

Suboptimal Choice by Pigeons: Evidence that the Value of the Conditioned Reinforcer Rather than its Frequency Determines Choice

Thomas R. Zentall · Jennifer R. Laude ·
Jessica P. Stagner · Aaron P. Smith

Published online: 19 March 2015
© Association for Behavior Analysis International 2015

Abstract In an analog of human gambling, pigeons prefer a suboptimal alternative that infrequently provides a signal that is reliably associated with reinforcement over the more optimal alternative that yields more reinforcement. We hypothesized that pigeons are sensitive to the conditioned reinforcement that accrues to the stimulus associated with the greatest probability of reinforcement, independent of the frequency of its occurrence. In the present experiment we tested that hypothesis in a between groups design. For Group 50/75, choice of the suboptimal alternative 50 % of the time resulted in presentation of a stimulus that was always associated with reinforcement. For Group 25/75, choice of the suboptimal alternative 25 % of the time resulted in presentation of a stimulus that was always associated with reinforcement. For both groups, choice of the optimal alternative always resulted in presentation of a stimulus associated with 75 % reinforcement. In support of our hypothesis, both groups chose suboptimally and at the same rate. These results suggest that the conditioned reinforcing properties of the signal for reinforcement are acquired through its reliability as a signal for reinforcement and the frequency of its occurrence is of little importance.

Keywords Suboptimal choice · Gambling · Conditioned reinforcer · Probability of reinforcement · Frequency of reinforcement · Pigeons

T. R. Zentall (✉) · J. R. Laude · A. P. Smith
Department of Psychology, University of Kentucky,
Lexington, KY 40506-0044, USA
e-mail: zentall@uky.edu

J. P. Stagner
University of Florida, Gainesville, FL, USA

Suboptimal Choice by Pigeons: Value of the Conditioned Reinforcer Determines Choice not the Frequency of Reinforcement

There is substantial evidence that pigeons prefer choices that produce discriminative stimuli over those that do not. This implies that pigeons prefer stimuli that provide “information” over those that do not even when the discriminative stimuli do not affect the probability of reinforcement. Specifically, they prefer choices that 50 % of the time result in a strong conditioned reinforcer (followed by reinforcement 100 % of the time) and 50 % of the time result in a strong conditioned inhibitor (never followed by reinforcement) over those that result in weak conditioned reinforcers (each followed by reinforcement 50 % of the time). They do so even though choice of either alternative would result in the same probability of reinforcement (Roper and Zentall 1999).

Surprisingly, pigeons also show a strong preference for an alternative associated with a lower probability of reinforcement if it involves discriminative stimuli. That is, they prefer an alternative associated with 20 % reinforcement that produces discriminative stimuli (20 % of the time a stimulus that reliably predicts reinforcement, 80 % of the time a different stimulus that reliably predicts the absence of reinforcement) over an alternative associated with two stimuli that both predict reinforcement 50 % of the time (see Fig. 1; Stagner and Zentall 2010).

Why do pigeons prefer an alternative associated with discriminative stimuli that provide an overall lower probability of reinforcement (20 %) over nondiscriminative stimuli associated with a higher probability of reinforcement (50 %)? One hypothesis is that stimuli associated with a high probability of reinforcement are better conditioned reinforcers (Dinsmoor 1983). Similarly, delay reduction theory (Fantino 1969) predicts that for any stimulus, the presence of which is associated

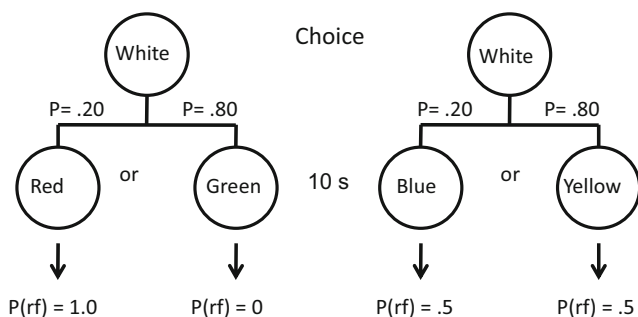


Fig. 1 Design of the Stagner and Zentall (2010) experiment. Pigeons chose between the two white keys (sides and colors were counterbalanced). Choice of the left key resulted in one of two discriminative stimuli. Choice of the right key resulted in one of two nondiscriminative stimuli. Pigeons preferred the left key although choice of the right key resulted in 2.5 times as much reinforcement

with a reduction in the delay to reinforcement (relative to its absence), will become a conditioned reinforcer.

Although it is clear that a conditioned stimulus that reliably predicts reinforcement should be preferred over a stimulus that predicts reinforcement only 50 % of the time (Stagner and Zentall 2010), it appears that the stimulus that was never associated with reinforcement and that occurred on 80 % of the trials when that alternative was chosen did not show compensating conditioned inhibition that should have countered the effect of the conditioned reinforcer.

If this analysis is correct, then it is the value of the conditioned reinforcer after the choice is made, rather than the frequency of reinforcement associated with that choice, that determines whether the pigeons will choose suboptimally. Differences in the *value* of the conditioned reinforcers associated with the two alternatives, rather than their frequencies, may explain differences in the results of these experiments.

When Gipson et al. (2009) gave pigeons a choice between a suboptimal alternative that provided a signal for reinforcement with a probability of 50 %, a signal for nonreinforcement with a probability of 50 %, and an optimal alternative that always provided a signal for 75 % reinforcement, they found a moderate preference for the suboptimal alternative. However, when Stagner and Zentall (2010) gave pigeons a choice between a suboptimal alternative that provided a signal for reinforcement with a probability of 20 % and provided a signal for nonreinforcement with a probability of 80 %, and an optimal alternative that always provided a signal for 50 % reinforcement; they found a much stronger preference for the suboptimal alternative. If one considers the relative difference in reinforcement associated with the two alternatives in the two experiments, the difference in results is difficult to explain. In the Gipson et al., study the optimal alternative provided 50 % more reinforcement than the suboptimal alternative, whereas in the Stagner and Zentall study the optimal alternative provided 150 % more reinforcement than the suboptimal alternative. Yet, preference for the suboptimal was

much stronger in the Stagner and Zentall study. However, if one considers the relative value of the conditioned reinforcers in the two studies, it may account for the difference in their results. In the Gipson et al., study, the conditioned reinforcer associated with the optimal alternative predicted reinforcement 75 % of the time, whereas the conditioned reinforcer associated with the suboptimal alternative predicted reinforcement 100 % of the time. Thus, the conditioned reinforcer associated with the suboptimal alternative was only somewhat more reliable (100 %) than the conditioned reinforcer associated with the optimal alternative (75 %). However, in the Stagner and Zentall study, the conditioned reinforcer associated with the optimal alternative predicted reinforcement only 50 % of the time, whereas the conditioned reinforcer associated with the suboptimal alternative predicted reinforcement 100 % of the time, a larger difference and, hence, a larger preference for the suboptimal alternative was found. Further support for the importance of the value of the conditioned reinforcer over its frequency was reported by Stagner et al. (2012) who found that pigeons did not prefer a conditioned reinforcer that occurred 20 % of the time over a conditioned reinforcer that occurred 50 % of the time, as long as when they occurred, they predicted reinforcement equally.

If it is the value of the conditioned reinforcer and not its frequency that determines the preference for the alternative that produces it, an interesting prediction follows. If a procedure is used in which the probability of reinforcement associated with the conditioned reinforcers that follow choice of the two alternatives is the same, then given the choice of either alternative, the probability of the occurrence of those conditioned reinforcers should be relatively unimportant.

Thus, the purpose of the present experiment was to further test the prediction that it is the value of the conditioned reinforcer that follows choice rather than the frequency of reinforcement associated with the choice itself that determines preference for the alternative. In the present experiment, we used a between groups design and varied the probability of reinforcement associated with the suboptimal alternative, while holding the probability of reinforcement associated with the optimal alternative constant. In the present experiment, two groups of pigeons were presented with a choice that involved a suboptimal option, which led to one of two stimuli, each associated with a different percentage of reinforcement (discriminative), and an optimal option which led to one of two stimuli, both of which were associated with the same percentage of reinforcement (nondiscriminative). Thus, for Group 25/75 choice of the suboptimal alternative meant that there was a 25 % chance of getting the conditioned reinforcer associated with 100 % reinforcement (and a 75 % chance of getting the stimulus associated with the absence of reinforcement), whereas for Group 50/75, choice of the suboptimal alternative meant that there was a 50 % chance of getting the conditioned reinforcer associated with 100 % reinforcement

(and a 50 % chance of getting the stimulus associated with the absence of reinforcement). For both groups, choice of the optimal alternative meant that they would always get a stimulus associated with 75 % reinforcement.

With this procedure, for Group 25/75, choice of the optimal alternative was associated with 200 % more reinforcement than the suboptimal alternative, whereas for Group 50/75, choice of the optimal alternative was associated with only 50 % more reinforcement than the suboptimal alternative. Thus, based on the probability of reinforcement associated with choice of each alternative, one might expect pigeons in both groups to prefer the alternative associated with optimal reinforcement. In light of earlier research, however, it might be expected that Group 50/75 would acquire a preference for the suboptimal choice alternative faster and to a higher level than Group 25/75 because for Group 25/75 there is a greater discrepancy between the optimal and suboptimal alternatives. That is, one might expect the larger discrepancy between the two alternatives for Group 25/75 would at least partially compensate for the attraction to the discriminative stimuli that follow the suboptimal choice. On the other hand, if it is the *value* of the conditioned reinforcer rather than its *frequency* that determines initial preference, as both conditioned reinforcers were associated with 100 % reinforcement, no difference in preference should be found.

Method

Subjects

The subjects were eight White Carneau pigeons that were retired breeders (8–10 years old). Throughout the experiment, the pigeons were maintained at 85 % of their free-feeding weight. They were individually housed in wire cages, with free access to water and grit, in a colony room that was maintained on a 12:12-h light:dark cycle.

Apparatus

The experiment was conducted in a standard three-key BRS/LVE (Laurel, MD) sound-attenuating standard operant test chamber (see Gipson et al. 2009, for details).

Procedure

Pretraining Each pigeon was first trained to peck each of six stimuli (red, yellow, green, blue, and a vertical- and horizontal-line orientation) for reinforcement on the left and right response keys. Each pigeon then was trained to peck each of four colors (red, green, blue, and yellow) on the left and right side keys on a fixed-interval 10-s schedule (the first response after 10 s was reinforced).

Discrimination Training: Group 50/75 Four pigeons were trained on the 50 % vs. 75 % discrimination (see Fig. 2a). On forced trials, either a vertical or horizontal line orientation was presented on either the left or right side key. The other side key remained dark. One peck to the illuminated key initiated a 10-s colored stimulus of fixed duration on the same key.

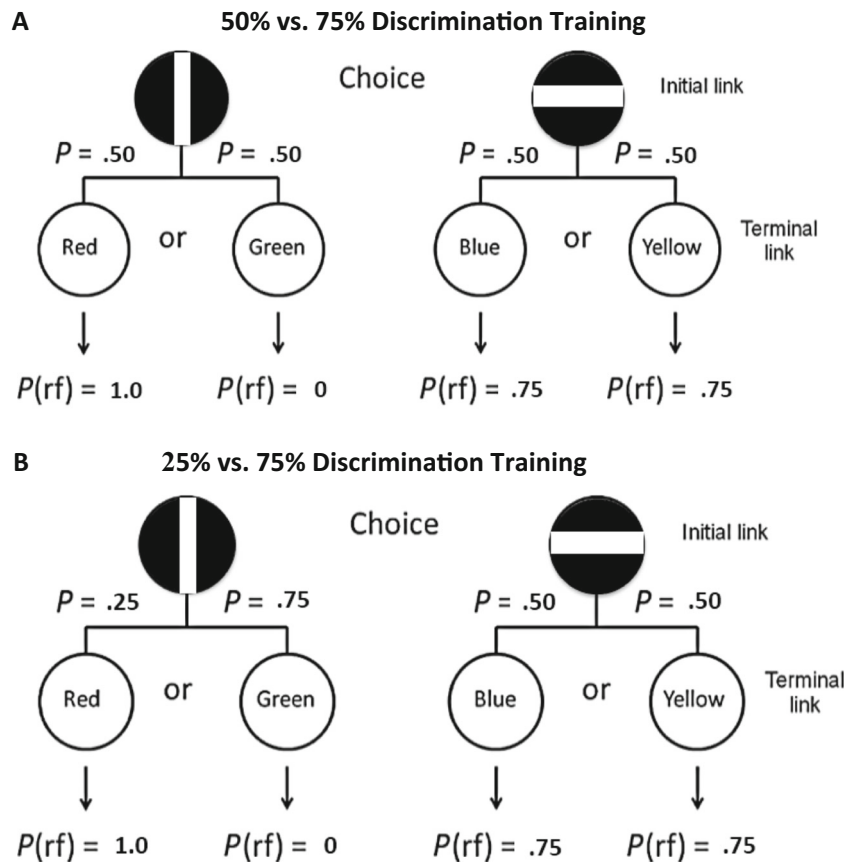
If the vertical stimulus was presented, for example, this indicated that the discriminative-stimulus schedule would be in effect. On 50 % of these trials, a peck to the vertical stimulus resulted in its replacement by the green stimulus, for example, and after 10 s, noncontingent reinforcement was provided. On the remaining 50 % of the trials, a peck replaced the vertical stimulus with the red stimulus, and after 10 s, the trial ended without reinforcement. Thus, choice of that alternative resulted in reinforcement 50 % of the time. There were 32 forced trials involving the vertical line initial stimulus in each session.

On the remaining forced trials, a horizontal line orientation was presented which indicated that schedules of reinforcement associated with nondiscriminative stimuli would be in effect. Thus, on 50 % of the trials, a peck replaced the horizontal line with the yellow stimulus and after 10 s, noncontingent reinforcement was provided with a probability of 75 %. On the remaining 50 % of the trials initiated by a horizontal line, a peck replaced the horizontal line with the blue stimulus and after 10 s, noncontingent reinforcement was provided with a probability of 75 %, as well. Thus, for the second alternative, reinforcement occurred 75 % of the time, regardless of the stimulus presented. There were 32 forced trials involving the horizontal line initial stimulus in each session.

On choice trials, the pigeons were presented with a choice between the vertical line orientation associated with 50 % reinforcement (the discriminative stimulus alternative) and the horizontal initial link associated with 75 % reinforcement (the nondiscriminative stimulus alternative). The location (left or right) of the vertical and horizontal lines was randomly determined. There were 32 choice trials in each session. The two kinds of forced trials and the choice trials were randomly interspersed throughout each session.

Discrimination Training Group 25/75 The remaining four pigeons were trained on the 25 % vs. 75 % discrimination (see Fig. 2b). The procedure was similar to that for the first four pigeons, with the exception that the discriminative stimulus alternative was associated with 25 % reinforcement. Specifically, the green stimulus appeared on 25 % of the vertical line initiated trials and after 10 s was always followed by food whereas the red stimulus appeared on 75 % of those trials and was never followed by food. Again, there were 32 forced trials involving each of the line orientation stimuli in each session.

Fig. 2 **a** Group 50/75: One alternative, associated with 50 % reinforcement, led to a signal for reinforcement on 50 % of the trials and a signal for the absence of reinforcement on 50 % of the trials (overall 50 % reinforcement). The other alternative led to a signal for 75 % reinforcement. **b** Group 25/75: One alternative, associated with 25 % reinforcement, led to a signal for reinforcement on 25 % of the trials and a signal for the absence of reinforcement on 75 % of the trials (overall 25 % reinforcement). The other alternative led to a signal for 75 % reinforcement



Once again, on choice trials, the pigeons were presented with a choice between the vertical line associated with 25 % reinforcement (the discriminative stimulus alternative) and the horizontal line associated with 75 % reinforcement (the nondiscriminative stimulus alternative). Again, there were 32 choice trials in each session. Again, the two kinds of forced trials and the choice trials were randomly interspersed throughout each session.

In both procedures, the contingencies associated with the two initial links (vertical and horizontal lines) were counterbalanced, as were the colors associated with the discriminative stimulus and nondiscriminative stimulus alternatives. For both groups each session involved a total of 96 trials. There were 25 sessions of training.

Results

The pigeons in both groups were close to indifferent between the two alternatives on the first training session but they showed a tendency to choose optimally over the first few sessions. The acquisition data for both groups is presented in Fig. 3. Although the difference from chance (50 %) did not reach statistical significance at the pigeons' lowest point of suboptimal choice, Session 4, $t(7)=1.65$, $p=0.14$, in other

experiments we have found similar tendencies early in training to choose optimally (Laude et al. 2014a, b; Zentall and Stagner 2011a, b). A two-way mixed effect ANOVA conducted on the data indicated that the main effect of sessions was significant, $F(24, 144)=4.57$, $p < 0.001$, but neither the effect

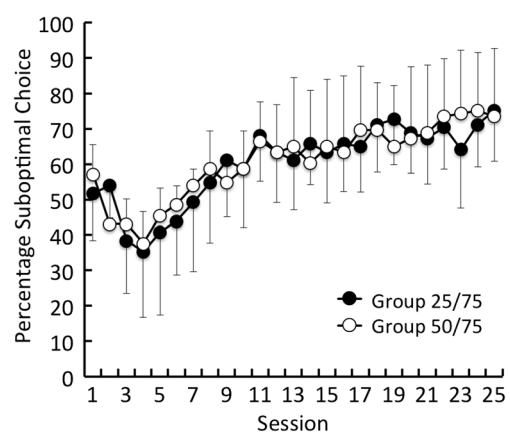


Fig. 3 Acquisition of the preference for an alternative that led to 25 % reinforcement with discriminative stimuli over an alternative that led to 75 % reinforcement with nondiscriminative stimuli (Group 25/75) and a preference for an alternative that led to 50 % reinforcement with discriminative stimuli over an alternative that led to 75 % reinforcement with nondiscriminative stimuli (Group 50/75). Data plotted are the means for each group. Error bars represent standard error of the mean

of group, $F(1, 6)=.003$, $p=0.958$, nor the Group \times Session interaction, $F(24, 144)=0.22$, $p=1.00$ was significant. Because there was no significant Group effect, the data from the two groups were pooled for further analyses. As can be seen in Fig. 3, after an initial dip in suboptimal choice, both groups began to choose suboptimally. When the data were pooled over the last five sessions (Sessions 21–25) the pigeons showed a significant preference for the suboptimal alternative as indicated by a one-sample t -test relative to 50 %, $t(7)=2.46$, $p=0.047$. Again, the difference in preference between the two groups was not statistically significant, $t(6)=0.18$, $p=0.86$.

As the ANOVA indicated that there was not a significant group effect in the acquisition of suboptimal choice, a Bayesian inference was applied using maximum likelihood functions to assess the evidence for the null and competing hypotheses (Gallistel 2009; Rouder et al. 2009). Specifically, the null hypothesis would predict that, over the 25 sessions of acquisition, there should be no group differences (0). The competing hypothesis, however, would suggest that a statistical difference should have been found (a Type II error) with group differences greater than 0. To assess this likelihood, assuming our sample variance with a mean 0 (the null hypothesis), a small group difference of 3.66 % as measured by η_p^2 (Cohen 1988), and the minimally required group difference to have achieved a statistical effect with our data of 36.6 %, Gaussian probability distributions were fit to the group differences at each session ($\bar{x}_1 - \bar{x}_2$) in the 2×25 mixed repeated measures ANOVA. The evidence for the null distribution was then divided by the evidence for the two competing hypotheses to determine likelihood ratios. Results indicated that the ratios were 1.16×10^{21} and 29.16 in favor of the null hypothesis for the minimally required group difference and small effect, respectively. Each test showed the null hypothesis to be over 100 and 29 times more likely than the alternative hypothesis, which indicates strong and decisive evidence, respectively, for the null hypothesis (Gallistel 2009). Thus, the null hypothesis is very likely.

Discussion

The purpose of the present experiment was to test the hypothesis that the relative difference in the probability of reinforcement between two alternatives would determine the preference for the alternative with the higher probability of reinforcement. If so, because the ratio of reinforcement associated with the two alternatives for Group 25/75 was greater than the ratio of reinforcement associated with the two alternatives for Group 50/75, one might have expected Group 25/75 to show a smaller preference for the suboptimal alternative. Alternatively, if the *value* of the conditioned reinforcer determined the preference for the suboptimal alternative, one

would have expected little difference in the acquisition and terminal preference for the suboptimal alternative for the two groups. Consistent with the second hypothesis, pigeons in the two groups acquired their preference for the suboptimal alternative at the same rate and to the same level.

Interpretation of the results of this experiment must be qualified by the fact that the effect found depends on accepting the null hypothesis. Is it possible that the design of the experiment was not sensitive enough to notice the difference in procedure for the two groups? There are three reasons to believe that this is not the case. First, the design was sensitive enough to find that there was a significant overall preference for the suboptimal alternative. That is, both groups replicated the results of earlier research (Gipson et al. 2009; Laude et al. 2014a; Stagner and Zentall 2010; Stagner et al. 2012; Zentall and Stagner 2011a). Second, the ratio of the probability of reinforcement associated with the two alternatives for the two groups was quite large. That is, the optimal alternative was only 50 % more than the suboptimal alternative for Group 50/75, whereas it was 200 % more than the suboptimal alternative for Group 25/75. Thus, if there were inherent differences between the two groups in choice of the suboptimal alternative, they should have appeared. In addition, the results of the present experiment were consistent with the results of Stagner et al. (2012) which, in a within-subject design, found that pigeons showed no preference for a discriminative stimulus alternative that provided a stimulus associated with 100 % reinforcement on 50 % of the trials over a discriminative stimulus alternative that provided a stimulus associated with 100 % reinforcement on 20 % of the trials. However, in the Stagner et al., experiment, the pigeons quickly chose the suboptimal alternative when the optimal reinforcement alternative was made nondiscriminative, as it was in the present experiment. Thus, the results of the present between-groups manipulation were consistent with the results of the earlier within-subject manipulation. Finally, a Bayesian analysis of the group difference indicated strong evidence to support the null hypothesis.

One limitation of the present experiment that should be noted is the experiment was conducted for only 25 sessions. Had the experiment been extended, it is possible that differences between the groups would have appeared. However, given the amount of overlap between the two groups, an emergent group difference seems unlikely.

What is responsible for the suboptimal choice found when it is followed occasionally by a strong conditioned reinforcer? One hypothesis is that the conditioned inhibitor (the stimulus that was never followed by reinforcement) that follows choice of the suboptimal alternative on many of the trials (75 % for Group 25/75, 50 % for Group 50/75) loses its inhibitory value with extended training. Laude et al. (2014a) tested this hypothesis by following the choice of the suboptimal alternative with a color as the conditioned reinforcer and a shape as the

conditioned inhibitor. They then tested the conditioned inhibitor for its ability to reduce responding to the conditioned reinforcer by superimposing the shape on the color. When they tested the pigeons early in training, they found that the conditioned inhibitor reduced responding to the conditioned reinforcer by about 40 %, whereas later in training, when the pigeons were choosing suboptimally, the conditioned inhibitor reduced responding to the conditioned reinforcer by only 20 %. The results of the present experiment are consistent with the hypothesis that it is the value of the conditioned reinforcer rather than its frequency that accounts for the suboptimal choice. In the present experiment, the conditioned reinforcer associated with the suboptimal alternative was always associated with 100 % reinforcement, whereas the conditioned reinforcer associated with the optimal alternative was always associated with 75 % reinforcement.

In the present experiment, acquisition of the suboptimal choice was somewhat slower than that reported by Stagner and Zentall (2010) (see Fig. 1 for the design). However, in that experiment, the difference in the value of the optimal alternative (50 % reinforcement) and the suboptimal conditioned reinforcer (100 %) was greater than the difference in the present experiment (100 % - 75 %), so slower acquisition would be expected.

The fact that it is the value of the conditioned reinforcer and not its frequency that appears to determine preference for the alternative producing it may also account for the results of several experiments that have reported otherwise paradoxical findings (Belke and Spetch 1994; Fantino et al. 1979; Mazur 1996; Spetch et al. 1990, 1994). In those experiments, pigeons were given a choice between two alternatives. In the signaled condition, one alternative, 50 % of the time provided a stimulus that always predicted reinforcement and 50 % of the time provided a stimulus that always predicted the absence of reinforcement. The other alternative always provided a stimulus that predicted reinforcement 100 % of the time. That is, the suboptimal alternative was associated with 50 % reinforcement and the optimal alternative was associated with 100 % reinforcement. Surprisingly, in all of those experiments some pigeons preferred the alternative that provided 100 % reinforcement, others preferred the suboptimal alternative that provided 50 % reinforcement, and still others were indifferent between the two alternatives. Furthermore, in general, the preference for the 50 % reinforcement alternative increased with increased duration of the signal (from 5 s to 90 s). However, what appear to be individual differences in optimal versus suboptimal choice may actually have been produced by an artifact. In all of those experiments, the fact that the choice was between two alternatives that were defined solely by their spatial location may have resulted in schedule indifference but the difference was masked by a spatial preference. That is, in the absence of a schedule preference, pigeons may revert to a spatial preference. Thus, the large individual differences in

schedule preference actually may have reflected a side bias, and, taken as a whole, those results are quite consistent with the current findings.

Although all of the research described here has manipulated the probability of reinforcement rather than reinforcer magnitude, we have also found similar suboptimal choice when magnitude of reinforcement was manipulated (Zentall and Stagner 2011a). In that experiment, pigeons were given a choice between two alternatives. Choice of one alternative was followed on 20 % of the trials by a stimulus that signaled ten pellets of food (analogous to a jackpot in human gambling) and on 80 % of the trials by a stimulus that signaled no food (a loss). Choice of the other alternative was followed by one of two stimuli each of which was followed by three pellets of food (analogous to not gambling). Once again, we found that the pigeons preferred the suboptimal alternative even though they could have received 50 % more food by choosing the optimal (3-pellet) alternative (see also Laude et al. 2014a, b). Thus, in the present experiment, suboptimal choice by both groups was probably not attributable to the uncertainty of outcome (only 75 % reinforcement) associated with the optimal alternative. With regard to the magnitude of reinforcement manipulation, it would be instructive to see if similar indifference could be found when the pigeons had a choice between a 20 % chance of receiving a stimulus that predicted the delivery of ten pellets of food and a 50 % chance of receiving a stimulus that predicted the delivery of ten pellets of food.

The present findings may also have implications for human gambling behavior because there is evidence to suggest that similar mechanisms may be at play. For example, people who gamble often do not attend to the odds of winning but rather primarily consider the outcome associated with winning (Blanco et al. 2000). This may also explain why, for many gamblers, losing has little effect on the future probability of gambling (until one runs out of money). For example, there is generally a large increase in the number of lottery tickets sold when the *value* of the winning ticket increases, whereas it is not clear that variability in the *probability* of winning plays an important role in the number of tickets sold. Although one could argue that it is difficult for most of us to fully understand the concept of gambling odds, problem gamblers who should have direct experience with the relation between odds and losing do not appear to be greatly affected by that experience.

References

- Belke, T. W., & Spetch, M. L. (1994). Choice between reliable and unreliable reinforcement alternatives revisited: preference for unreliable reinforcement. *Journal of the Experimental Analysis of Behavior*, 62, 353–366.

- Blanco, C., Ibáñez, A., Sáiz-Ruiz, J., Blanco-Jerez, C., & Nunes, E. (2000). Epidemiology, pathophysiology and treatment of pathological gambling. *CNS Drugs*, *13*, 397–407.
- Cohen, J. (1988). *Statistical power analysis for the behavioral sciences* (2nd ed.). Hillsdale: Lawrence Erlbaum Associates.
- Dinsmoor, J. A. (1983). Observing and conditioned reinforcement. *The Behavioral and Brain Sciences*, *6*, 693–728.
- Fantino, E. (1969). Choice and rate of reinforcement. *Journal of the Experimental Analysis of Behavior*, *12*, 723–730.
- Fantino, E., Dunn, R., & Meck, W. (1979). Percentage reinforcement and choice. *Journal of the Experimental Analysis of Behavior*, *32*, 335–340.
- Gallistel, C. R. (2009). The importance of proving the null. *Psychological Review*, *116*, 439–453.
- Gipson, C. D., Alessandri, J. D., Miller, H. C., & Zentall, T. R. (2009). Preference for 50 % reinforcement over 75 % reinforcement by pigeons. *Learning & Behavior*, *37*, 289–298.
- Laude, J. R., Beckmann, J. S., Daniels, C. W., & Zentall, T. R. (2014a). Impulsivity affects gambling-like choice by pigeons. *Journal of Experimental Psychology: Animal Behavior Processes*, *40*, 2–11.
- Laude, J. R., Stagner, J. P., & Zentall, T. R. (2014b). Suboptimal choice by pigeons may result from the diminishing effect of nonreinforcement. *Journal of Experimental Psychology: Animal Behavior Processes*, *40*, 12–21.
- Mazur, J. E. (1996). Choice with certain and uncertain reinforcers in an adjusting delay procedure. *Journal of the Experimental Analysis of Behavior*, *66*, 63–73.
- Roper, K. L., & Zentall, T. R. (1999). Observing behavior in pigeons: the effect of reinforcement probability and response cost using a symmetrical choice procedure. *Learning and Motivation*, *30*, 201–220.
- Rouder, J. N., Speckman, P. L., Sun, D., Morey, R. D., & Iverson, G. (2009). Bayesian t tests for accepting and rejecting the null hypothesis. *Psychonomic Bulletin & Review*, *16*, 225–237.
- Spetch, M. L., Belke, T. W., Barnet, R. C., Dunn, R., & Pierce, W. D. (1990). Suboptimal choice in a percentage-reinforcement procedure: effects of signal condition and terminal link length. *Journal of the Experimental Analysis of Behavior*, *53*, 219–234.
- Spetch, M. L., Mondloch, M. V., Belke, T. W., & Dunn, R. (1994). Determinants of pigeons' choice between certain and probabilistic outcomes. *Animal Learning & Behavior*, *22*, 239–251.
- Stagner, J. P., Laude, J. R., & Zentall, T. R. (2012). Pigeons prefer discriminative stimuli independently of the overall probability of reinforcement and of the number of presentations of the conditioned reinforcer. *Journal of Experimental Psychology: Animal Behavior Processes*, *38*, 446–452.
- Stagner, J. P., & Zentall, T. R. (2010). Suboptimal choice behavior by pigeons. *Psychonomic Bulletin & Review*, *17*, 412–416.
- Zentall, T. R., & Stagner, J. P. (2011a). Maladaptive choice behavior by pigeons: an animal analog of gambling (sub-optimal human decision making behavior). *Proceedings of the Royal Society B: Biological Sciences*, *278*, 1203–1208.
- Zentall, T. R., & Stagner, J. P. (2011b). Sub-optimal choice by pigeons: failure to support the Allais paradox. *Learning and Motivation*, *42*, 245–254.