Journal of Comparative Psychology

Don’t Let the Pigeon Chair the Search Committee: Pigeons (Columba livia) Match Humans’ (Homo sapiens) Suboptimal Approach to the Secretary Problem

Walter T. Herbranson, Hunter Pluckebaum, Jaidyanne Podsobinski, and Zachary Hartzell

CITATION

Herbranson, W. T., Pluckebaum, H., Podsobinski, J., & Hartzell, Z. (2021, November 22). Don’t Let the Pigeon Chair the Search Committee: Pigeons (Columba livia) Match Humans’ (Homo sapiens) Suboptimal Approach to the Secretary Problem. *Journal of Comparative Psychology*. Advance online publication. http://dx.doi.org/10.1037/com0000304

AMERICAN PSYCHOLOGICAL ASSOCIATION

Journal of Comparative Psychology

ISSN: 0735-7036

https://doi.org/10.1037/com0000304

Don't Let the Pigeon Chair the Search Committee: Pigeons (*Columba livia*) Match Humans' (*Homo sapiens*) Suboptimal Approach to the Secretary Problem

Walter T. Herbranson, Hunter Pluckebaum, Jaidyanne Podsobinski, and Zachary Hartzell
Department of Psychology, Whitman College

The secretary problem is a notorious mathematical puzzle in which one attempts to hire the best available candidate from a pool of known size. Under specific constraints, the problem has an ideal solution, but it is difficult for humans to solve. In particular, humans generally consider too few options from the available pool and in doing so make inferior hires. Three experiments investigated pigeons' and humans' choices on a version of the secretary problem. Pigeons performed suboptimally by choosing too soon, but suffered only limited costs to their rate of earned reinforcement. Depending on the instruction set, human participants approximated either prior suboptimal human results or current pigeons' results. These results may provide some insight into what makes the problem difficult to solve and how the secretary problem connects with decisions in the real world.

Keywords: decision making, choice, optimal stopping, pigeons, humans

Social scientists often assume that decisions are made in ways that maximize gains and minimize losses (Becker, 1976; Skinner, 1938). Classic behavioral research has shown that human decision makers often do just that (Friedman & Savage, 1948; Herrnstein, 1970). However, some subsequent research also reveals that this is not a universal pattern. Decision makers sometimes fail to maximize gains and minimize losses. More importantly, these failures are not random, and many occur in specific, predictable patterns (Tversky & Kahneman, 1974; Ariely, 2010).

This lack of consistently rational behavior also seems to be the case in the world of nonhuman animals. Consider, for example, the ephemeral reward task. In this task, an animal is given a choice between two alternatives, A and B, both of which lead to an immediate reward of similar magnitude (Bshary & Grutter, 2002). However, there is an important difference between the two options. If A is selected, the trial ends following the reward, whereas if B is selected, the animal can subsequently select A and thus receive a second reward on that same trial. Although the difference constitutes an obvious advantage to selecting B, not all animals do so. Wrasse fish quickly learn to choose the option that yields two rewards, whereas chimpanzees, orangutans, and capuchins do not (Salwiczek et al., 2012). The difference is not easily accounted for by general learning ability or brain size, both of which would seem to favor primates over fish. Instead, the difference may have evolved origins: cleaner wrasses in their natural environment forage by removing parasites and other materials from the gills and mouths of other fish. Those client fish are either residents, who remain at the reef (analogous to Option A), or visitors, who are likely to swim away if not immediately engaged (analogous to Option B). Initial selection of a visitor allows for the subsequent cleaning of a remaining resident, whereas initial selection of a resident likely results in the visitor leaving before it too can be cleaned. If the task accurately captures this aspect of wrasse ecology, it makes sense that they outperform primates. Note, however, that parrots cast some doubt on this particular evolutionary explanation by choosing more like wrasses than primates (Pepperberg & Hartsfield, 2014).

Given such variations in behavior across species, it may be informative to seek out situations in which suboptimal performance has been demonstrated in multiple species or in which the performances of different species vary. For instance, gambling is a suboptimal human behavior, in that it generally (though perhaps not always) carries a negative financial expectation, yet many humans pursue it enthusiastically, to negative and sometimes disastrous effect (Clark et al., 2013). Similarly, pigeons engage in choice behavior that parallels human gambling (Zentall, 2011) and is suboptimal in some of the same ways. Thus, there may be some factor that underlies the

---

Walter T. Herbranson https://orcid.org/0000-0001-5858-0712
Jaidyanne Podsobinski https://orcid.org/0000-0002-2297-4880
Zachary Hartzell https://orcid.org/0000-0001-5795-1235

This research was supported in part by a Louis B. Perry research award from Whitman College.

Walter T. Herbranson served as lead for data curation, formal analysis, project administration, resources, software, supervision, methodology, and writing–original draft, review, and editing. Hunter Pluckebaum served in a supporting role for methodology, project administration, writing–review and editing. Jaidyanne Podsobinski served in a supporting role for conceptualization, methodology, and writing–review and editing. Zachary Hartzell served in a supporting role for conceptualization, investigation, methodology, writing–review and editing. Walter T. Herbranson and Hunter Pluckebaum contributed equally to conceptualization and investigation.

Correspondence concerning this article should be addressed to Walter T. Herbranson, Department of Psychology, Whitman College, 345 Boyer Avenue, Walla Walla, WA 99362, United States. Email: herbrawt@whitman.edu

persistence of gambling that applies to multiple species. Similarly, the sunk cost effect is an increased tendency to continue an effort after an irreversible investment of time, work, money, or some other resource has been made. Humans (Arkes & Blumer, 1985) demonstrate the sunk cost effect by, for example, choosing to attend more plays if they have paid more for the same season tickets. Pigeons, chimpanzees, and rhesus monkeys also show sunk cost effects, continuing to work toward a reward more persistently if the choice is made after an investment has already been spent (Pattison et al., 2012; Watzek & Brosnan, 2020).

Different species, however, do not always behave in the same ways, and sometimes those differences are reflected in markedly different levels of success or optimality. Humans, for example, are generally better than pigeons at delaying reinforcement, often preferring a larger delayed reinforcement over a smaller immediate reinforcement, even if the larger reinforcement is delayed by months (Green et al., 2005). Pigeons on the other hand show a much shorter delay discounting horizon, with preference for a smaller reward even if the delay to a larger one is only seconds (Vanderveldt et al., 2016).

On the other hand, pigeons quickly learn to choose optimally in the Monty Hall dilemma (MHD), whereas humans do not despite their ostentatious ape brains (Herbranson & Schroeder, 2010). The difference may be due to various cognitive traps that ensnare humans, preventing them from learning from experience, whereas pigeons may be freer to change their behavior based on trial feedback (Herbranson, 2012).

Given that there is no consistent pattern of superiority between pigeons and humans, further comparative investigation of optimality in these two species should be informative, especially in those situations where one or both tend to choose suboptimally. The *secretary problem* (SP) is an interesting candidate for further research, in that humans' approach to it has been established as consistently and stubbornly suboptimal. Though similar scenarios had been described as early as 1613 by notables such as the astronomer Johannes Kepler (see Koestler, 1960), the SP was first described in its present form in print by Martin Gardner in his *Mathematical Games* column of *Scientific American* (Gardner, 1960). The problem may be presented in many formats that all have the same mathematical structure (variously known as, for example, *the game of googol, the marriage problem*, or *the best choice problem*), but it is usually articulated as a hiring decision in which an employer wishes to hire the best available applicant for a job (such as a secretary) under the following constraints. First, the pool of applicants is finite and of a known size (n). Second, applicants vary in desirability, though the specific value or quality of the best is unknown. Third, applicants are considered one at a time and in a random sequence. Finally, after each individual applicant is considered, the employer can opt either to hire that applicant on the spot (thus ending the search immediately) or reject them forever (thus continuing the search) until reaching the last candidate in the pool. At that point, the final applicant must be hired.

The SP is compelling because it is easy to understand and flexible enough to be used to simulate real-world decisions such as choosing a mate (Todd, 1997) or hunting for an apartment (Zwick, et al., 2003). At the same time, it remains mysterious because the best solution is not readily apparent. Despite widespread interest, the SP remained unsolved for years. Many strategies are possible, yielding a range of possible levels of success (Stein et al., 2003). The most effective strategy for the problem, as articulated earlier (derived by Bruss, 1984) is to interview and reject the first n/e applicants (where n is the applicant pool size, and e is the base of the natural logarithm) and then select the first applicant that ranks higher than any that have been previously considered. This is known as a cutoff rule, and if applied, the rate of selecting the best candidate is 37%, or $1/e$, regardless of the size of the applicant pool or the range of applicants' scores.

For example, one might wish to select the best available pigeon from a sample of 10 historical columbid heroes under the constraints of the SP. Figure 1 shows 10 individual pigeons of varying quality (scale rating) in an arbitrary sequence. For a pool of this size, n/e =

Figure 1
Illustration of the Secretary Problem

Applicant								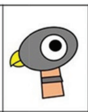	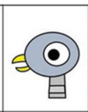	
Name	Cher Ami	William of Orange	Martha	Paddy	Mocker	GI Joe	President Wilson	Commando	Winky	Mary of Exeter
Ordinal Position	1	2	3	4	5	6	7	8	9	10
Scale Rating	57	160	61	50	65	46	202	78	82	155
Rank	8	2	7	9	6	10	1	5	4	3

Note. In this example, scale ratings are arbitrarily based on the sum of letters in a pigeon's name using an A1Z26 cipher. The optimal cutoff solution is to consider and reject the first four available options (based on ordinal position; Cher Ami, William of Orange, Martha, and Paddy), and subsequently hire the next applicant with a higher scale rating. In this case, Applicant 7, President Wilson, ranks higher with a rating of 202 than any of the first four (with a maximum of 160) and would be hired. See the online article for the color version of this figure.

$10/e \approx 3.68$, so the optimal cutoff approach is to consider the ratings of but automatically reject the first four applicants (Cher Ami, William of Orange, Martha, and Paddy) and wait for one with a rating >160 (corresponding to the rating for William of Orange, the best of the rejected applicants). In this example, that would be the seventh candidate, President Wilson, with a rating of 202. Note that President Wilson also happens to be the best in the pool, so in this case, the optimal cutoff strategy yields the best possible result. However, that is not a necessary outcome, as the best candidate could be among the first four options (e.g., if President Wilson's rating had been <160), or a better candidate could fall between the fourth ordinal candidate and the best ranked option (e.g., if Mocker or GI Joe had rated between 160 and 202).

While this cutoff solution maximizes the chance of hiring the best applicant in a given pool, human decision makers generally fail to adopt it in practice. More specifically, empirical investigations of the SP show that human decision makers select their hires too early (before the appropriate cutoff value) and in doing so make inferior hires. Seale and Rapoport (1997) asked participants to complete 100 trials of an SP. Their participants viewed information about applicants sequentially on a computer screen. Specifically, they were shown the current rank (first, second, third, etc.) of each applicant relative to those already viewed, but not a numerical rating of overall quality or an overall rank that included unviewed applicants. Their analyses considered various possible decision rules and found that participants tended to utilize a relatively simple cutoff rule, in which a decision maker selects any top-ranked candidate after a set number has been viewed and rejected. Note that the optimal strategy described earlier is a version of a cutoff rule, in which the hiring threshold is equal to the best of the first *n*/*e* candidates. Seale and Rapoport's participants however, stopped significantly earlier than that policy would dictate.

Lee (2006) used an optimal stopping scenario similar to the classical SP, in which participants were presented with a finite list of numbers and tasked with choosing the highest valued number. In this "full information" variation on the task, in which participants know the possible range of scores and are shown precise numerical scores (rather than their ordinal positions relative to other previously viewed candidates), the additional information leads to a different optimal strategy that increases the likelihood of choosing the highest number. In this case, the chance of winning is maximized by implementing a threshold that depends on the size of the applicant pool and the number of applicants remaining. Specifically, the ideal threshold decreases with each viewed candidate (Gilbert & Mosteller, 1966). With a pool of 10 applicants, this optimal strategy increases the probability of winning to approximately 60.9%, outperforming the cutoff rule described earlier. Despite the difference in how participant quality was presented and the accompanying optimal strategy, Lee's participants behaved similarly to Seale and Rapoport's: they applied a rule that stopped too soon.

In addition, Lee observed no change in strategy over time, indicating that participants did not improve on their approach with experience and feedback about wins and losses. Subsequent research has replicated the lack of learning in full information SP experiments (Campbell & Lee, 2006; Guan & Lee, 2018; Guan et al., 2020). Goldstein et al.'s (2020) participants however, showed evidence of learning and with experience approached the optimal solution. Presumably, this indicates that learning based on feedback is possible and may be dependent on specific experimental details.

Experiment 1

Although humans tend to use the wrong approach to the SP, that might not be the case for all animals. Pigeons are sensitive to rates of reinforcement (Shimp, 1973) and sometimes outperform humans on probability-based tasks (Herbranson & Schroeder, 2010). Thus, it could be informative to arrange a task in an operant chamber that matches the mathematical structure of the SP. Pigeons can be trained to associate colors with specific probabilities of reinforcement and to preferentially select colors linked to higher probabilities. Furthermore, although the solution to the SP is complex, the contributing choices are simple: Each is a decision to select or reject an immediately available option. Thus, we trained pigeons to view sequences of up to ten color stimuli, with different colors each linked to a specific probability of reinforcement. Birds made a decision after seeing each color in a sequence to either accept or reject it, and subsequent reinforcement was dependent on the selected color. In this task then, pigeons were faced with a self-terminating series of choices between a specific probability (associated with the displayed color) and another, yet unseen possibility.

Birds' performance on this task can be compared to various benchmarks corresponding to the strategies described earlier. Following studies of human choices on the SP by Lee et al (2004), we outline three specific strategies. Given that the colors birds see represent specific probabilities of reinforcement and do not indicate their ordinal rank in a pool, the task is most similar to the "full information" version of the SP (Lee, 2006). Thus, the best approach birds could use is the optimal decreasing threshold strategy, in which they adjust their selection threshold with each subsequent option. Gilbert and Mosteller (1966) indicate that with a pool of 10, the optimal threshold strategy would lead to a 60.9% chance of winning.

A second strategy worth considering is a simpler approach using a threshold that remains constant across choices rather than decreasing with each option. With a sample of 10 ranging in quality from .0 to 1.0, the most effective fixed threshold is .85, and this approach would win on about 51.7% of trials (Gilbert & Mosteller, 1966). Note that the considerably simpler approach leads to a relatively small decrease in win rate.

A third possibility is a cutoff strategy, in which one selects any best option after a fixed number has been considered. Recall that the most effective application of a cutoff requires that one reject the first *n*/*e* options and subsequently select the next that ranks highest. With a pool of 10, the most effective version of a cutoff strategy is to automatically reject the first four options and subsequently select any color associated with a greater rate of reinforcement than the maximum of those previously viewed and rejected. Using this approach, and a pool of 10 options, one would be expected to "win" by selecting the best option 39.9% of the time (Gilbert & Mosteller, 1966).

Although the strategies described earlier are complex enough to elude human participants, and thus one might be skeptical that pigeons would adopt any of them, even with experience, keep in mind that even if birds do not fully conform, they might still skew away from randomness and in a direction that increases their chance of winning. Thus, a fourth, random strategy is another useful comparison as a lower boundary. Consider, for example, a completely ambivalent bird that makes each choice arbitrarily, with no preference between selecting and rejecting. Such a bird would have a .5 chance of selecting and a .5 chance of rejecting any displayed option. Under the rules of this experiment, that strategy would result in a mean of two options considered before a selection and a 10% chance of selecting the best option from the pool.

If pigeons learn the probabilities of reinforcement associated with each color and use that information appropriately, we expect that they will do better than the random model and select the best option on more than 10% of trials. Pigeons could make optimal use of the available information, in which case they should match the optimal model and select the best option in about 61% of trials. In between these two extremes are other possibilities. As two specific examples, pigeons might use a fixed threshold strategy or an *n*/*e* cutoff strategy, in which case they would select the best option in up to 51% and 40% of trials, respectively.

Method

Methods for Experiment 1 were reviewed and approved by the Whitman College Institutional Animal Care and Use Committee.

Animals

Eight White Carneau pigeons (*Columba livia*) were purchased from Double T Farm (Glenwood, IA). Birds were maintained at approximately 80% of free-feeding weight. One bird stopped responding during the early stages of the experiment and did not complete pretraining. Thus, all presented analyses are based on the remaining seven birds.

Apparatus

Sessions were run in three-key operant chambers (BRS/LVE, Laurel, MD), each interfaced to a personal computer via USB I/O relays (Acces I/O, San Diego, CA) and controlled by code written in Visual C++ and compiled in Visual Studio (Microsoft, Redmond, WA).

Pretraining

Pigeons were pretrained in sessions consisting of habituation, hand shaping, and autoshaping until they responded consistently on each of three response keys. Stimuli during autoshaping were back-projected key lights in one of six colors: white, red, green, blue, cyan, and magenta.

Following autoshaping, pigeons were trained with different probabilities of reinforcement for pecking different colored key lights. Each of five colors (red, green, blue, cyan, and magenta) was assigned a different probability of reinforcement (.2, .4, .6, .8. or 1.0), with specific color-probability assignments unique for each bird. On individual trials, a color was back-projected onto one of the three keys (color and location each randomly selected on each trial), and after a single peck to the lit key, that key was immediately darkened and reinforcement was presented with the associated probability. Reinforcement consisted of an approximately 3-s access to mixed grain from the chamber's food hopper. Access times varied slightly from bird to bird to maintain individual running weights.

After 42 days of probability training, adequacy of learning was assessed by a choice procedure over 6 days. On each of 100 daily trials, two different colors were randomly selected from the training set and presented simultaneously on two of the three response keys (key locations were randomly selected and the third key remained dark). After a single peck to one of the illuminated keys, all keys were darkened, and reinforcement was presented according to the probability associated with the chosen color (even if it was the color linked to the lower of the two probabilities). Birds chose the higher probability color on 77.0% of trials (95% confidence interval [CI] [74.0, 80.0]). That level is greater than a chance performance of 50% $t(6) = 21.90$, $p < .001$, $d = 8.27$, $BF_{10} > 100$, indicating that birds had learned to preferentially peck keys with higher probabilities.

Procedure

Following pretraining, an SP-inspired procedure began. Birds were run 5 days per week for 100 days, and each daily session consisted of 80 trials. Each trial was either a choice trial or a maintenance trial ($p = .5$ for each, randomly determined on a trial-by-trial basis and described in the following text). Following completion of the final trial in a daily session, all keys were darkened, the house light was turned off, and a bird was returned to its home cage. Over the first 60 days, some experimental parameters varied, such as required viewing time, number of daily maintenance trials, and intertrial interval (ITI) duration. Thereafter, all details were set as described in the following text and remained unchanged for the remainder of the experiment.

Choice Trials. At the beginning of each choice trial, the computer generated a sequence of 10 colors (described in the following text). Subsequently, those colors were presented one at a time with a required decision to select or reject each, until one was selected. Following the selection of a color, a pigeon could receive access to mixed grain according to the relevant probability.

Stimulus Generation. At the beginning of each choice trial, the computer pseudorandomly generated a pool of 10 color options, each associated with one of five probabilities of reinforcement (p = .2, .4, .6, .8, or 1.0). Each color (red, green, blue, cyan, or magenta) was linked to a different specific probability, using the same assignments each bird saw during pretraining. Random generation of the sample was constrained in the following way: First, a superior option was randomly generated with $p = .25$ for each of .4, .6, .8, or 1.0. Then, the other nine options were pseudorandomly generated from all possible values with the constraint that each must have a lower probability of reinforcement than the initially generated superior option for that trial. Because there were five colors/probability values and a pool size of 10, some values appeared more than once in the pool generated for a given trial. The best option, however, was always unique (e.g., if the best option was associated with a probability of .8, only one such instance was in the pool and all nine of the others had a p less than .8). The sample was then shuffled into a random sequence, with the "best" equally likely to occur in any of the 10 ordinal positions. Figure 2 shows a possible sequence of colors and associated probabilities for one trial. The relative ranks and the ordinal location of the best option parallel those of the candidates shown in Figure 1, but following the described constraints of Experiment 1.

Trial Structure. Figure 3 depicts the structure of a trial and is based on the sample and sequence shown in Figure 2. Each trial began with a cross displayed on each of the three keys. Each cross indicates that the next available option can be viewed by pecking the cross (Figure 3, top row). The first peck to any key immediately replaced the displayed cross on that key with the first color in the generated pool (Figure 3, second row). A pigeon could then choose

Figure 2
Example Applicant Pool in Experiment 1

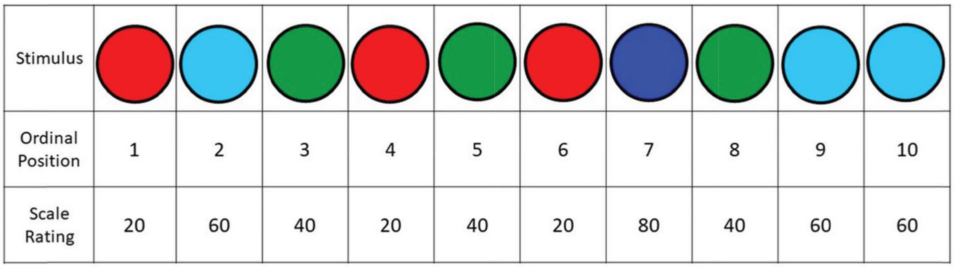

Stimulus	🔴	🩵	🟢	🔴	🟢	🔴	🔵	🟢	🩵	🩵
Ordinal Position	1	2	3	4	5	6	7	8	9	10
Scale Rating	20	60	40	20	40	20	80	40	60	60

Note. Color options have the same relative ranks as those of the candidates depicted in Figure 1. See the online article for the color version of this figure.

to see another option by pecking one of the remaining crosses. If so, the color associated with the next option immediately replaced the pecked cross. At the same time, the color linked to the previously available option disappeared and was replaced by a new cross (Figure 3, third row). Alternatively, a pigeon could select the displayed option by refraining from pecking any cross for 10 s. Pecks to the colored key had no effect, but any peck to the cross on either of the other two keys initiated a new display and 10-s interval. Following 10 s with no pecks to a displayed cross, the entire display was cleared, and reinforcement was presented with the probability linked to the displayed color.

If a pigeon rejected nine options and the 10th and final one was displayed, no crosses were projected onto the remaining keys, and thus the bird was effectively forced to choose that 10th option, as no crosses were available to be pecked.

If reinforcement was obtained on a trial, the food hopper was raised, allowing the bird access to mixed grain for approximately 3 s. Finally, a 5-s ITI preceded the next trial. During the ITI, all keys were dark, and the house light remained on.

Maintenance Trials. On each day, approximately one half of trials were maintenance trials, intended to provide regular experience with each color and reinforcement probability, regardless of how birds responded during experimental choice trials. On maintenance trials, only one key was illuminated with a color (randomly chosen from the available set), and pecks to keys had no consequences. Following 10 s, reinforcement was provided according to the reinforcement probability associated with the displayed color. Thus, maintenance trials were the same as choice trials after reaching the 10th option.

Results

Results presented here summarize all trials from all seven birds over the final 10 days of the experiment.

Wins

In the secretary problem, a "win" is defined as selection of the best available option within the pool. Pigeons selected the best color available on 20.63% of trials, 95% CI [14.46, 26.81], as

Figure 3
Schematic of a Trial in Experiment 1

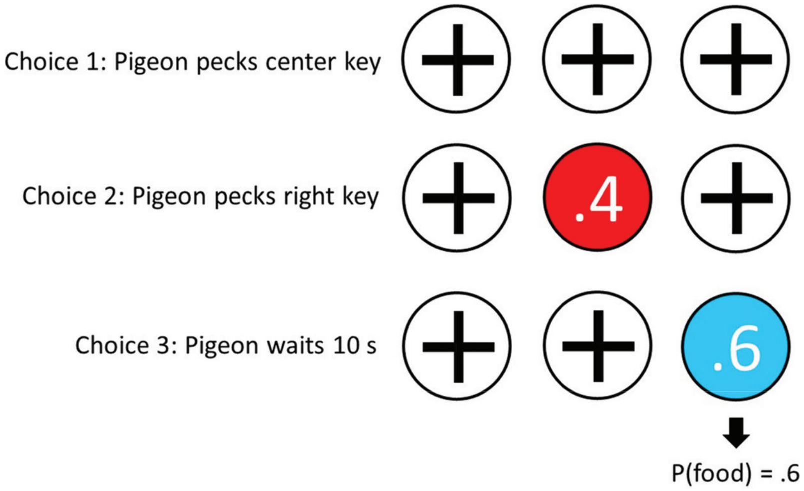

Note. The probability value on color-filled keys is for clarity in the figure, and was not displayed to birds. In this example trial, a pigeon rejected the first displayed option (red, $p = .4$) by pecking the rightmost key, and selected the second (cyan, $p = .6$) by refraining from pecking either cross. See the online article for the color version of this figure.

shown in Figure 4 (left bar). For comparison, reference lines indicate the expected performance of the four possible decision strategies described earlier: optimal, threshold, cutoff, and random. Table 1 further outlines how pigeons' performance compared to the expectations for each of these strategies. As indicated by *t* tests, pigeons performed reliably better than the random strategy, but reliably worse than each of the three other strategies. Also shown in Table 1 are Bayes Factors (BF_{10}) for each, expressing the ratio of evidence for the alternative hypothesis to evidence for the null hypothesis (see Vandekerckhove et al., 2018; Wagenmakers et al. 2016) . Here, a $BF_{10} > 1$ would indicate that there is more evidence that the probability of winning was different from that of the decision strategy under consideration. In contrast, a $BF_{10} < 1$ would indicate more evidence that the probability of winning matches that of the decision strategy. Using guidelines from Jeffreys (1961), there is strong evidence that pigeons' win probability was greater than the random strategy ($BF_{10} > 10$), and decisive evidence that pigeons' win probability was less than each of the other strategies (all $BF_{10} > 100$). Thus, by the metric of wins, pigeons appear to have performed somewhere between the expectations of random selection and the mathematically derived strategies.

Number of Options Considered

Another potentially informative measure of performance is the number of options considered per trial before making a selection. Possibilities range from one to 10, and each strategy would yield a different expectation (indicated in the row headers of Table 1). Pigeons viewed a mean of 3.18 options on each trial before making a selection, 95% CI [1.63, 4.73], as shown in Figure 5 (left bar). Again, predictions for each of the four strategies under consideration are shown as reference lines, and a summary of statistical tests appears in Table 1. In this case, pigeons' selections appear to be congruent with a random model ($p > .05$; $BF_{10} = 1.1$), but not with any of the other strategies (all $p < .001$; all $BF_{10} \geq 48.2$). Specifically, pigeons consistently made their selections before the optimal, threshold, or cutoff strategies would dictate.

Reinforcers Earned

Although use of an effective strategy increases the chance of selecting the best option, neither strategy use nor selection of the best option guarantees that a pigeon will receive grain at the end of a trial. Thus, another potentially informative measure of performance is the percent of reinforced trials. Pigeons were reinforced on 46.23% of trials, 95% CI [41.95, 50.51], as shown in Figure 6 (left bar). Again, expected values for each of the four strategies under consideration appear as reference lines, and statistical comparisons to each are summarized in Table 1. In terms of reinforcers earned, pigeons appear to have performed better than a random strategy ($p = .004$; $BF_{10} = 12.5$), and not appreciably different from the cutoff strategy ($p = .18$; $BF_{10} = 0.8$), but worse than the optimal or threshold strategies (all $p \leq .002$; all $BF_{10} \geq 26.4$).

Use of Strategies

Figure 7 (left cluster of bars) shows the percentage of trials on which a pigeon's responses were consistent with each of the mathematically derived strategies. The optimal strategy for the full information SP is to use progressively decreasing thresholds. For a sample size of 10, those thresholds are [.92, .91, .89, .88, .86, .82, .78, .69, .50, .00] (from Gilbert & Mosteller, 1966). Using the closest values in the experiment, the series of thresholds becomes [1.0, 1.0, .8, .8, .8, .8, .8, .6, .6, .0]. This approximation of the optimal strategy will be used to assess performance in subsequent analyses because it consists of thresholds that match the actual choices available in the experiment. Birds responded in a manner consistent with this optimal decreasing threshold strategy on 16.37% of trials (95% CI [1.41, 31.33]). The best single-threshold strategy for a sample size of

Figure 4

Percent of Winning Trials in Experiments 1, 2, and 3

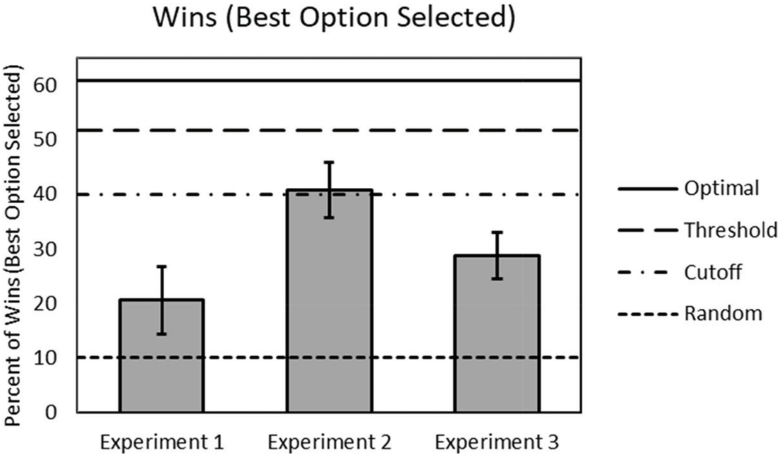

Note. Error bars depict 95% confidence intervals. A win means any trial on which the best option was selected. Reference line shows expected performance by four possible strategies: optimal (solid line), threshold (long dashes), cutoff (alternating dashes), and random (short dashes).

Table 1

Statistical Summary of Responses Consistent With Four Decision Strategies in Experiments 1, 2, and 3

	Experiment 1				Experiment 2				Experiment 3			
Strategy	$t(6)$	p	d	BF_{10}	$t(44)$	p	d	BF_{10}	$t(44)$	p	d	BF_{10}
Wins												
Optimal (60.9)	16.00	<.001	6.03	>100	8.01	<.001	1.19	>100	15.10	<.001	2.26	>100
Threshold (51.7)	12.30	<.001	4.65	>100	4.36	<.001	0.65	>100	10.80	<.001	1.61	>100
Cutoff (39.9)	7.64	<.001	2.89	>100	−0.33	.74	0.05	0.2	5.26	<.001	0.78	>100
Random (10.0)	−4.22	.006	1.59	10.6	−12.20	<.001	1.82	>100	−8.79	<.001	1.31	>100
Options												
Optimal (7.3)	6.51	<.001	2.46	59.5	11.60	<.001	1.73	>100	14.30	<.001	2.13	>100
Threshold (7.1)	6.19	<.001	2.34	48.2	10.50	<.001	1.56	>100	13.30	<.001	1.99	>100
Cutoff (7.7)	7.14	<.001	2.70	88.9	13.90	<.001	2.07	>100	16.20	<.001	2.41	>100
Random (2.0)	−1.86	.11	0.70	1.1	−18.60	<.001	2.78	>100	−10.60	<.001	1.59	>100
Reinforcers												
Optimal (62.6)	9.36	<.001	3.54	>100					5.30	<.001	0.79	>100
Threshold (55.6)	5.36	.002	2.02	26.4					2.27	.02	0.34	1.6
Cutoff (48.9)	1.53	.18	0.58	0.8					0.64	.53	0.10	0.2
Random (38.5)	−4.42	.004	1.67	12.5					−5.14	<.001	0.77	>100

Note. Reinforcement was only provided on Wins in Experiment 2. Consequently, cells corresponding to Reinforcers in Experiment 2 would not correspond to expected values in the row headers and are left blank.

10 corresponds to a threshold of .85 (Gilbert & Mosteller, 1966). In this experiment, the closest value is .8. A mean of 17.21% of trials (95% CI [12.59, 21.83]) resulted in a sequence of responses consistent with use of that threshold strategy. Finally, the cutoff strategy involves setting a threshold based on the first *n*/*e* options. With a sample size of 10, that equates to a threshold equal to the best of the first four options viewed on a given trial. A mean of 14.91% of trials (95% CI [5.07, 24.75]) resulted in a sequence of responses consistent with this cutoff strategy. Although the means are numerically close, individual birds showed different patterns: Four pigeons responded in a manner consistent with the threshold strategy more frequently than the others (maximum 24.0% of trials); two conformed to the optimal strategy more (maximum 47.7% of trials); and one matched the cutoff strategy most (26.8% of trials).

Effect of Strategies on Wins and Reinforcers Earned

There was variability among individual birds, with some performing better than others on all measures. Figure 8 summarizes data from individual birds characterizing the relationship between use of strategies and two measures of performance: reinforcers and wins. Because many trials yielded responses that were consistent with more than one of the strategies, they are combined into a single measure of strategy use: The x axis displays the percent of trials on which a bird's response was consistent with at least one of the mathematically derived strategies. Variability along the x axes indicates that some birds' responses matched at least one of the strategies more frequently than others. More importantly, the data in Figure 8 also show that the use of these effective strategies correlated positively with wins (left panel) and with reinforcers (right panel).

Figure 5

Mean Number of Options Considered per Trial in Experiments 1, 2, and 3

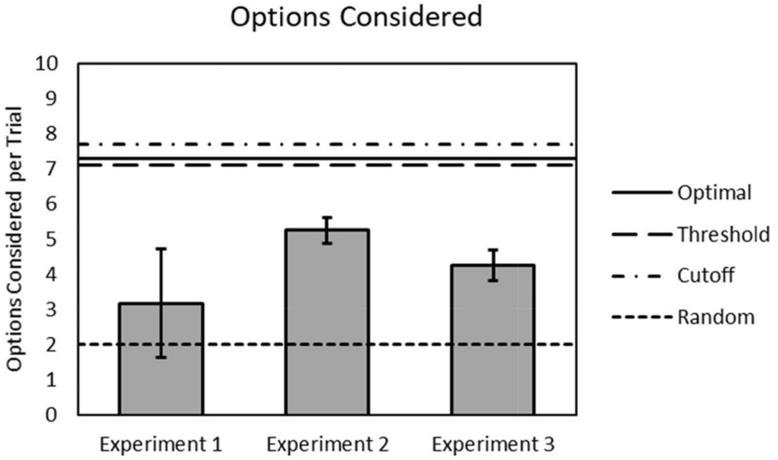

Note. Error bars depict 95% confidence intervals. Reference line shows expected performance by four possible strategies: optimal (solid line), threshold (long dashes), cutoff (alternating dashes), and random (short dashes).

Figure 6

Percent of Reinforced Trials in Experiments 1, 2, and 3

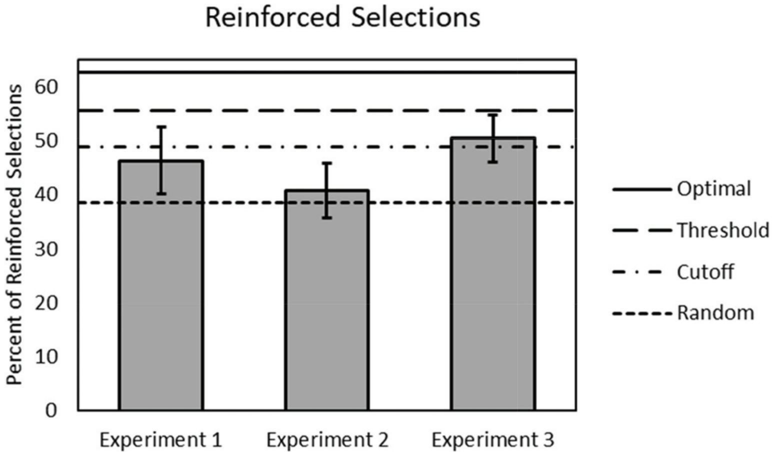

Note. Error bars depict 95% confidence intervals. Reference line shows expected performance by four possible strategies: optimal (solid line), threshold (long dashes), cutoff (alternating dashes), and random (short dashes).

Discussion

In Experiment 1, pigeons chose suboptimally by considering too few options per trial, and that matched the general pattern produced by humans in prior investigations of the SP. Pigeons' suboptimal responses may not be particularly surprising, given the complexity of the ideal solution. However, note that although suboptimal, pigeons outperformed a random response strategy in terms of wins and in doing so earned reinforcement on nearly half of the trials, approaching the expected performance of the cutoff strategy, which is the best solution under certain assumptions. These mixed results show the importance of considering the specifics of how exactly a strategy falls short of optimal. Thus, let us consider birds' performance in more detail.

First, pigeons outperformed the random model, indicating that they did learn something that increased their chances of gaining grain access, even if they still fell short of optimal. The random strategy considered here was merely the simplest model possible, using the assumption that accepting and rejecting were equiprobable, at $p = .5$ each. Thus, it serves as a comparison point for what a blindly selecting individual might do with no regard for the quality of a presented option, rather than as a model of pigeons' actual cognitive processes. Future research into what those cognitive processes might be should consider other unequal values and even

Figure 7

Percent of Trials That Were Consistent With Optimal, Threshold, and Cutoff Decision Strategies

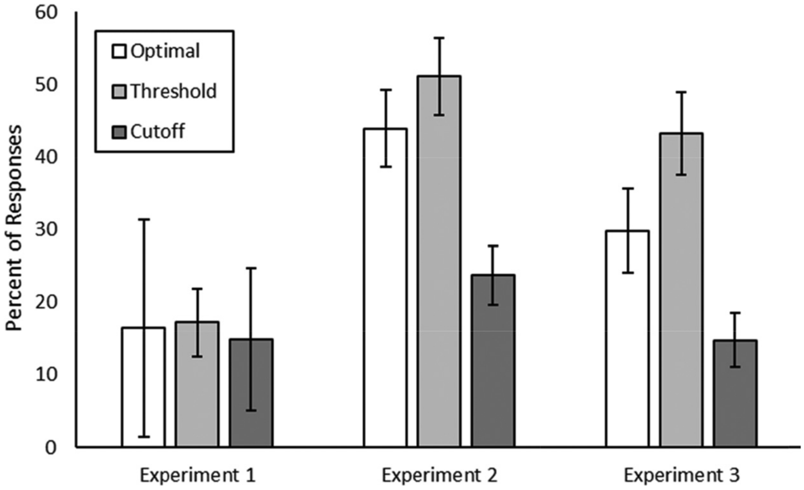

Note. Error bars depict 95% confidence intervals.

Figure 8

Scatterplots Comparing Wins (Left Panel) and Reinforced Trials (Right Panel) With Response Strategy Use in Experiment 1

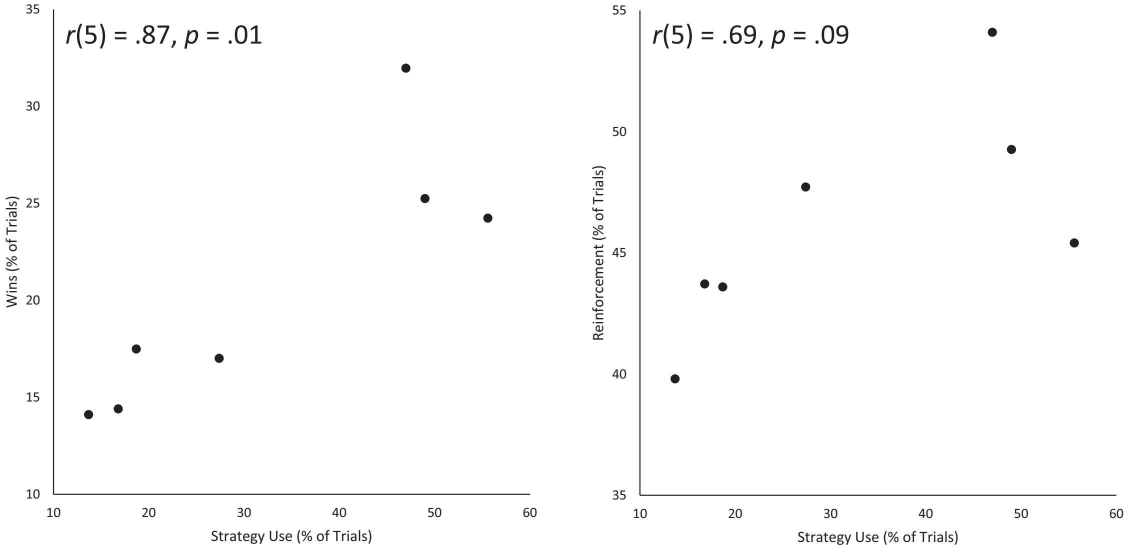

that the probabilities of accepting and rejecting might change on each trial, independent of the presented stimuli (Lee & Courey, 2021).

Birds generally made their selections too early, matching the suboptimal pattern produced by humans in prior research on the SP (Lee, 2006; Seale & Rapoport, 1997). Those results could be a consequence of a general human tendency toward risk aversion (Kahneman & Tversky, 1984). In the context of the SP, a human participant may not want to risk rejecting a quality applicant only to get stuck with an inferior one and experience regret later, especially as end of the pool approaches and there are fewer options remaining. It is plausible that pigeons were similarly swayed by high-quality options (colors linked to high probabilities of reinforcement) appearing early in a trial. Our pigeons were maintained at a fairly standard 80% of free-feeding weight (Poling et al., 1990). Presumably this motivated them to seek high-probability choices, but we should consider food deprivation level as a possible moderator. Zentall et al. (2014) demonstrated a less-is-more effect in pigeons, whereby food-deprived pigeons chose a more-preferred food item by itself over the same item plus a less-preferred item. However, this nonoptimal choice went away in another condition, in which pigeons were more food-restricted. Thus, optimal choice was at least partly dependent on motivation level. Their interpretation was that higher levels of motivation increased the value of the less preferred food item, making the two-item option relatively more appealing as starvation approached. By the same logic, the value of available options in the SP might also change along with food restriction. One might predict that birds would become more impulsive and choose even sooner under more severe food restriction.

On a related note, the possible values (i.e., reinforcement probabilities) that were available could also have influenced pigeons' choices. In this experiment, the number (five) and range (from .2 to 1.0) of options were selected with the goal of maximizing discriminability. But of course, this was only one of an infinite number of possible arrangements. For example, alternate versions could be generated in which the probabilities are all relatively low or all relatively high. A small increment in win probability from an already high chance might not carry the same weight as the same increment from a low win probability. If so, results could look quite different and might interact with the previously mentioned variable of food deprivation level. Thus, the present results should be interpreted with caution and not yet be taken as a general pattern. Nevertheless, they do indicate that at least under some conditions, pigeons' responses to an SP-based task approximately match the consistently suboptimal tendency of humans.

Another important feature of the present results is that there was considerable between-subjects variability on all measures of performance. Not all birds were equally successful, and use of effective strategies was an important determinant of success (see Figure 8). Birds that responded according to the cutoff, threshold, or optimal strategies more frequently won trials by selecting the best available option, and also gained more reinforcers. This confirms the validity of the present procedure and its link to the classic SP, in that judicious strategy use does indeed produce more wins and more earned reinforcers. Nevertheless, pigeons chose poorly, and there are many different ways to choose poorly. Pigeons' tendency overall was to choose too soon, and this particular pattern was consistent across birds. None of the pigeons reliably chose too late by considering too many options. Over the final 10 days of the experiment, pigeons saw seven or more options (the lowest expectation for the effective strategies considered) on only 12% of trials and saw all 10 options on only 5% of trials.

Perhaps the most informative result is that despite failing to use the optimal strategy, pigeons' suboptimal approach still led to frequent reinforcement. This would seem to reflect the distinction between process and outcome optimality (Dunwoody, 2009). The former refers to use of an effective strategy, whereas the latter refers to a desirable outcome, and our pigeons confirmed that the two are not always the same. Birds seem to have achieved fairly good outcomes on average, although not closely following the optimal process. They earned close to the same number of reinforcers they would have if they had adopted the optimal cutoff strategy, and only 26% less than if they had adopted the overall optimal strategy. This is likely due in part to the fact that pigeons could gain reinforcement on any trial, even if they did not choose the best available option. In other words, trials that constituted "losses" by the terminology of the SP might have been received the same way as "wins" by a lucky pigeon that was nevertheless contentedly eating from the food hopper. Similarly, selection of the best option constitutes a "win" by the rules of the SP but in many instances would have resulted in no reinforcement even though it maximized the probability of reinforcement on that trial.

Although reinforcement probability was the intended independent variable, birds' choices may also have been affected by costs relating to time. A small increment in probability of reinforcement could be outweighed by the time and effort required to sample additional options. Although 10 s were required to make any selection, and rejection led to immediate presentation of a new option, the time and effort to work through an entire pool of 10 may have been too much to invest, given that even the worst options had a nonzero probability of producing reinforcement. Seale and Rapoport (1997) similarly suggested that humans' early selections could have been due to the accumulation of search time costs. This is somewhat inconsistent with previously cited research on delay discounting, which emphasizes humans' greater willingness to tolerate a delay for a larger reward. Thus, it seems that something about the SP may encourage impatience in both humans and pigeons. Delay discounting is complex, and has many contributing factors (Vanderveldt et al., 2016), but perhaps the element that is most relevant to the SP is uncertainty. Just as a larger, later reward might not retain its value to an individual who experiences an unexpected windfall (or death) in the interim, a better, later option might not exist at all in the remaining pool of an SP experiment.

On a related note, the impulsivity pigeons show in delay discounting tasks may have been countered by the selection process in the current experiment. Recall that to select an option, pigeons had to wait 10 s without pecking either of the crosses ("reject" keys). Thus, they could not impulsively select an option or "give up" and immediately skip to the end of the trial; they were required to wait 10 s sooner or later. Though selection by omission tends to be difficult for pigeons (see Zentall & Smith, 2016), in this case, there was no requirement for birds to completely inhibit responding to a displayed color. For the most part, they did not: Pecks to colored key lights were not recorded and had no programmed consequences, but most pigeons spent the entire 10 s interval pecking the selected color. Thus, very few selections would have been due to something other than choice (such as inattention or frustration).

Given that much of the intrigue connected to the SP lies in the established tendency of humans to choose suboptimally, a solid comparison to prior human research is necessary context for Experiment 1. Our pigeons and previously studied humans performed similarly, in that both failed to consistently use the optimal strategy, and were suboptimal in the same predictable way. Given that it took mathematicians considerable time and effort to derive the optimal strategy, neither failure may be particularly surprising. However, we urge caution when comparing these pigeon data directly to previously collected human data. First, the SP is a broadly defined puzzle that accommodates considerable methodological variability. We implemented some aspects of our procedure based on available instrumentation, and those aspects did not always match the way research has been conducted with human participants. In particular, recall that the definition of a win in the SP (selection of the best option) correlated with, but did not perfectly match, reinforcement (grain access). In addition, Experiment 1 used just five values to maximize discriminability. This limited set is not incompatible with the SP but could at least be considered atypical, given that the problem usually incorporates continuously varying values and works the same way even with infinite possible values. Thus, a human version that more closely parallels those particular details of Experiment 1 would provide valuable context.

Experiment 2

Experiment 2 is a replication of Experiment 1 with human participants, using the same available pool size and probability options. To disguise the experiment from participants who may have already been familiar with the SP, it was presented as a horse race prediction game. Note that in a race, the winner is the top finishing rank, as in the original formulation of the SP. Only if a bettor chooses the first place finisher is the outcome considered a win. As in Experiment 1, there is a single winner, and ties among losers may be possible but are irrelevant. Most importantly for our purposes, the quality of options can be precisely matched to those appearing in Experiment 1, and multiple iterations can be performed, in the manner of the multiple daily races at a track, without concern for trial-to-trial carryover effects.

Although the SP has many variants that all generally yield the same pattern of results, note that some similar probability puzzlers can be significantly affected by the instruction set. The MHD, for example, when articulated differently, sometimes results in different choices. For example, Burns and Wieth (2004) recast the MHD as a competition, in which one of three boxers is superior and cannot lose. Although the underlying mathematical structure of the problem was the same as in the classical formulation, as was the optimal solution, their human participants were more likely to use the optimal approach than when presented with the classic three-door version. Thus, Experiment 2 has the added benefit of testing whether humans respond the same way when the SP is presented as an athletic competition, rather than as a hiring decision.

Method

Methods for Experiments 2 and 3 were reviewed and approved by the Whitman College Institutional Review Board prior to participant recruitment and data collection.

Participants

Forty-seven undergraduate students responded to a request for participants through the Whitman College subject pool. Some received course participation credit. Two did not finish the experiment, leaving 45 complete response sets on which the presented results are based.

Apparatus

Participants completed a single experimental session administered via Qualtrics survey software (Seattle, WA). Participants could complete the experiment on any Internet connected computer or mobile device.

Procedure

At the beginning of a session, informed consent was requested. If a participant provided their consent and opted to continue, they moved on to the experiment.

Following informed consent, each participant was randomly assigned to one of two conditions (corresponding to Experiment 2 or Experiment 3), and the appropriate program was initiated. Subsequently, instructions appeared on screen and participants responded by clicking or tapping the relevant portion of the screen.

The first trial served as practice trial and was not included in the results. On the practice trial feedback was provided as normal, and a participant could repeat the practice trial as many times as they chose.

Each session consisted of 12 trials (not including the practice trial), and each trial consisted of up to 10 decisions. At the start of each trial, the computer generated a sample of 10 numbers, drawn from a set of five possible values matching those in Experiment 1 (20, 40, 60, 80, and 100). These numbers were presented to the participant sequentially, in a scenario involving simulated horse races. The following instructions appeared on screen to set up the choices requested later:

> The following experiment will ask you to try to pick the winner of a horse race, among 10 horses. Your goal is to pick as many winners as you can from several races. On each race, you will be provided with information about an individual horse, and then asked if you think the horse will win. If you think a horse will win, you can choose that horse and then see the outcome of the race. If you do not think a horse will win, you will be provided with information about another horse in the race and asked if you think that horse will win. You can continue considering horses one at a time until you have one that you think will win, or until you have seen information about all 10 horses in the race. But be careful: you cannot go back and choose a previously considered horse. That means that if you choose to see all 10, you'll be forced to choose the 10th horse.
>
> In order to help you decide which horse to choose, I will provide you with a speed rating for each horse. Speeds range from 0 – 100, with 100 being the fastest. You should look for horses with higher speed ratings, since they will beat horses with lower ratings. However, keep in mind that each race has different horses, so speed is relative. A given race might feature many horses with low ratings, or many horses with high ratings.
>
> Let's do a practice trial that won't be counted in your results. On this trial, you should reject at least one horse before picking one, so that you have a chance to see what happens. Ready?

On each trial, one of the previously drawn sample of numbers was presented to the participant: "Horse #1 has a speed rating of [XX]. Would you like to choose Horse #1, or see another?" In place of [XX] was the numerical speed rating of the relevant horse (20, 40, 60, 80, or 100).

Each time a participant rejected a horse, the procedure moved on to the next number in the sample and displayed the same message, with a new horse number and speed rating.

After a participant chose a horse, the software announced on screen the winner of the race (always the horse from the sample with the highest speed rating), and whether or not it corresponded to the horse chosen by the participant. At the end of each trial, it also displayed the number of trials completed out of 12, and the accumulated number of wins.

Results

Wins

Participants selected the winning horse on 40.74% of trials, 95% CI [35.67, 45.81], shown in Figure 4 (middle bar). As in Experiment 1, performance can be compared to the expected win rates for four possible decision strategies, indicated by reference lines. Table 1 also summarizes statistical comparisons of mean wins in Experiment 2 with expected wins for each strategy. In Experiment 2, win rate was better than predicted by a random strategy ($p < .001$; $BF_{10} > 100$), and not appreciably different from the best cutoff strategy ($p = .74$; $BF_{10} = 0.2$), but worse than the optimal or threshold strategies (all $p < .001$; all $BF_{10} > 100$).

Compared to the pigeons in Experiment 1, the humans in Experiment 2 won a significantly larger percentage of trials, $t(50) = 3.09$, $p = .003$, $d = 1.26$, $BF_{10} = 11.3$.

Number of Options Considered

Participants viewed a mean of 5.26 horses on each trial before making a selection, 95% CI [4.91, 5.62], as shown in Figure 5 (middle bar). Table 1 summarizes comparisons with each of the four strategies, and indicates that the number of options considered was consistently greater than the random strategy ($p < .001$; $BF_{10} > 100$), and consistently less than the other three strategies (all $p < .001$; all $BF_{10} > 100$).

In comparison to the pigeons in Experiment 1, human participants viewed significantly more options per trial, $t(50) = 4.12$, $p < .001$, $d = 1.68$, $BF_{10} > 100$.

Reinforcers Earned

In Experiment 2, the only way to win was by selecting the winning horse. Thus, earned reinforcers match wins exactly, at 40.74%, 95% CI [35.67, 45.81]. For the sake of comparisons with Experiments 1 and 3, that performance is reproduced in Figure 6 (middle bar). Note that the reference lines in Figure 6 corresponding to expected reinforcers for the four strategies apply to Experiments 1 and 3, but not to Experiment 2. Thus, those comparisons are absent from Table 1.

Despite outperforming pigeons from Experiment 1 in some other ways, Experiment 2 participants' choices led to approximately the same percent of reinforced trials, $t(50) = .85$, $p = .40$, $d = 0.35$, $BF_{10} = 0.5$.

Use of Strategies

Figure 7 (middle cluster of bars) shows the percentage of trials on which responses were consistent with each of the mathematically derived strategies. The largest percent of trials were consistent with the threshold strategy ($M = 51.1$, 95% CI [45.8, 56.4]), followed by the optimal ($M = 43.9$, 95% CI [38.6, 49.2]) and cutoff ($M = 23.7$, 95% CI [19.7, 27.7]) strategies.

Discussion

In this implementation of the SP, humans outperformed the pigeons of Experiment 1 by earning wins on a larger percentage of trials. This was presumably due to the larger percentage of their responses that were consistent with the high win-probability optimal or threshold strategies. Despite that success, they employed the optimal strategy on less than half of all trials, and consistently chose too soon, just like humans in previous research on the SP, and pigeons in Experiment 1. While they performed better than the pigeons of Experiment 1 in terms of wins and options viewed, they did not earn significantly more reinforcers.

Taken by themselves, these primary results of Experiment 2 are largely consistent with previous research on the SP in humans (Lee, 2006; Seale & Rapoport, 1997). Thus, the variations introduced here in terms of pool size, absolute ranks, and instructions did not seem to change the difficulty of the problem, or the general pattern of humans' responses. As can be inferred from the myriad of ways the SP has been presented, it is apparently quite resilient in the face of modification, provided the basic structure of the puzzle remains intact.

While humans performed suboptimally by choosing too soon, they did outperform the pigeons from Experiment 1 in some ways. Thus, it is possible that humans may be better suited (though still not very well) to solving the SP than pigeons. Alternatively, their superiority might reflect necessary differences in the experimental methods used. Each trial in both experiments ended with presentation or denial of reinforcement (either grain, or a feedback message). However, it is important to acknowledge that reinforcement is not necessarily the same as a "win" in the tradition of the SP. The SP specifies a win as selection of the top option, and human participants in Experiment 2 received feedback consistent with that definition. Pigeons in Experiment 1 on the other hand worked for grain access, which was imperfectly correlated with selection of the best option. That is, they could be reinforced on trials in which they selected an inferior option, and could be denied reinforcement on trials in which they selected the best option. That was not the case in Experiment 2. Human participants were never reinforced on trials in which they did not select the best available horse (whether viewed or not), and never went unreinforced if they chose the best horse.

To their credit, our human participants in Experiment 2 won over 40% of all trials, and matched the expected wins for the cutoff strategy (the best strategy for the SP under some assumptions), even though they consistently chose too soon. We do not believe that their win rate was merely a consequence of luck, as they significantly outperformed the expectations for a random strategy. Seale and Rapoport (1997) might have predicted this when they concluded that the optimal strategy is somewhat insensitive to moderate deviations from the optimal approach: choosing a bit too early or a bit too late reduces win rate, but not always by an extreme amount. Thus, moderate variations on the best strategy can still do quite well, and were presumably reflected in our participants' relatively high win rate despite consistently premature selections.

Experiment 3

Two important measures of performance, wins and reinforced trials, were necessarily identical in Experiment 2, but not in Experiment 1. If choices are driven by reinforcers rather than wins, we might well see different results, even if other details are equal. We can easily test this possibility by reinforcing human participants in a probabilistic manner that more closely parallels the way that grain was presented to pigeons in Experiment 1. In other words, we might expect humans to perform differently (and perhaps match pigeons' performance more closely) if success on a trial is not dependent on selecting the absolute best possible option, but instead scales up with the quality of the selected option. Note that this may be more akin to the kinds of scenarios that produce reinforcement for pigeons and humans in their natural environments. A hungry pigeon or human is motivated to find the most fruitful foraging location, but can still gain food in other places, and should still be motivated to choose a better location, even if it is not the best of all possibilities.

The method from Experiment 2 can be easily modified into one in which win probability correlates with option quality, but is not solely dependent on selection of the best option. To this aim, Experiment 3 uses the same ranks as in Experiments 1 and 2 to predict a probabilistic, binary outcome: the successful hatching of an egg. The important difference is that while higher ranking options are still more likely to win (hatch), any of them can, even lower ranking ones. Similarly, a higher ranking option can lose even if it is the best available. In either case, the optimal solution remains the same, and selecting well increases the chance of reinforcement on any trial.

Method

Methods for Experiments 2 and 3 were reviewed and approved by the Whitman College Institutional Review Board.

Participants

45 undergraduate students responded to a request for participants through the Whitman College subject pool. Some received course participation credit. All completed the experiment.

Apparatus and Procedure

Apparatus and procedure were the same as in Experiment 2, with the following exceptions. On each trial, a pseudorandomly generated sample of numbers was presented in a scenario involving collection and incubation of eggs. The following instructions appeared on screen to set up the choices requested later, and paralleled those from Experiment 2:

> The following experiment will ask you to try to choose eggs that are most likely to hatch into birds. Your goal is to accumulate as many hatchlings as you can from several nests. From each nest, you will be provided with information about an individual egg, and then asked if you think the egg will hatch. If you think an egg will hatch, you can choose that egg and then see the outcome of incubation. If you do not think an egg will hatch, you will be provided with information about another egg in the nest and asked if you think that egg will hatch. You can continue considering eggs one at a time until you have one that you think will hatch, or until you have seen information about all 10 eggs in the nest. But be careful: you cannot go back and choose a previously considered egg. That means that if you choose to see all 10, you'll be forced to choose the 10th egg.
>
> In order to help you decide which egg to choose, I will provide you with a hatching probability for each egg. Probabilities range from 0 – 100. You should look for eggs with higher probabilities, since they are more likely to hatch. However, keep in mind that each nest has different eggs, so probabilities are relative. A given nest might feature many eggs with low probabilities, or many eggs with high probabilities.

> Let's do a practice trial that won't be counted in your results. On this trial, you should reject at least one egg before picking one, so that you have a chance to see what happens. Ready?

On each trial, one value from the sample of probabilities was presented to the participant. "Egg #1 has a hatching probability of [XX]. Would you like to choose egg #1, or see another?"

Each time the participant rejected an egg, the procedure moved along to the next number in the sample and displayed the same message, with an updated egg number and probability.

After a participant chose an egg, the computer randomly generated a number between 0 and 99. If the number fell below the chosen egg's rating, the computer announced that the egg had successfully hatched. If the number was higher than the egg's rating, the computer announced that it did not hatch. It also displayed the number of trials completed out of 12, and the accumulated number of hatchlings.

Results

Wins

Participants selected an egg that hatched on 28.70% of trials, 95% CI [24.42, 32.99], as shown in Figure 4 (right bar). Again, we can compare this performance to the expected win rates for four possible decision strategies, indicated by reference lines in the figure. Table 1 also summarizes statistical comparisons of mean wins in Experiment 3 with expected wins for each strategy. Win rate was better than that of the random strategy ($p < .001$; $BF_{10} > 100$), but worse than the optimal, threshold, and cutoff strategies (all $p < .001$; $BF_{10} > 100$).

Compared to those in Experiment 2, participants in Experiment 3 won fewer trials, $t(88) = 3.65$, $p < .001$, $d = 0.77$, $BF_{10} = 60.7$. However, their mean number of wins was comparable to those obtained by pigeons in Experiment 1, $t(50) = 1.46$, $p = .15$, $d = 0.59$, $BF_{10} = 0.8$.

Number of Options Considered

Participants considered a mean of 4.26 eggs on each trial before making a selection, 95% CI [3.83, 4.69], as shown in Figure 5 (right bar). Table 1 summarizes comparisons with each of the four strategies, and indicates that the number of options considered was consistently greater than the random strategy ($p < .001$; $BF_{10} > 100$), and significantly less than the other three strategies (all $p < .001$; all $BF_{10} > 100$).

Participants in Experiment 3 considered fewer options than those in Experiment 2, $t(88) = 3.64$, $p < .001$, $d = 0.77$, $BF_{10} = 57.9$. However, they considered approximately the same number as the pigeons in Experiment 1, $t(50) = 1.83$, $p = .07$, $d = 0.74$, $BF_{10} = 1.2$.

Reinforcers Earned

Participants in Experiment 3 were reinforced on 50.37% of trials, 95% CI [45.72, 55.02], as shown in Figure 7 (right bar). Table 1 summarizes comparisons with each of the four strategies, and indicates that the number of reinforcers was consistently greater than the random strategy ($p < .001$; $BF_{10} > 100$), congruent with the cutoff ($p = .53$; $BF_{10} = .2$) and possibly the threshold ($p = .02$, $BF_{10} = 1.6$) strategies, and significantly less than the optimal strategy ($p < .001$; $BF_{10} > 100$).

Participants earned more reinforcers than those in Experiment 2, $t(88) = 2.82$, $p = .006$, $d = 0.60$, $BF_{10} = 6.7$. They earned a comparable number to the pigeons in Experiment 1, $t(50) = .70$, $p = .49$, $d = 0.28$, $BF_{10} = 0.4$.

Use of Strategies

Figure 7 (right cluster of bars) shows the percentage of trials that were consistent with each of the mathematically derived strategies. The largest percent of trials were consistent with the threshold strategy ($M = 43.3$, 95% CI [37.6, 49.1]), followed by the optimal ($M = 29.8$, 95% CI [24.0, 35.7]) and cutoff ($M = 14.8$, 95% CI [11.08, 18.55]) strategies.

Discussion

Humans in Experiment 3 did not conform to the optimal strategy, similar to humans in Experiment 2, pigeons in Experiment 1, and other previous investigations of the SP. However, in some important ways, they most closely resembled the choices of pigeons in Experiment 1 and may provide some insight into the cognitive traps that make the SP a challenge.

In the classic SP a win is defined as selection of the best possible option. However, this kind of omniscient feedback is generally not available to behaving organisms in the real world. Individual decision makers are often best served by ignoring whether they obtained the best available outcome, and focusing instead on whether or not reinforcement was attained (Schwartz et al., 2002). In decision making research, this approach is sometimes called *satisficing* (Simon, 1955). Presumably this is what pigeons in Experiment 1 and humans in Experiment 3 did, as they both earned a healthy rate of reinforcement that reliably outperformed random choice. They did so by choosing options with high absolute ratings, presumably knowing that they had a high probability of gaining reinforcement regardless of the quality of options remaining in the pool. Given a task that allows for this approach, it leads to better outcomes in the currency of the real world, reinforcement. Participants in Experiment 2 more closely approximated the optimal strategy, but this was a pyrrhic victory given that they actually garnered fewer reinforcements.

Based on the results of Experiments 2 and 3, it appears that instruction set can be a determinant of success in the SP, just as it is in the MHD (Burns & Wieth, 2004). Here, we suggest that the important part of the instruction set is what determines a win: relative rank, or absolute rating. Experiment 3 is more similar to the "instructions" that organisms face in the real world: reinforcement is not perfectly predictable, but a well-selected strategy, even if not ideal, can lead to an improved rate of reinforcement. Experiment 2 is truer to the classic formulation of the SP: an unforgiving scenario that demands selection of the best option, a feat that can only be accomplished on a fraction of trials even when using the best possible approach. These are among the ecologically invalid features that make the SP a psychologically interesting puzzle in the first place. Note though that although unrealistic, it bears some similarity to another suboptimal behavior that is seen in the real world: gambling. Coincidental, but perhaps telling is that the scenario used in Experiment 2 was horse racing, the preferred hobby of many habitual gamblers (Valleur et al., 2016). Recent research has shown that "near-miss" events in a gambling simulation increase gambling propensity and activate some of the same brain areas as monetary wins (Clark et al., 2009). Thus, one can still learn from events that deny reinforcement, as presumably participants in Experiment 2 did when they selected a horse that was rated only slightly lower than the winner. Those participants more closely approximated the optimal solution, despite earning fewer reinforcers. Therefore, an important trap in the case of the SP, seems to be the desire to obtain the best possible outcome and designating other results as failures.

General Discussion

Across three experiments, pigeons and humans failed to consistently use the optimal solution to the SP. Most importantly, both species failed to do so in the same way: They made selections too soon, before considering the ideal number of options. There were, however, some meaningful patterns across experiments. If reinforcement probability scaled up with the quality of the selected option, both pigeons (Experiment 1) and humans (Experiment 3) earned reinforcement at a comparable rate. Humans could be pushed toward use of the optimal strategy by defining wins in a way that required selection of the best possible option, but doing so reduced their rate of earned reinforcement (Experiment 2). We argue that human participants in Experiment 2 outperformed pigeons in Experiment 1 because of the method of providing reinforcement, rather than by virtue of being human, because the humans in Experiment 3 behaved more like pigeons, earning more reinforcers while making fewer superior selections.

Nevertheless, there were some differences between pigeons and humans that merit consideration. Although Experiments 1 and 3 demonstrate remarkable similarity across the two species in terms of wins, options considered, and reinforcement rate (Figures 4–6), there remained a conspicuous difference in how frequently pigeons and humans responded in ways that matched the cutoff, threshold, and optimal strategies. Figure 7 shows that humans were most likely to respond in a manner consistent with the threshold strategy, followed by the optimal, and this was true in both Experiments 2 and 3. Pigeons on the other hand, show a more even distribution of response sequences among the three.

Table 2 further explores possible cross-species differences in use of threshold-based strategies. It displays the percent of responses in Experiments 1, 2, and 3 that were consistent with multiple fixed threshold and decreasing threshold strategies. The top portion of the table considers strategies with fixed thresholds corresponding to each of the available elements in the stimulus set. In each case, the strategy assumes that a participant will select the first option viewed that is equal to or greater than the specified threshold. Note that .8 is the fixed threshold strategy that would win most frequently and corresponds to the threshold strategy considered in earlier analyses. Here, there is a notable difference between the choices of the pigeons of Experiment 1 and the humans of Experiments 2 and 3. Pigeons most frequently responded in a manner consistent with a threshold of .4, a value lower than the most effective fixed threshold. Humans on the other hand, most frequently responded in a manner consistent with a threshold of 1.0, a value higher than the most effective fixed threshold.

Table 2

Percent of Responses in Experiments 1, 2, and 3 That Were Consistent With Use of Different Fixed Threshold and Decreasing Threshold Strategies

Strategy	Experiment 1	Experiment 2	Experiment 3
Fixed threshold			
.2	.369	.103	.119
.4	.433	.133	.135
.6	.327	.365	.317
.8	.172	.522	.433
1.0	.061	.575	.435
Decreasing thresholds			
Optimal	.164	.454	.298
Round up	.136	.369	.230
Round down	.250	.583	.461
Shift up	.130	.341	.219
Shift down	.302	.613	.520
Shift early	.198	.518	.365
Shift late	.150	.391	.265

Note. Decreasing threshold strategies in Table 2 correspond to the following sequences of thresholds. Optimal: [1.0, 1.0, .8, .8, .8, .8, .8, .6, .6, .0]; Round up: [1.0, 1.0, 1.0, 1.0, 1.0, 1.0, .8, .8, .6, .0]; Round down: [.8, .8, .8, .8, .8, .8, .6, .6, .4, .0]; Shift up: [1.0, 1.0, 1.0, 1.0, 1.0, 1.0, 1.0, .8, .8, .0]; Shift down: [.8, .8, .6, .6, .6, .6, .6, .4, .4, .0]; Shift early: [1.0, .8, .8, .8, .8, .8, .6, .6, .6, .0]; Shift late: [1.0, 1.0, 1.0, .8, .8, .8, .8, .8, .6, .0].

Similarly, the bottom portion of Table 2 shows the percent of choices that were consistent with the optimal strategy, as well as several variations. The optimal strategy is based on the mathematical analysis of Gilbert and Mosteller (1966), with thresholds rounded to the nearest value present in the stimulus set, as specified earlier. Two variations presented in Table 2 (round up and round down) round all thresholds to the next higher or lower value in the stimulus set, respectively. Two additional variations shifted each threshold in the optimal sequence up one value (e.g., from .4 to .6) or down one value (e.g., from .8 to .6). The final two variations (shift early and shift late) shift each threshold one position earlier (e.g., from third to second), or later (e.g., from third to fourth) in the sequence. Again, some cross-species differences are worthy of note. First, humans' responses were better described by decreasing threshold strategies than by fixed threshold strategies. Pigeons showed the opposite pattern, with fixed threshold strategies accounting for a greater proportion their choices. This suggests that pigeons may have been more inclined to choose solely based on color, without adjusting for the current position in the sequence. Note that use of a fixed threshold could simply be a consequence of associative strength, with each color acquiring a different association with reward during pretraining and maintenance trials. By this interpretation, those colors that were most frequently followed by food (.4 or greater) could have been more appealing than the alternatives ("reject keys") because those alternatives were never immediately followed by food. In contrast, the remaining color in the set may have been so weakly associated with food (.2) that it was less appealing than the alternatives and their secondary link to food via subsequently presented options.

Despite these differences, Table 2 also shows some similarities in how humans and pigeons may have used thresholds. Variations on the optimal decreasing threshold strategy that resulted in lower thresholds (round down, shift down, and shift early) all accounted for more choices than the optimal strategy, and this was true in all three experiments. In contrast, variations that resulted in higher thresholds (round up, shift up, and shift late) all accounted for fewer choices than the optimal strategy. Thus, it may be that the tendency to choose too soon resulted from the use of thresholds that were generally lower than the optimal strategy prescribes, and this applies to both humans and pigeons. Nevertheless, note that the highest values in Table 2 only account for 40% to 60% of all responses, indicating that even the best fitting strategies considered here fail to account for a large proportion of choices.

These results resonate with primary conclusions made in prior empirical investigations of the SP in humans: choices are generally made too soon, but the consequences of myopia are mitigated by the fact that the optimal solution is somewhat tolerant of improper stopping points (Seale & Rapoport, 1997). If the additional time investment to consider more options outweighs the small incremental gain from using the optimal strategy, individuals might naturally be led to select too soon, rather than too late, or even optimally. We see no reason that this reasoning would not apply to pigeons and humans equally, and our results are consistent with that interpretation. Optimal foraging theory (Stephens & Krebs, 1986) may provide additional context for why pigeons generally select too soon. It posits that animals adopt strategies that push them to both maximize their gains and minimize costs. The satisficing strategy apparently used by the pigeons in Experiment 1, although not optimal, may allow them to maximize the rate of food gained while also minimizing the time and energy expended. However, it should be noted that optimal foraging theory emerged from observations of animals' behavior in their natural environment and it should not be assumed that the same cognitive strategies that evolved from real-world foraging translate to a laboratory decision making task. The set of three experiments considered together reinforce the resilience of the SP and its optimal solution: It appears to be the same challenging puzzle (a) for pigeons and humans, (b) when presented with a limited set of numerical values, (c) whether success is determined by numerical value or relative rank, and (d) when presented as a novel scenario with the same mathematical structure, but not involving hiring from a pool of applicants.

This set of experiments brings the SP into the realm of comparative cognition as a useful method for investigating the complexities of decision making, by demonstrating suboptimal choice in multiple species (placing it alongside such procedures as the MHD and the ephemeral reward task). Although some variables that influence performance were identified, additional research should explore and further define such factors. These may include the instruction set, the size of the applicant pool, and the range and distribution of numerical scores. In addition, some constraints imposed on Experiment 1 as a starting point could be relaxed to open up more possibilities. In particular, a scalar version of the task might be possible and informative, with video displays or LEDs allowing for infinite color gradations through the range of short to long wavelengths visible to pigeons. However, note that color perception does not scale the same way in pigeons and humans (Wright & Cumming, 1971). Furthermore, most available LEDs and video monitors are calibrated to human rather than pigeon vision, and thus may miss some wavelengths that would be ideal for a pigeon implementation. Thus, selection of color values would need to be done with care, and treating color as a nominal variable as was done in Experiment 1 could still be the best approach.

This research was motivated by questions about optimality in animal behavior, and the data have provided another curious case where both humans and pigeons fail to embrace a strategy that would increase their chances of success. This resonates with some other research programs such as those on gambling and the sunk cost effect and leads to the same question of why we sometimes fail to learn from experience. Possible answers to that question might come from a comparison of the methods used in Experiments 2 and 3. Humans in Experiment 2 were better able to utilize the information available to win trials, even though they earned fewer reinforcers in the process. Thus, the important difference might lie in how wins are defined. In Experiment 2, a win required selection of the best option available, a consistent and unambiguous criterion. In Experiment 3 (and Experiment 1), success was determined probabilistically and thus chance played an important role, matching the way that consequences happen in the natural world where wise choices often go unrewarded, and foolish ones sometimes lead to success.

Although designed and conducted in the tradition of comparative cognition, the present experiments have the potential to provide some connection to and inspiration for other areas of psychology. In addition to the direct links to cognitive psychology and decision making, it might also link to important concepts from social psychology. Consider, for example, social comparison theory (Festinger, 1954), by which an individual defines their own success in relation to that obtained by relevant peers. Such comparisons can lead to escalating expectations, with the goal not just of succeeding, but of outperforming all others in the comparison group. In the terminology of the SP, social comparison norms may guide individuals away from a focus on numerical scores and toward a preoccupation with relative ranks. The current results would argue that such an approach is ill-advised, in that it can lead to a decrease in overall reinforcement rate. The SP might also connect with contemporary research on fairness. For example, monkeys (Brosnan & de Waal, 2003) and dogs (Brucks et al., 2016) respond differently after observing another conspecific receiving a more desirable outcome for performing the same behavior. These research traditions converge on the notion that decision makers are swayed not only by the value of the reinforcers they obtain, but also by the ranks of those reinforcements relative to others that are apparently available.

Finally, the results may have implications for how best to utilize pigeons' cognitive abilities for real-world applications. We propose that pigeons are best deployed on tasks that match their empirically verified talents, such as those required for radiology (Levenson et al., 2015), telecommunications (Wallraff, 2001), medical manufacturing (Verhave, 1966), game show participation (Herbranson & Schroeder, 2010), and aerial combat (Skinner, 1960). Similarly, pigeons should be discouraged from tasks that reflect their weaknesses, such as gambling (Zentall, 2011), public transportation (Willems, 2003; though see Gibson et al. 2012), and we now argue, human resources.

References

Ariely, D. (2010). *Predictably irrational: The hidden forces that shape our decisions*. Harper Perennial.

Arkes, H. R., & Blumer, C. (1985). The psychology of sunk cost. *Organizational Behavior and Human Decision Processes*, *35*(1), 124–140. https://doi.org/10.1016/0749-5978(85)90049-4

Becker, G. S. (1976). *The economic approach to human behavior*. University of Chicago Press. https://doi.org/10.7208/chicago/9780226217062.001.0001

Brucks, D., Essler, J. L., Marshall-Pescini, S., & Range, F. (2016). Inequity aversion negatively affects tolerance and contact-seeking behaviours towards partner and experimenter. *PLoS ONE*, *11*(4), Article e0153799. https://doi.org/10.1371/journal.pone.0153799

Brosnan, S. F., & De Waal, F. B. (2003). Monkeys reject unequal pay. *Nature*, *425*(6955), 297–299. https://doi.org/10.1038/nature01963

Bruss, F. T. (1984). A unified approach to a class of best choice problems with an unknown number of options. *Annals of Probability*, *12*(3), 882–891. https://doi.org/10.1214/aop/1176993237

Bshary, R., & Grutter, A. S. (2002). Asymmetric cheating opportunities and partner control in a cleaner fish mutualism. *Animal Behaviour*, *63*(3), 547–555. https://doi.org/10.1006/anbe.2001.1937

Burns, B. D., & Wieth, M. (2004). The collider principle in causal reasoning: Why the Monty Hall dilemma is so hard. *Journal of Experimental Psychology: General*, *133*(3), 434–449. https://doi.org/10.1037/0096-3445.133.3.434

Campbell, J., & Lee, M. D. (2006). The Effect of Feedback and Financial Reward on Human Performance Solving ‘Secretary’ Problems. *In Proceedings of the Annual Meeting of the Cognitive Science Society*, *28*. https://escholarship.org/uc/item/24z2s2ff

Clark, L., Averbeck, B., Payer, D., Sescousse, G., Winstanley, C. A., & Xue, G. (2013). Pathological choice: The neuroscience of gambling and gambling addiction. *The Journal of Neuroscience*, *33*(45), 17617–17623. https://doi.org/10.1523/JNEUROSCI.3231-13.2013

Clark, L., Lawrence, A. J., Astley-Jones, F., & Gray, N. (2009). Gambling near-misses enhance motivation to gamble and recruit win-related brain circuitry. *Neuron*, *61*(3), 481–490. https://doi.org/10.1016/j.neuron.2008.12.031

Dunwoody, P. T. (2009). Introduction to the special issue: Coherence and correspondence in judgment and decision making. *Judgment and Decision Making*, *4*(2), 113–115.

Festinger, L. (1954). A theory of social comparison processes. *Human Relations*, *7*(2), 117–140. https://doi.org/10.1177/001872675400700202

Friedman, M., & Savage, L. J. (1948). The utility analysis of choices involving risk. *Journal of Political Economy*, *56*(4), 279–304. https://doi.org/10.1086/256692

Gardner, M. (1960). Mathematical games. *Scientific American*, *202*(2), 150–154. https://doi.org/10.1038/scientificamerican0260-150

Gibson, B., Wilkinson, M., & Kelly, D. (2012). Let the pigeon drive the bus: Pigeons can plan future routes in a room. *Animal Cognition*, *15*(3), 379–391. https://doi.org/10.1007/s10071-011-0463-9

Gilbert, J. P., & Mosteller, F. (1966). Recognizing the maximum of a sequence. *American Statistical Association Journal*, *61*(313), 35–73. https://doi.org/10.1080/01621459.1966.10502008

Goldstein, D. G., McAfee, R. P., Suri, S., & Wright, J. R. (2020). Learning when to stop searching. *Management Science*, *66*(3), 1375–1394. https://doi.org/10.1287/mnsc.2018.3245

Green, L., Myerson, J., & Macaux, E. W. (2005). Temporal discounting when the choice is between two delayed rewards. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *31*(5), 1121–1133. https://doi.org/10.1037/0278-7393.31.5.1121

Guan, M., & Lee, M. D. (2018). The effect of goals and environments on human performance in optimal stopping problems. *Decision*, *5*(4), 339–361. https://doi.org/10.1037/dec0000081

Guan, M., Stokes, R., Vanderckhove, J., & Lee, M. D. (2020). A cognitive modeling analysis of risk in sequential choice tasks. *Judgment and Decision Making*, *15*(5), 823–850.

Herbranson, W. T. (2012). Pigeons, humans, and the Monty Hall dilemma. *Current Directions in Psychological Science*, *21*(5), 297–301. https://doi.org/10.1177/0963721412453585

Herbranson, W. T., & Schroeder, J. (2010). Are birds smarter than mathematicians? Pigeons (*Columba livia*) perform optimally on a version of the Monty Hall Dilemma. *Journal of Comparative Psychology*, *124*(1), 1–13. https://doi.org/10.1037/a0017703

Herrnstein, R. J. (1970). On the law of effect. *Journal of the Experimental Analysis of Behavior*, *13*(2), 243–266. https://doi.org/10.1901/jeab.1970.13-243

Jeffreys, H. (1961). *Theory of probability* (3rd ed.). Oxford University Press.

Kahneman, D., & Tversky, A. (1984). Choices, values, and frames. *American Psychologist*, *39*(4), 341–350. https://doi.org/10.1037/0003-066X.39.4.341

Koestler, A. (1960). *The watershed: A biography of Johannes Kepler*. Anchor Books.

Lee, M. D. (2006). A hierarchical Bayesian model of human decision-making on an optimal stopping problem. *Cognitive Science*, *30*(3), 1–26. https://doi.org/10.1207/s15516709cog0000_69

Lee, M. D., & Courey, K. A. (2021). Modeling optimal stopping in changing environments: A case study in mate selection. *Computational Brain & Behavior*, *4*(1), 1–17. https://doi.org/10.1007/s42113-020-00085-9

Lee, M. D., O’Connor, T. A., & Welsh, M. B. (2004). Decision-making on the full-information secretary problem. In K. Forbus, D. Gentner & T. Regier (Eds.), *Proceedings of the 26th annual conference of the cognitive science society* (pp. 819–824). Erlbaum.

Levenson, R. M., Krupinski, E. A., Navarro, V. M., & Wasserman, E. A. (2015). Pigeons (*Columba livia*) as trainable observers of pathology and radiology breast cancer images. *PLoS ONE*, *10*(11), Article e0141357. https://doi.org/10.1371/journal.pone.0141357

Pattison, K. F., Zentall, T. R., & Watanabe, S. (2012). Sunk cost: Pigeons (*Columba livia*), too, show bias to complete a task rather than shift to another. *Journal of Comparative Psychology*, *126*(1), 1–9. https://doi.org/10.1037/a0023826

Pepperberg, I. M., & Hartsfield, L. A. (2014). Can grey parrots (*Psittacus erithacus*) succeed on a “complex” foraging task failed by nonhuman primates (*Pan troglodytes*, *Pongo abelii*, *Sapajus apella*) but solved by wrasse fish (*Labroides dimidiatus*)? *Journal of Comparative Psychology*, *128*(3), 298–306. https://doi.org/10.1037/a0036205

Poling, A., Nickel, M., & Alling, K. (1990). Free birds aren’t fat: Weight gain in captured wild pigeons maintained under laboratory conditions. *Journal of the Experimental Analysis of Behavior*, *53*(3), 423–424. https://doi.org/10.1901/jeab.1990.53-423

Salwiczek, L. H., Prétôt, L., Demarta, L., Proctor, D., Essler, J., Pinto, A. I., Wismer, S., Stoinski, T., Brosnan, S. F., & Bshary, R. (2012). Adult cleaner wrasse outperform capuchin monkeys, chimpanzees and orang-utans in a complex foraging task derived from cleaner—Client reef fish cooperation. *PLoS ONE*, *7*(11), Article e49068. https://doi.org/10.1371/journal.pone.0049068

Schwartz, B., Ward, A., Monterosso, J., Lyubomirsky, S., White, K., & Lehman, D. R. (2002). Maximizing versus satisficing: Happiness is a matter of choice. *Journal of Personality and Social Psychology*, *83*(5), 1178–1197. https://doi.org/10.1037/0022-3514.83.5.1178

Seale, D. A., & Rapoport, A. (1997). Sequential decision making with relative ranks: An experimental investigation of the “secretary problem”. *Organizational Behavior and Human Performance*, *69*(3), 221–236. https://doi.org/10.1006/obhd.1997.2683

Shimp, C. P. (1973). Probabilistic discrimination learning in the pigeon. *Journal of Experimental Psychology*, *97*(3), 292–304. https://doi.org/10.1037/h0034134

Simon, H. A. (1955). A behavioral model of rational choice. *Quarterly Journal of Economics*, *69*, 99–118. https://doi.org/10.2307/1884852

Skinner, B. F. (1938). *The behavior of organisms: An experimental analysis*. Appleton-Century.

Skinner, B. F. (1960). Pigeons in a pelican. *American Psychologist*, *15*(1), 28–37. https://doi.org/10.1037/h0045345

Stein, W. E., Seale, D. A., & Rapoport, A. (2003). Analysis of heuristic solutions to the best choice problem. *European Journal of Operational Research*, *151*(1), 140–152. https://doi.org/10.1016/S0377-2217(02)00601-X

Stephens, D. W., & Krebs, J. R. (1986). *Foraging theory*. Princeton University Press.

Todd, P. M. (1997). Searching for the next best mate. In R. Conte, R. Hegselmann, & P. Terna (Eds.), *Simulating social phenomena: Lecture notes in economics and mathematical systems* (Vol. 456, pp. 419–436). Springer. https://doi.org/10.1007/978-3-662-03366-1_34

Tversky, A., & Kahneman, D. (1974). Judgment under uncertainty: Heuristics and biases. *Science*, *185*(4157), 1124–1131. https://doi.org/10.1126/science.185.4157.1124

Valleur, M., Codina, I., Vénisse, J.-L., Romo, L., Magalon, D., Fatséas, M., Chéreau-Boudet, I., Gorsane, M.-A., Guilleux, A., Grall-Bronnec, M., & Challet-Bouju, G., & Groupe JEU. (2016). Towards a validation of the three pathways model of pathological gambling. *Journal of Gambling Studies*, *32*(2), 757–771. https://doi.org/10.1007/s10899-015-9545-y

Vandekerckhove, J., Rouder, J. N., & Kruschke, J. K. (2018). Editorial: Bayesian methods for advancing psychological science. *Psychonomic Bulletin and Review*, *25*(1), 1–4. https://doi.org/10.3758/s13423-018-1443-8

Vanderveldt, A., Oliveira, L., & Green, L. (2016). Delay discounting: Pigeon, rat, human—does it matter? *Journal of Experimental Psychology: Animal Learning and Cognition*, *42*(2), 141–162. https://doi.org/10.1037/xan0000097

Verhave, T. (1966). The pigeon as a quality-control inspector. *American Psychologist*, *21*(2), 109–115. https://doi.org/10.1037/h0023273

Wagenmakers, E. J., Morey, R. D., & Lee, M. D. (2016). Bayesian benefits for the pragmatic researcher. *Current Directions in Psychological Science*, *25*(3), 169–176. https://doi.org/10.1177/0963721416643289

Wallraff, H. G. (2001). Navigation by homing pigeons: Updated perspective. *Ethology Ecology and Evolution*, *13*(1), 1–48. https://doi.org/10.1080/08927014.2001.9522786

Watzek, J., & Brosnan, S. F. (2020). Capuchin and rhesus monkeys show sunk cost effects in a psychomotor task. *Scientific Reports*, *10*, Article 20396. https://doi.org/10.1038/s41598-020-77301-w

Willems, M. (2003). *Don't let the pigeon drive the bus*. Hyperion Books for Children.

Wright, A. A., & Cumming, W. W. (1971). Color-naming functions for the pigeon. *Journal of the Experimental Analysis of Behavior*, *15*(1), 7–17. https://doi.org/10.1901/jeab.1971.15-7

Zentall, T. R. (2011). Maladaptive "gambling" by pigeons. *Behavioural Processes*, *87*(1), 50–56. https://doi.org/10.1016/j.beproc.2010.12.017

Zentall, T. R., Laude, J. R., Case, J. P., & Daniels, C. W. (2014). Less means more for pigeons but not always. *Psychonomic Bulletin and Review*, *21*(6), 1623–1628. https://doi.org/10.3758/s13423-014-0626-1

Zentall, T. R., & Smith, A. P. (2016). Delayed matching-to-sample: A tool to assess memory and other cognitive processes in pigeons. *Behavioural Processes*, *123*, 26–42. https://doi.org/10.1016/j.beproc.2015.07.002

Zwick, R., Rapoport, A., Lo, A. K. C., & Muthukrishnan, A. V. (2003). Consumer sequential search: Not enough or too much? *Marketing Science*, *22*(4), 503–519. https://doi.org/10.1287/mksc.22.4.503.24909

Received February 25, 2021
Revision received September 16, 2021
Accepted September 19, 2021 ■