrial interval, the next trial followed reinforcement immediately, thus making a repetitive peck-S1-eat-peck-S1 movement, without a break, quite possibly because the location of S1 could be anticipated. Furthermore, interruption of that sequence by the absence of reinforcement may have served as a salient cue to switch responding to the other location. Under these conditions, it is possible that the pigeons were using their postural orientation or their physical location in the operant chamber as an important basis of their choice. Of course, it should have been possible for the pigeons to use the same overt cues when the intertrial interval was 5 s or 10 s, but as Rayburn-Reeves and Cook (2016) noted, during typical intertrial intervals of 5 s or 10 s, pigeons often move in the operant box, thus making those postural cues at the start of each trial less effective.

Although it is possible that the reduction in the intertrial interval to 1.5 s shortened the session and thereby made timing more accurate, the fact that pigeons that had a visual color

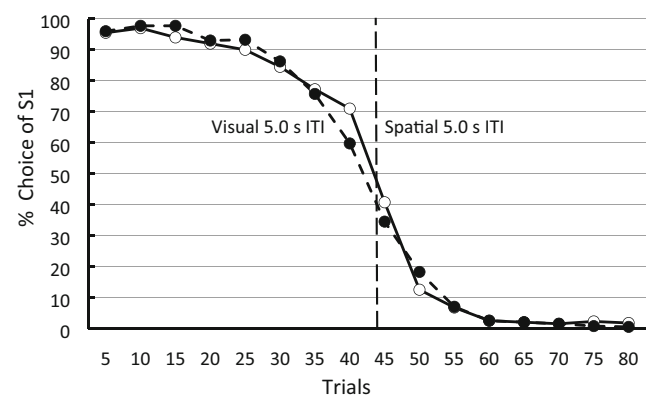


Fig. 6 Comparison of the visual versus the spatial midsession reversal task with 5-s intertrial intervals with pigeons (after Rayburn-Reeves, Stagner et al., 2013). Choice of S1 was correct for the first 40 trials. Choice of S2 was correct for the last 40 trials




Fig. 7 The spatial midsession reversal task with pigeons as a function of trial number, with intertrial intervals of 1.5 s, 5 s, and 10 s (after Rayburn-Reeves, Stagner et al., 2013)

discrimination with short, 1.5-s intertrial intervals showed no such facilitation (Laude et al., 2014; Fig. 8) makes less likely the hypothesis that shorter time from the start of the session to the reversal was responsible for the accurate performance by pigeons that had the spatial discrimination with short intertrial intervals (Rayburn-Reeves, Stagner et al., 2013).

Midsession reversal learning by rats

Most of the research on the midsession reversal task has been done with pigeons, but several experiments have been conducted with rats. Because rats tend to have poor vision they have been tested with a spatial discrimination. When rats were tested on the midsession reversal task with a left versus right lever-press response and a 5-s intertrial interval, they

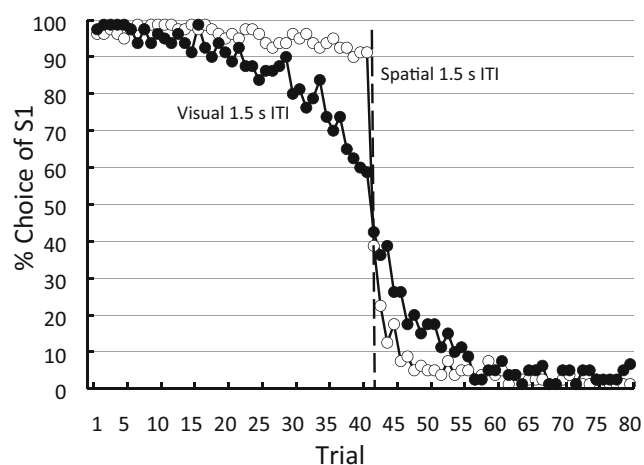


Fig. 8 The midsession reversal task with pigeons. Comparison of the spatial and visual discriminations with 1.5-s intertrial intervals. S1 was correct for the first 40 trials. S2 was correct for the last 40 trials (after Rayburn-Reeves, Stagner et al., 2013)

showed close to win-stay/lose-shift accuracy (Rayburn-Reeves, Stagner et al., 2013; see Fig. 1). Similarly, rats trained on the midsession reversal task with a left versus right nose-poke response showed similar accuracy (Smith, Pattison, & Zentall, 2016). Pigeons, however, when tested similarly with a spatial discrimination and a 5-s intertrial interval, continued to show a large number of anticipatory and perseverative errors (Rayburn-Reeves, Laude, & Zentall, 2013) unless a redundant visual discrimination is added (McMillan & Roberts, 2012)

To test the hypothesis that the rats may have been using their postural orientation or their physical location in the operant chamber as an important basis of their choice, McMillan, Kirk, and Roberts (2014) tested rats on the midsession reversal task in a T maze with a spatial discrimination. By removing the rats from the maze after each response and returning them to the start box, they could avoid postural mediators between trials. Under these conditions, they found poor accuracy, quite similar to that of the pigeons. Furthermore, much like the pigeons, when the point of reversal was varied over sessions accuracy did not improve. Thus, the difference between rats and pigeons is likely that, with the spatial midsession reversal task with 5-s intertrial intervals, the pigeons were more active during the intertrial interval than the rats and the rats were better able to use their spatial orientation between trials as a cue for their next choice than the pigeons.

Memory for the stimulus last chosen and the results of that choice

Smith, Beckmann, and Zentall (2017) hypothesized that the reason that pigeons used the time from the start of the session to the reversal may have been because they were not always able to remember the events from the preceding trial. That is, to be able to make an appropriate choice, the pigeons would have to have remembered what stimulus they chose and the outcome of that choice. Although the ability to retrieve that information from the preceding trial should have been relatively easy, Randall and Zentall (1997) found that pigeons can show considerable forgetting with a procedure not very different from the midsession reversal.

In the Randall and Zentall (1997) experiment, pigeons were required to peck a white light that was presented either on the left or on the right key. The result of that peck was either reinforcement or the absence of reinforcement. However, the spatial location of the peck, together with the outcome of the peck, could then serve as a cue for choice of a left or right comparison stimulus. For example, if a left peck was followed by reinforcement, a peck to the left comparison stimulus would be reinforced, whereas if a left peck was not followed by reinforcement, a peck to the right comparison stimulus would be reinforced. Similarly, if a

right peck was followed by reinforcement, a peck to the right comparison stimulus would be reinforced, whereas if a right peck was not followed by reinforcement a peck to the left comparison stimulus would be reinforced. Effectively, these pigeons were trained to win-stay/lose-shift. The pigeons had no problem learning this task when there was no delay between the initial reinforcement (or its absence) and comparison choice. When a delay as short as 4 s was inserted between the initial reinforcement (or its absence) and comparison choice, however, accuracy dropped to about 65% correct. Thus, the pigeons had difficulty remembering where they pecked and the outcome of that peck when the delay was just 4 s.

Based on the Randall and Zentall (1997) finding, Smith et al. (2017) tested the hypothesis that the pigeons' inability to remember the preceding response and its outcome may have contributed to their inability to avoid making anticipatory and perseverative errors. To do this, they presented pigeons with "reminder" cues during the intertrial interval. Specifically, if the red stimulus had been chosen on the preceding trial, a houselight located at the top of the response panel was turned on throughout the intertrial interval. If the green stimulus had been chosen on the preceding trial, however, a houselight located in the middle of the ceiling of the operant box was turned on throughout the intertrial interval. In addition, if the response on the preceding trial had been correct, following 1.5 s of reinforcement, the feeding tray was lowered but the feeder light remained on for the remainder of the intertrial interval.

Smith et al. (2017) found that providing the pigeons with a reminder of the stimulus chosen and the outcome of that choice significantly improved the pigeons' accuracy, but the pigeons still made a significant number of anticipatory errors prior to the reversal and perseverative errors following the reversal (see Fig. 9).




Fig. 9 The midsession reversal task with pigeons. S1 was correct for the first 40 trials. S2 was correct for the last 40 trials. During the 5-s intertrial interval, the experimental group was given reminders of the stimulus they had chosen on the preceding trial and whether that choice had been reinforced or not (after Smith et al., 2017)

Cuing and miscuing which stimulus to choose

A different approach to providing pigeons with an external cue to the correct stimulus was attempted by Rayburn-Reeves, Qadri, Brooks, Keller, and Cook (2017). They trained pigeons on a spatial midsession reversal task, but on some sessions, during the intertrial interval, the two halves of the session were signaled by a background color. For example, the background color was blue during the intertrial interval for the first 40 trials and it changed to yellow for the last 40 trials. Cued and un-cued sessions were alternated. Rayburn-Reeves et al. found, not surprisingly, that cuing the pigeons virtually eliminated both anticipatory and perseverative errors compared with un-cued sessions. Rayburn-Reeves et al. then introduced occasional probe trials on which the cue presented during the intertrial interval was incorrect. That is, the pigeons were occasionally mis-cued during the session. They did this to determine the effect of the miscue at different points in the session. They found that miscues presented either early or late in the session had little effect on the pigeons' choice. That is, control of the pigeons' choice early and late in the session appeared to be based on the passage of time from the start of the session. Miscues presented as the reversal approached and immediately after the reversal, however, produced a large decrease in accuracy. Thus, in the temporal region around the reversal, where the temporal cue was somewhat ambiguous, the pigeons' choice appeared to be controlled by the intertrial interval cues.

As mentioned earlier, it is noteworthy that competition between the response strength to S1 and S2 that occurs in the region around the reversal and that appears to be under temporal control does not result in an increase in response latency. Both Rayburn-Reeves and Cook (2016) and Smith et al. (2018) have found a relatively flat function when response latency is plotted as a function of trial in the session. This finding suggests that the increase in response strength for choice of S2 over the course of the first half of the session, and the decline in response strength of S1 over the last half of the session, are independent and have little effect on the latency of choice of the two stimuli. This result belies the intuitive notion that it requires more time to make difficult choices.

Effect of asymmetries between S1 and S2 on midsession reversal accuracy

In the free operant psychophysical procedure (Bizo & White, 1995), animals are given a choice between two stimuli. For the first half of each trial, responding to one of the stimuli (S1) is reinforced on a variable interval schedule (the first response after a variable amount of time is reinforced), whereas during the second half of each trial, responding to the other stimulus is reinforced, also on a variable interval schedule. With the free operant psychophysical procedure, the psychophysical

function relating the proportion of responses to S1 during a trial is very similar to the function relating the proportion of choices of S1 relative to trial number in the midsession reversal task. In the free operant psychophysical procedure, pigeons choose the S1 alternative early in each trial and the S2 alternative late in each trial, and they choose a mixture of S1 and S2 in the middle of the trial. Interestingly, consistent with Kacelnik et al.'s (2011) sequential choice model, although the distribution of responses varies as a function of time in the trial, the combined number of responses to S1 and S2 were about the same in the region of the reversal as they were both early and late in the trial (Bizo & White, 1995). The main difference between the midsession reversal task and the free operant psychophysical procedure is that in the free operant psychophysical procedure, the reinforcement feedback following choice of either alternative is quite intermittent because most responses to S1 and S2 are not reinforced.

Furthermore, Bizo and White (1995) found, as one might expect, when the schedule of reinforcement to the S1 stimulus was made richer than to the S2 stimulus, the psychophysical function shifted in the direction of S1 and when the schedule of reinforcement to the S2 stimulus was made richer than to the S1 stimulus, the psychophysical function shifted in the direction of S2. The more likely choice of the S1 stimulus (relative to the S2 stimulus) was reinforced, the longer from the start of the trial the pigeons tended to choose S1, and the more likely choice of the S2 stimulus (relative to the S1) was reinforced, the sooner from the start of the trial the pigeons tended to choose S2.

In the midsession reversal task, if one were to manipulate the relative probability of reinforcement for correct choice of the S1 and S2 stimuli, one would expect a similar shift in the psychophysical function, a bias in the direction of the richer schedule. When Santos, Soares, Vasconcelos, and Machado (2019) reduced the probability of reinforcement for correct choice of the S1 stimulus from 100% to 20%, they found the expected shift in the psychophysical function; earlier choice of the S2 stimulus resulted in an increase in anticipatory errors. When they reduced the probability of reinforcement for correct choice of the S2 stimulus from 100% to 20% (correct choice of the S1 stimulus remained at 100%), however, the results were quite surprising. Instead of finding a general shift in the probability of choice of the S1 stimulus resulting in a decrease in anticipatory errors and an increase in perseverative errors, they found a decrease in the total number of errors. Although they did find a decrease in the number of anticipatory errors, they found no increase in the number of perseverative errors. In fact, the number of perseverative errors also appeared to decrease.

Santos et al. (2019) reasoned that when the probability of reinforcement for correct choice of the S1 stimulus was only 20%, the pigeons were using the passage of time, biased by reward following (more reinforcement associated with choice

of S2). When the probability of reinforcement for correct choice of the S1 stimulus was 100% but was only 20% for correct choice of the S2 stimulus, however, the pigeons appeared to rely primarily on local cues, a strategy that resulted in close to win-stay/lose-shift accuracy. Thus, paradoxically, decreasing the probability of reinforcement for correct choice of S2 had the effect of increasing overall task accuracy. It appears that the reduction in the probability of reinforcement for correct choice of the S2 stimulus decreased the pigeons' use of the passage of time from the start to the session and increased their use of the feedback from their choice of the preceding trial(s). But why did this occur?

Zentall, Andrews, Case, and Peng (2020) were able to replicate the effect reported by Santos et al. (2019; see Fig. 10). They proposed that under the typical conditions in which there was 100% reinforcement for correct choice of both the S1 and S2 stimuli, competition between the S1 and S2 stimuli was responsible for both anticipatory and perseverative errors. By reducing the validity of the feedback from choice of the S2 stimulus, however, the competition between S1 and S2 was reduced. Although the feedback from choice of the S2 stimulus was of reduced validity because nonreinforcement could occur for incorrect first half choice of the S2 stimulus as well as correct second half choice of the S2 stimulus, the feedback from choice of the S1 stimulus alone was sufficient to develop an efficient win-stay/lose shift strategy. That is, it may be that the bias to choose the S1 stimulus during the entire first half of the session made nonreinforcement for choice of the S1 stimulus at the time of the reversal more discriminable.

It may be that the anticipatory errors experienced under the typical 100% correct choice of both the S1 and S2 stimuli produced an effect similar to that of a schedule of partial




Fig. 10 Performance of pigeons on the midsession reversal task in which S1 was correct for the first 40 trials and S2 was correct for the last 40 trials. For the experimental group, correct choices of S1 were reinforced 100% of the time but correct choices of S2 were reinforced only 20% of the time. For the control group, correct choices of both S1 and S2 were reinforced 100% of the time (after Zentall, Andrews, Case, & Peng, 2020)

reinforcement. That is, as the reversal approached, the tendency to choose the S2 stimulus resulted in an increasing number of nonreinforced, anticipatory errors (partial reinforcement). But when correct S2 choice is rarely reinforced, the resulting bias to choose the S1 stimulus throughout the first half of the session should result in what would become a schedule of continuous reinforcement. If this analysis is correct, something akin to a partial reinforcement extinction effect (Humphreys, 1939) may result. Thus, the sudden nonreinforcement at the time of the reversal may have created a transition not unlike that of extinction following continuous reinforcement, when compared to the slower extinction typically found following partial reinforcement.

Of course, when all correct choices are reinforced 100% of the time, responses to either the S1 or the S2 stimulus provides perfect feedback as to where in the session the animal is. However, could this be a procedure in which there is too much information? That is, could the pigeons be choosing the S2 stimulus before the reversal to avoid missing any reinforcers associated with choice of the S2 stimulus? But if correct choice of the S2 stimulus is reinforced only 20% of the time, could the pigeons be learning to ignore the feedback from choice of the S2 stimulus? And by learning to ignore the feedback from choice of the S2 stimulus, could the pigeons' attention to the feedback from choice of the S1 stimulus be sufficient to develop a pattern of choice close to win-stay/lose-shift?

But why does the reverse contingency (20% reinforcement for correct choice of the S1 stimulus and 100% reinforcement for correct choice of the S2 stimulus) not result in a similar reduction in errors? Unfortunately, during the first half of the session, choice of the S2 stimulus, but not the S1 stimulus, provides perfect feedback, but getting that feedback requires that the pigeon makes an error. Thus, with 20% reinforcement of correct S1 choices, one would expect a bias to choose S2; however, this bias would result in many more anticipatory errors, which is exactly what Santos et al. (2019) found.

Zentall, Andrews, et al. (2020) proposed that when only 20% of correct S2 choices are reinforced, choice of S2 does not provide useful information. For this reason, the pigeons attend primarily to the feedback from choice of S1. When choice of S1 stops providing reinforcement, the pigeons begin to avoid S1. According to the Zentall, Andrews, et al. attentional interpretation of the paradoxical finding reported by Santos et al. (2019), any bias that is introduced that favors the S1 stimulus relative to the S2 stimulus, should produce a similar beneficial result – a decrease in anticipatory errors but virtually no increase in perseverative errors. Zentall, Andrews, et al. tested this hypothesis by increasing the response requirement for all choices of the S2 stimulus from one to ten pecks but leaving the response requirement for all choices of the S1 stimulus at one peck. In support of the hypothesis, pigeons that were required to make ten pecks for choice of the S2

stimulus, but only one peck for choice of the S1 stimulus, made significantly fewer anticipatory errors and made no more perseverative errors than pigeons in the control group that were required to make a single peck for choice of either the S1 or S2 stimulus.

Furthermore, to test the hypothesis that the pigeons in the experimental group were learning to avoid S1 (as well as learning to choose S2), Zentall, Andrews, et al. (2020) tested the pigeons by replacing the S2 color with a novel shape. If pigeons in the experimental group were learning to avoid the S1, the novel S2 stimulus should have been less disruptive than for the control group. On the novel S2 test session, the pigeons in the control group that had been trained with 10 peck for choice of the S2 stimulus made 75% errors on the test session following the reversal, rarely choosing the novel S2 stimulus. Pigeons in the experimental group made only 36% errors.

Recently, Mueller, Halloran, and Zentall (submitted) further tested the attentional hypothesis by manipulating the magnitude of reinforcement for correct choice of the S1 and the S2 stimulus. For the experimental group, correct choice of the S1 stimulus produced five pellets of food, whereas correct choice of the S2 stimulus produced only one pellet. For the control group, correct choice of either the S1 or S2 stimulus produced three pellets of food. Relative to the control group, they found fewer anticipatory errors and no increase in perseverative errors for the experimental group. These results provide further support for the attentional hypothesis.

Learning to select or learning to reject?

Manipulation of the number of stimuli in the S2 set

If an asymmetry between the consequences of choice of the S1 and S2 stimuli, favoring the S1 stimulus, biases the pigeons to attend to the S1 stimulus and reduces the interference from the S2 stimulus, resulting in reduction in the total number of errors, could the same result be accomplished by manipulating the number of S2 stimuli that could appear over trials in a session? Zentall, Peng, House, and Yadev, (2020) tested this hypothesis by allowing the S2 stimulus to vary from trial to trial. In that study, there were four S2 stimuli, one per trial, that appeared equally often over trials, but there was only one S1 stimulus. The idea was by varying the S2 stimulus from trial to trial, it would presumably make it more difficult for the pigeons to learn to choose the S2 stimulus, which might result in a bias to choose the S1 stimulus. And if this hypothesis is correct, the bias to choose the S1 stimulus should result in a decrease in anticipatory errors, and no increase in perseverative errors.

Zentall, Peng, House, and Yadev, (2020) conducted two experiments. In the first experiment the stimuli were colors:

red, green, yellow, blue, and white. Because Zentall, Peng, et al. were concerned that there might be some stimulus generalization between one of the S2 stimuli and the S1 stimulus (which was red for half of the pigeons and green for the remaining pigeons), they conducted a second experiment with stimuli consisting of national flags that differed from each other in color, shape, and pattern. The surprising finding was inconsistent with the attentional hypothesis proposed by Zentall, Andrews, et al. (2020). In both the color and national-flag experiments, instead of finding a decrease in anticipatory errors relative to the control group that had a single S1 stimulus and a single S2 stimulus, they found a significant increase in anticipatory errors (see Fig. 11). They also found no difference between the experimental and control groups in perseverative errors. Zentall, Andrews, et al. proposed three hypotheses that might account for the results.

First, it was possible that stimulus generalization between one of the S2 stimuli and the S1 stimulus made the S1–S2 discrimination more difficult. Zentall, Andrews, et al. (2020) rejected this hypothesis because, had the effect been produced by stimulus generalization, one would have expected the difference between the experimental and control groups to have been a more general effect. That is, it should have led to an increase in both anticipatory and perseverative errors.

Second, Zentall, Andrews, et al. (2020) proposed that the variability in the S2 stimuli may have drawn attention to the S2 stimulus during the first half of each session rather than to the S1 stimulus. This hypothesis came from Beckmann and Young's (2007) finding that, under certain conditions, variability can serve as an abstract feature (see also Macphail & Reilly, 1989). If so, it might increase the difficulty in learning to choose S1 during the first 40 trials. Drawing attention to the variable S2 stimuli would account for the poorer accuracy for the experimental group during the first 40 trials of the session.

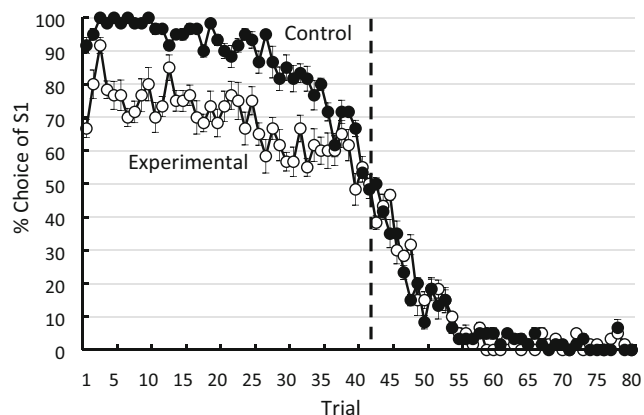


Fig. 11 Performance of pigeons on the midsession reversal task in which S1 was correct for the first 40 trials and S2 was correct for the last 40 trials. For the experimental group there was one S1 stimulus on all trials but the S2 stimulus changed from trial to trial. For the control group there was a single S1 and a single S2 stimulus on all trials (after Zentall, Peng, House, & Yadev, 2020)

However, drawing attention to the variable S2 stimuli during the last 40 trials of the session would predict there should have been an advantage over the last 40 trials of the session relative to the single-S2 control group, but that was not found.

A third mechanism that could be responsible for increased anticipatory errors when the S2 stimuli change from trial to trial has to do with how learning occurs in the midsession reversal task. As noted earlier, McMillan et al. (2015) trained pigeons on a single stimulus version of the midsession reversal in which S1 and S2 were presented successively over trials (go/no-go), and they found that the successive discrimination consisted of learning to inhibit responding to S2 during the first half of each session and learning to inhibit responding to S1 during the second half of each session. They concluded that learning to inhibit was also how the pigeons learned the more typical simultaneous discrimination version of the task. They concluded further that the problem that pigeons had with the midsession reversal task prior to the reversal was learning to inhibit responding to S2 as the reversal approached rather than learning to continue responding to S1. They suggested that in the midsession reversal task pigeons learn independent rules about the S1 and S2 stimuli (see also Smith et al., 2018).

Manipulation of the number of stimuli in the S1 set

In the Zentall, Peng, et al. (2020) experiment, if the pigeons had to learn to inhibit responding to the varying S2 stimuli, because there were multiple S2 stimuli, it would have made learning to inhibit responding to the S2 stimuli more difficult, thus leading to an increase in anticipatory errors. It would not have affected perseverative errors, however, because for both groups, the S- would now be represented by the single S1 stimulus.

Zentall, Halloran, and Peng (submitted) tested this hypothesis. Instead of increasing the number of S2 stimuli over trials, they increased the number of S1 stimuli. The stimuli were the same as those used by Zentall, Peng, et al. (2020). All that changed was their designation as S1 or S2.

Zentall, Peng, et al. (2020) had described three hypotheses to account for their results. What would those hypotheses predict for the results of the Zentall, Halloran, et al. (submitted) procedure? If stimulus generalization made the discrimination between the S1 and S2 stimuli more difficult, one might expect the results of Zentall, Halloran, et al. to be similar to the results of Zentall, Peng, et al., or more likely it should result in a general increase in both kinds of errors.

If, however, in the Zentall, Peng, et al. (2020) experiment, attention to the variability of the S2 stimuli drew attention to the S2 stimulus, when the number of S1 stimuli was increased, attention to the variability of the S1 stimuli should result in fewer anticipatory errors. That is, the results of increasing attention to the S1 stimuli should be similar to what Santos et al. (2019) and Zentall, Andrews, et al. (2020) found when

the probability of reinforcement for correct choice of the S2 stimulus was reduced to 20%.

Finally, if in the Zentall, Peng, et al. (2020) experiment the pigeons were learning to inhibit choice of the multiple S2 stimuli as McMillan et al. (2015) and Smith et al. (2018) proposed, a very different prediction would follow. According to the learning-to-inhibit hypothesis, learning the simultaneous discrimination requires pigeons to learn to inhibit responding to the S2 stimulus during the first 40 trials and learning to inhibit responding to the S1 stimulus during the last 40 trials. If pigeons had to learn to inhibit responding to four S2 stimuli during the first 40 trials, that should increase the number of anticipatory errors in the Zentall, Peng, et al. experiment, which it appeared to do. One would expect, however, that when the number of S1 stimuli was increased, as it was in the Zentall, Halloran, et al. (submitted) experiment, it should have little effect on anticipatory errors, but it should result in a greater number of perseverative errors. And that is exactly what Zentall, Halloran, et al. found (see Fig. 12).

The finding that variability produced by multiple S2 stimuli results in increased anticipatory errors, while variability in the S1 stimulus result in increased perseverative errors, suggests that with the midsession reversal, pigeons may learn to reject the S- stimulus, rather than select the S+. Furthermore, it also may be that in a simple simultaneous discrimination, pigeons learn to reject the S- stimulus, rather than select the S+ stimulus. There is, in fact, some support for this hypothesis (Blackmore, Temple, & Foster, 2016; Mandler, 1970; Navarro & Wasserman, 2019; Newport, Wallis, Temple, & Siebeck, 2013; but see Mandler, 1973; Scienza, Pinheiro de Carvalho, Machado, Moreno, Biscassi, & de Souza, 2019).

It may be that procedural differences between the midsession reversal task and the simple simultaneous

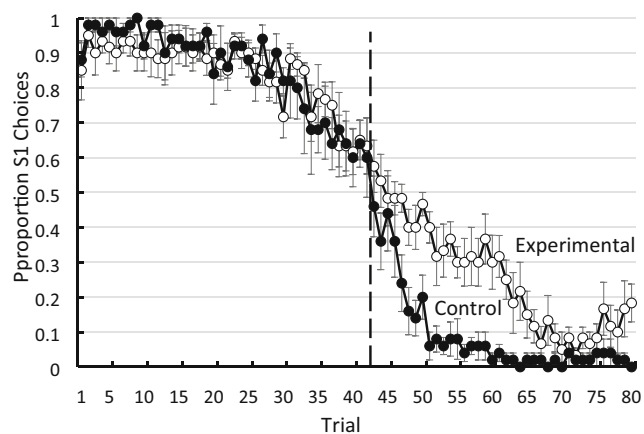


Fig. 12 Performance of pigeons on the midsession reversal task in which S1 was correct for the first 40 trials and S2 was correct for the last 40 trials. For the experimental group there was one S2 stimulus on all trials but the S1 stimulus changed from trial to trial. For the control group there was a single S1 and a single S2 stimulus on all trials (after Zentall, Halloran, & Peng, 2020)

discrimination may account for the mixed results with the simple simultaneous discrimination. Unlike in the midsession reversal experiments, in which many sessions of training are provided, in the simple simultaneous discrimination, tests of select and reject control are typically conducted by replacing either the S+ or S- and determining the degree to which the discrimination is disrupted. It is possible, however, that in some cases the tendency to avoid novel stimuli (neophobia to the replacement stimulus) played a role in the failure to find support for learning to reject the S-. That is, when the S- was replaced, neophobia could have caused subjects to select the original S+ because it was familiar; however, when the S+ was replaced, it could have caused subjects to fail to reject the original S- because they were avoiding the novel S+. Whether animals learn to select the S+ and/or reject the S- in a simple simultaneous discrimination remains an open question, but the results of research with the midsession reversal task suggest that these effects warrant further study.

Conclusions

This line of research started with what appeared to be a relatively simple question: Can organisms learn a simple simultaneous discrimination in which, in each session, one stimulus is correct for the first half of the session and the other stimulus is correct for the last half of the session. With practice, humans and rats appear to learn the task efficiently, developing a win-stay/lose-shift strategy, thereby making close to a single error each session. Pigeons, however, appear to attempt to time when the reversal will occur. By doing this, they make many anticipatory errors, beginning to choose the stimulus that will be correct prior to the reversal, as well as many perseverative errors, continuing to choose the stimulus that was correct prior to the reversal. That pigeons are using the passage of time as a cue to reversal can be readily shown by shortening or lengthening the intertrial interval. Attempts to discourage timing, by making unpredictable the point in the session at which the reversal occurred, proved to be relatively unsuccessful.

But why were the pigeons using the passage of time, a relatively poor cue, rather than the feedback from their choice on the preceding trial? Under the assumption that the pigeons were having difficulty remembering both the stimulus that they had selected on the previous trial and the outcome of that choice, pigeons were provided with reminder cues during the intertrial interval. The effect of these reminder cues was to increase accuracy, but it did not result in the development of win-stay/lose-shift performance.

Apparently, competition between the two stimuli in the region around the reversal was responsible for the large number of anticipatory and perseverative errors. When competition between the two stimuli was reduced by devaluing the stimulus that was to be correct during the last half of the

session (S2), near win-stay/lose-shift accuracy was found. Devaluation could take the form of (1) a reduction in the probability of reinforcement for correct choice of the S2 stimulus, (2) a smaller magnitude of reinforcement for correct choice of S2 relative to correct choice of S1, or (3) requiring a higher response ratio to the S2 stimulus than to the S1 stimulus.

Surprisingly, however, an attempt to lessen the competition between the S1 and S2 stimuli by increasing the number of S2 stimuli over trials did not have the same effect. Instead, it had the effect of *increasing* the number of anticipatory errors. This result raised the interesting possibility that rather than learning to select S1 during the first half of the session and learning to select S2 during the last half of the session, the pigeons were actually learning to reject S2 during the first half of the session and learning to reject S1 during the last half of the session. We tested this hypothesis by increasing the number of S1 stimuli rather than the number of S2 stimuli. In support of the hypothesis that the pigeons were learning which stimulus to reject, when there was an increase in the number of S1 stimuli, it resulted in no increase in anticipatory errors but a significant increase in the number of perseverative errors.

The results of increasing the number of stimuli in the set of S1 or S2 stimuli suggest the possibility that in a simple simultaneous discrimination, rather than learning which of the two stimuli to select, pigeons instead learn which of the two stimuli to reject. Learning to reject, rather than learning to select, may be a general feature of simple discrimination learning but it may be difficult to assess using the typical procedure in which the effect of replacement of the original S+ or S- is assessed because pigeons tend to be somewhat neophobic. In any case, variations in the performance of animals on the midsession reversal task have generated important hypotheses, not only concerning how animals learn discrimination reversals, originally thought to provide a measure of cognitive flexibility, but also how animals learn simple discriminations.

Open Practices Statement All the data and materials are available from the first author at zentall@uky.edu

References

- Beckmann, J. S. & Young, M. E. (2007). The feature positive effect in the face of variability: Novelty as a feature. *Journal of Experimental Psychology: Animal Behavior Processes*, *33*, 72–77.
- Bitterman, M. E. (1965). The evolution of intelligence. *Scientific American*, *212*, 92–100.
- Bizo, L. A., & White, K. G. (1994). The behavioral theory of timing: Reinforcer rate determines pacemaker rate. *Journal of the Experimental Analysis of Behavior*, *61*, 19–31.
- Bizo, L. A., & White, K. G. (1995). Biasing the pacemaker in the behavioral theory of timing. *Journal of the Experimental Analysis of Behavior*, *64*, 225–235.
- Blackmore, T. L., Temple, W., & Foster, T. M. (2016). Selective attention in dairy cattle. *Behavioural Processes*, *129*, 37–40.
- Cheng, K., Westwood, R. & Crystal, J. D. (1993). Memory variance in the peak procedure of timing in pigeons. *Journal of Experimental Psychology: Animal Behavior Processes*, *19*, 68–76.
- Cook, R. G., & Rosen, H. A. (2010). Temporal control of internal states in pigeons. *Psychonomic Bulletin & Review*, *17*, 915–922.
- Humphreys, L. G. (1939). Acquisition and extinction of verbal expectations in a situation analogous to conditioning. *Journal of Experimental Psychology*, *25*, 294–301.
- Kacelnik, A., Vasconcelos, M. & Monteiro, T. (2011). Darwin's "tug-of-war" vs. starlings' "horse-racing": how adaptations for sequential encounters drive simultaneous choice. *Behavioral Biology and Sociobiology*, *65*, 547–558.
- Laude, J. R., Stagner, J. P., & Zentall, T. R. (2014). Suboptimal choice by pigeons may result from the diminishing effect of nonreinforcement. *Journal of Experimental Psychology: Animal Behavior Processes*, *40*, 12–21.
- Levine, M. (1975). *A Cognitive Theory of Learning: Research on Hypothesis Testing*. Oxford, UK: Lawrence Erlbaum.
- Macphail, E. M., & Reilly, S. (1989). Rapid acquisition of a novelty versus familiarity concept by pigeons (*Columba livia*). *Journal of Experimental Psychology: Animal Behavior Processes*, *15*, 242–252.
- Mandler, J. M. (1970). Two-choice discrimination learning using multiple stimuli. *Learning and Motivation*, *1*, 261–266.
- Mandler, J. M. (1973). Multiple stimulus discrimination learning. III. What is learned? *Quarterly Journal of Psychology*, *25*, 112–123.
- Mackintosh N. J. & Cauty, A. (1971). Spatial reversal learning in rats, pigeons, and goldfish. *Psychonomic Science*, *22*, 281–282.
- McMillan, N., & Roberts, W. A. (2012). Pigeons make errors as a result of interval timing in a visual, but not a visual-spatial, midsession reversal task. *Journal of Experimental Psychology: Animal Behavior Processes*, *38*, 440–445.
- McMillan, N., Kirk, C. R., & Roberts, W. A. (2014). Pigeon and rat performance in the midsession reversal procedure depends upon cue dimensionality. *Journal of Comparative Psychology*, *128*, 357–366.
- McMillan, N., Sturdy, C. B., & Spetch, M. L. (2015). When is a choice not a choice? Pigeons fail to inhibit incorrect responses on a go/no-go midsession reversal task. *Journal of Experimental Psychology: Animal Learning and Cognition*, *41*, 255–265.
- Mueller, Halloran, and Zentall (submitted). Midsession reversal learning by pigeons: Effect of greater magnitude of reinforcement for correct choice of the first correct stimulus.
- Navarro, V. & Wasserman, E. A. (2019, April). *Select- and reject-control in learning rich associative networks*. Paper presented at the meeting of the Comparative Cognition Society. Melbourne, FL.
- Newport, C., Wallis, G., Temple, S. E., & Siebeck, U. E. (2013). Complex, context-dependent decision strategies of archerfish. *Animal Behaviour*, *86*, 1265–1274.
- Randall, C. K., & Zentall, T. R. (1997). Win-stay/lose-shift and win-shift/lose-stay learning by pigeons in the absence of overt response mediation. *Behavioural Processes*, *41*, 227–236.
- Rayburn-Reeves, R. M., & Cook, R. G. (2016). The organization of behavior over time: Insights from midsession reversal. *Comparative Cognition and Behavior Reviews*, *11*, 103–125.
- Rayburn-Reeves, R. M., Laude, J. R., & Zentall, T. R. (2013). Pigeons show near optimal win-stay/lose-shift performance on a simultaneous-discrimination, midsession reversal task with short intertrial intervals. *Behavioural Processes*, *92*, 65–70.
- Rayburn-Reeves, R. M., Molet, M., Zentall, T. R. (2011). Simultaneous discrimination reversal learning in pigeons and humans: Anticipatory and perseverative errors. *Learning & Behavior*, *39*, 125–137.

- Rayburn-Reeves, R. M., Qadri, M. A. J., Brooks, D. I., Keller, A. M., & Cook, R. G. (2017). Dynamic cue use in pigeon mid-session reversal. *Behaviour Processes*, *137*, 53–63.
- Rayburn-Reeves, R. M., Stagner, J. P., Kirk, C. R., & Zentall, T. R. (2013). Reversal learning in rats (*Rattus norvegicus*) and pigeons (*Columba livia*): Qualitative differences in behavioral flexibility. *Journal of Comparative Psychology*, *127*, 202–211.
- Restle, F. (1962). The selection of strategies in cue learning. *Psychological Review*, *69*, 329–343. doi:<https://doi.org/10.1037/h0044672>
- Roberts, S. (1981). Isolation of an internal clock. *Journal of Experimental Psychology: Animal Behavior Processes*, *7*, 242–268.
- Romanes, G. J. (1882). *Animal Intelligence*. London: Kegan Paul, Trench, & Co.
- Santos, C., Soares, C., Vasconcelos, M., & Machado, A. (2019). The effect of reinforcement probability on time discrimination in the midsession reversal task. *Journal of the Experimental Analysis of Behavior*, *111*, 371–386.
- Scienza, L., de Carvalho, M. P., Machado, A., Moreno, A. M., Biscassi, N., & de Souza, D. G. (2019). Simple discrimination in stingless bees (*Melipona quadrifasciata*): Probing for select- and reject-stimulus control. *Journal of the Experimental Analysis of Behavior*, *112*, 74–87.
- Smith, A. P., Beckmann, J. S., & Zentall, T. R. (2017). Gambling-like behavior in pigeons: ‘jackpot’ signals promote maladaptive risky choice. *Nature: Scientific Reports*, *7*, 6625. doi:<https://doi.org/10.1038/s41598-017-06641-x>
- Smith, A. P., Pattison, K. F., & Zentall, T. R. (2016). Rats’ midsession reversal performance: The nature of the response. *Learning & Behavior*, *44*, 49–58.
- Smith, A. P., Zentall, T. R., & Kacelnik, A. (2018). Midsession reversal task with pigeons: Parallel processing of alternatives explains choices. *Journal of Experimental Psychology: Animal Learning and Cognition*, *44*(3), 272–279.
- Zentall, T. R., Andrews, D. M., Case, J. P., & Peng, D. N. (2020). Less information results in better midsession reversal accuracy by pigeons. *Journal of Experimental Psychology: Animal Learning and Cognition*.
- Zentall, T. R., Halloran, & Peng (2020). Midsession reversal learning: Pigeons learn what stimulus to avoid. *Journal of Experimental Psychology: Animal Learning and Cognition*.
- Zentall, T. R., Peng, D. N., House, D. C., & Yadev, R. (2020). Midsession reversal learning by pigeons: Effect on accuracy of increasing the number of stimuli associated with one of the alternatives. *Learning and Behavior*.

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.