

Optimal Stopping Theory: Behavior of Pigeons (*Columba livia*) in the Secretary Problem

by

Hunter Pluckebaum and Jaidyanne Podsobinski

A thesis submitted in partial fulfillment of the requirements
for graduation with Honors in Psychology.

Author Note

Advisor to this senior thesis project was Walter T. Herbranson, Department of Psychology, Whitman College. Correspondence concerning this article should be addressed to Walter Herbranson, Department of Psychology, Whitman College, 345 Boyer Ave, Walla Walla, WA 99362.

E-mail: herbrawt@whitman.edu

Whitman College
2017

Certificate of Approval

This is to certify that the accompanying thesis by Hunter Pluckebaum and Jaidyanne Podsobinski has been accepted in partial fulfillment of the requirements for graduation with Honors in Psychology.

Walter Herbranson

Whitman College
May 10, 2017

Table of Contents

Acknowledgements.....	iv
Abstract.....	v
List of Figures and Tables.....	vi
Introduction.....	1
Methods.....	13
Results.....	16
Discussion.....	24
References.....	34

Acknowledgements

We thank our coauthor Zachary Hartzell who was an incredible thesis partner throughout the entire year. Additionally, we thank Walter Herbranson for his unwavering support and assistance throughout this experience and for giving us access to his Comparative Cognition Lab. We would also like to show our gratitude to the Psychology Department at Whitman College for all they have done for us over the past four years.

Abstract

Optimal stopping theory is a mathematical theorem concerned with selecting the optimal choice when presented with a series of options. Stemming from mathematical derivations, this theorem puts forth a set of guidelines intended to maximize rewards and mitigate loss. The Secretary Problem is one example to which optimal stopping theory can be applied. This basic procedure involves a decision-maker viewing a finite list of randomized candidates with the goal of selecting the best one. The decision-maker is permitted to view only one candidate at a time and can either select or reject a candidate on the spot; each decision is final. Many studies have shown that humans perform suboptimally at this task, usually by selecting a candidate too early. Previous research has shown that animals sometimes outperform humans on other probability-based tasks, and act in ways that maximize their potential gains. This exploratory study investigated pigeons' (*Columba livia*) performance on a multi-trial variation of the Secretary Problem task. Success within a trial was measured by whether or not the pigeon followed the solution laid out by optimal stopping theory. Results indicate that, similarly to humans, pigeons perform suboptimally by stopping too early. Through the analysis of the Secretary Problem, this study provided additional context for interpreting suboptimal performance in both humans and nonhuman animals.

Keywords: pigeon, the secretary problem, optimal stopping theory, optimal foraging theory

List of Figures and Tables

Table 1: Results of One Sample t-test Compared to the Test Value = 0.5	17
Table 2: Mean Acceptance Rates and Standard Error for Candidates.....	18
Table 3: Mean Performance and Standard Error Between Blocks	19
Figure 1: Interaction Between Blocks and Candidates for Acceptance Rate for the Five Quality of Candidates	20
Figure 2: Interaction Between Blocks and Days for the Average Number of Candidates Seen.....	21
Table 4: Mean Number of Candidates Seen and Standard Error between Each Block	22
Table 5: Mean and Standard Deviations for Application of the Optimal Solution	23
Figure 3: Average Percent of Wins During the Ten Days Constituting Blocks 1, 2, and 3.....	24

Optimal Stopping Theory: Behavior of Pigeons (*Columba livia*) in the Secretary Problem

Animals, both human and nonhuman, are constantly faced with decisions, most of which involve uncertain outcomes. The uncertainty in many of these real world decisions stems from their sequential nature. More specifically, a series of choices are made available over time, and it is often efficient and necessary to make a choice without waiting for all of the alternatives to be presented. When deciding on a restaurant in an unfamiliar town, for example, people often select one from a sequence of choices, without knowing the quality of subsequent restaurants. People aim to select the highest quality restaurant without having all the potentially relevant information. These real world decision-making scenarios have the same essential features as problems in the mathematical field of optimal stopping. Optimal stopping theory is a mathematical theorem concerned with determining the best way to select the ideal choice among a finite series of options. Stemming from mathematical derivations, this theorem puts forth a set of guidelines intended to maximize rewards and mitigate loss for individual decision-makers.

The Secretary Problem

One prominent application of optimal stopping theory is most commonly known as the Secretary Problem (SP). The SP is a recreational problem that arose in the late 1950's and early 1960's and was initially studied within the domain of mathematics (see Ferguson, 1989, for an overview). This hypothetical scenario symbolizes a variety of real life situations to which optimal stopping theory can be successfully applied. The basic procedure of the SP involves a decision-maker viewing

a finite list of randomized candidates with the goal of picking the best candidate. It is commonly articulated as a hiring scenario, as when interviewing job applicants for a secretarial position (hence the name), but the essence of the problem is its mathematical structure, and need not be tied to any specific scenario. While there are many variations to this problem, in its simplest form it has the following features (Seale & Rapoport, 2000):

1. Only a single candidate can be selected.
2. The number of candidates, n , is finite and is known before the search starts.
3. The candidates are seen sequentially in random order, with each possible order being equally likely. The decision-maker can only view one candidate at a time.
4. It is possible to rank all the candidates from best to worst. The decision to accept or reject a candidate is dependent on the relative rank of the candidate with respect to previously seen candidates.
5. Once rejected, a candidate cannot be recalled.
6. A decision-maker receives a payoff of 1, if he or she selects the best candidate, and 0, otherwise.

This basic problem as stated above has a remarkably simple solution: decision-makers should always reject the first n/e candidates (e being the base of the natural logarithm) and then stop at the first candidate that is better than all previously viewed candidates (meaning it has the relative rank of 1). If this rule is applied in the SP then the rate of selecting the best candidate is 37%, or $1/e$, regardless of the number of candidates.

Problems like the SP occupy a useful place in the study of problem solving and decision-making. For one, the optimal stopping problem is suited to controlled laboratory studies, making it of interest not just to mathematicians, but also to behavioral scientists (Seale & Rapoport, 1997). Secondly, the optimal stopping problem incorporates sufficient complexity and uncertainty to make it applicable to real world decision-making situations. The SP is interesting because it requires people to reason under uncertainty, a characteristic of most real-world decisions. It is rare to have a situation in which all relevant information is given prior to making a decision. Therefore, unlike other laboratory studies of decision-making, the SP reflects the uncertainty present in most real-world scenarios. Finally, because the SP has a known optimal solution, researchers can assess both the decision-makers' rate of achieving optimal outcomes as well as the strategies they utilize via specific quantitative analyses.

Human Performance on the Secretary Problem

While it may be surprising that humans deviate from optimal choice in systematic ways, behavioral economists have called into question the assumption that humans always make decisions in ways that are rational and that maximize utility (Tversky & Kahneman, 1973). In fact, humans fail to perform optimally in many decision-making scenarios and the SP is just one of them. While there has not yet been extensive research around the SP and optimal stopping theory in humans, preliminary findings indicate that humans generally do not adhere to the optimal stopping solution on the SP task (Lee, 2006; Seale & Rapoport, 1997). This suboptimal performance aligns with other research, indicating that despite the complexity of human thinking,

the heuristics and biases employed by humans lead to suboptimal performances on decision-making tasks (Granberg & Brown, 1995).

Seale and Rapoport (1997) conducted the initial empirical experimental investigation of the SP. In this study, 50 participants completed 100 trials of the SP, with every trial consisting of either 40 or 80 candidates. Participants viewed applicants sequentially on a computer screen and were shown the relative rank of each applicant. As the participants moved through each of the candidates, they could also see the relative (but not absolute) ranks of all candidates previously viewed, with 1 being the highest relative rank. For example, if three candidates were viewed, the relative ranks displayed were either (1, 2, 3), (1, 3, 2), (2, 1, 3), (2, 3, 1), (3, 1, 2), or (3, 2, 1). This process continued until a candidate was selected. The computer then revealed the absolute rank, informing the participant about whether or not the best possible candidate had been selected. After completing 100 independent trials, participants filled out a questionnaire about their decision strategies. Analyses of participants' strategies specifically looked at three particular decision rules that participants could have used. The first was a cutoff rule, whereby the decision-maker chooses to reject a certain number of applicants and then to select the next candidate that ranks higher than any of those previously seen. Note that the optimal strategy described earlier is a form of the cutoff rule, with the cutoff equal to n/e . The second was a candidate count rule, whereby the decision-maker chooses a candidate after encountering a certain number of competitive candidates (relative rank = 1). For example, a decision-maker would decide to select the 10th candidate with the relative rank of 1. The final was a successive non-candidate rule, whereby the decision-maker chooses the first

competitive candidate (relative rank = 1) that follows a certain number of non-competitive applicants (relative rank $\neq 1$). For example, a decision-maker could decide that after seeing 10 non-competitive applicants, the decision-maker would select the next candidate with the relative rank of 1. The results overwhelmingly found that participants tended to utilize the cutoff rule, but in doing so stopped earlier than the optimal policy would dictate.

In another study conducted by Seale and Rapoport (2000), participants performed a similar implementation of the SP, but with a slight variation. They changed the task so that participants were unaware of the population size. While this affected the probability of success, it did not change the optimal policy of the original SP. Once again, the results found that a majority of participants adhered to a cutoff rule. Furthermore, the participants again stopped too early, meaning that the cutoff was too low. These two studies are important because they show that without any guidance, people tend to operate under heuristics that share some characteristics with optimal stopping policies, but are not, in fact, optimal (Seale & Rapoport, 1997; Seale & Rapoport, 2000).

Bearden, Murphey, and Rapoport (2005) conducted a study with more realistic human decision-making scenarios by implementing additional dimensions from which each candidate in the SP could be assessed. In other words, rather than being told the relative rank of each candidate, participants in this multi-attribute SP had to determine the relative rank of each candidate by viewing multiple sub-rankings on different, uncorrelated attributes. By implementing this adaptation, they essentially created a scenario that more closely resembles the complex, multi-dimensional choices that

people face in everyday life. Consistent with previous literature in this area (Seale & Rapoport, 1997; Seale & Rapoport, 2000), the results suggest that participants made their selections prematurely, and as a result, the average overall ranking of the selected candidates was suboptimal and relatively low. The authors noted that this pattern of decision-making could be explained through the participants' application of a modified satisfying rule, by which a selection is made when a candidate has a minimum acceptable ranking on all attributes. However, participants failed at consistently setting a strict level of acceptability from which to judge candidates and, in doing so, had a tendency to select suboptimal candidates. A subsequent study conducted by Lee (2006) included an optimal stopping scenario that was very similar to the classical SP in order to gain insight into the possible underlying cognitive processes and strategies utilized by human participants. Participants were presented with a finite list of sequential numbers and their task was to choose the number with the highest value. Again, the results suggest that one possible strategy that decision-makers apply when completing a task that is solvable by the optimal stopping method is to use threshold-based models (i.e., a satisfying rule). Additionally, Lee found that there was no change in strategy over time, indicating that participants were not learning the optimal strategy through a trial and error process.

Despite the apparent simplicity of the optimal stopping theory solution, the current research shows that humans consistently perform suboptimally and do not improve over multiple trials in situations represented in the SP. Furthermore, this research suggests that suboptimal performance is caused by imperfect decision-making strategies. In the face of uncertainty, humans often rely on heuristic principles to

reduce a complex task into a simpler one (Tversky & Kahneman, 1973). In general, heuristics are quite useful in that they conserve time and effort, but they can also result in systematic errors. In the case of the SP, results show that humans tend to use a cutoff rule or some kind of threshold rule to select a candidate. Note that this fits the general pattern seen in most heuristic strategies: they simplify the problem and make the SP less taxing but can also lead to suboptimal performances. However, there are still gaps in the literature around the SP and why humans perform so poorly. Progress towards a clearer understanding may arise through a comparison between humans' decision-making strategies and decision-making strategies of other species that might or might not use similar kinds of heuristics.

Decision-Making in Nonhuman Animals

In the wild, animals are constantly required to make decisions under conditions of uncertainty; these situations vary immensely with respect to contextual factors and levels of sophistication. Thus, animals have had to adapt specific techniques for how to approach these circumstances; for in some cases these decisions can mean the difference between life and death. It follows that species frequently confronted with dilemmas that could jeopardize their survival would benefit from evolving cognitive capabilities that allow for better decisions. Comparative psychologists have been at the forefront of investigations into nonhuman decision-making, both in the laboratory and in the wild.

A fundamental problem faced by any human or nonhuman decision-maker is the compromise between swift decisions and good decisions. In general, the longer a decision-maker takes to make a decision, the better the result (Stephens & Krebs,

1986). Yet, there can also be costs associated with prolonged decision-making: time is a finite resource, and the longer it takes to make a decision, the less time a decision-maker has to allocate to other behaviors. This trade-off between speed and accuracy can be clearly illustrated in the decisions made during foraging. Understanding the rules that shape the foraging behaviors of animals has been a central focus of behavior analysis (Pyke, 1984). For instance, optimal foraging theory has provided a framework for predicting how predators may behave if the goal is to maximize the rate of energy gained while foraging (Stephens & Krebs, 1986). Optimal foraging theory posits that evolved cognitive processes push animals toward strategies that maximize their gains (i.e., food collected) while minimizing costs (i.e., time and energy expended).

For any animal, prey profitability (energy gained divided by prey handling time) is an essential variable for estimating the optimal diet, and by definition illustrates the trade-off between speed and accuracy. For example, the common shrew constantly faces foraging decisions that keep it only hours away from starvation. Because of its small size and fast metabolism, the shrew must satiate its voracious appetite with protein-rich insects, a high quality resource (Barnard & Brown, 1981). The shrew must decide whether or not to settle for readily available but suboptimal choices (i.e. smaller, less nutritious food items) in hopes of finding larger and more nutrient-dense prey that will yield a much higher payoff. Therefore, the shrew constantly has to make decisions on whether or not to pass up on the smaller and more accessible grubs in hopes that it will find termites, which better suit its high metabolism (Barnard & Brown, 1981).

Barnard and Brown (1985) found that shrews alter their foraging behaviors depending on states of hunger. They found that starving shrews are more likely to choose the higher risk but higher payoff sites, sites with less readily available food but potentially more nutrient-dense prey. While risky, the shrew opts for the location that could potentially give it the highest chance of success, while passing up on a more reliable but lower-payoff site. Given the conditions of starvation this risky strategy is optimal because opting for lower quality sites would not be enough to sustain the shrew in the state of survival. This strategy falls in line with an extension of optimal foraging known as risk sensitive foraging that suggests that under different circumstances predators respond differently to reward rates at various feeding sites (Barnard & Brown, 1985). In the case of shrews, the more critical the threat of starvation becomes, the more risk-prone they become in selecting a site to feed. In this instance, despite the potential for great loss (i.e. starvation), the shrew's preference for the riskier option is optimal given the necessity for a high payoff. Studies like this provide evidence that cognitive processes involved in decision-making can play a critical role in animals' foraging decisions and that these strategies are not fixed patterns of behavior but adaptive. As seen in the study of shrews above, optimal foraging theory demonstrates that like humans, animals are often faced with decisions that involve uncertainty and they use sophisticated strategies to maximize their chances of success.

However, animals' decision-making is not always optimal and can be subject to systematic biases just like human decision-making. For example, pigeons have been shown to choose an option that provides reinforcement at an average rate of 50% over a choice that provides them with a higher 75% rate of reinforcement (Laude, Pattison,

& Zentall, 2012). The researchers found that, when hungry, the pigeons opted for the choice that would provide a reliable signal for food half the time with a blue light and a reliable signal for the absence of food the other half with a yellow light (leading to food 50% of the time overall). This 50% reinforcement option was preferred over a different option that would lead to overall reinforcement 75% of the time regardless if the stimulus presented was the blue or yellow light. However, when the pigeons were not in a state of hunger they were more inclined to choose the 75% reinforcement choice. This research shows insight into the shortcomings of optimal foraging theory; while it is assumed animals should always behave optimally (e.g. demonstrate a preference for a higher rate of reinforcement), different internal and external factors can alter decision-making and may lead to suboptimal preferences. Further research with pigeons found that these effects are not limited to hunger. Zentall (2014), using a similar procedure, found that environmental factors like satiation or group housing could mitigate suboptimal performance on the choice task. These two studies again suggest that animals, like humans, can be susceptible to suboptimal decision-making.

Because of numerous differences in environment and physical abilities, cognitive processes vary between species. Therefore, some animals may perform optimally while others perform suboptimally, even on the same decision-making tasks. An example of this is the Monty Hall Dilemma (MHD), a probability puzzle that is notoriously difficult for humans, while pigeons generally perform optimally (Herbranson & Schroeder, 2010). The MHD is a recreational probability puzzle, anecdotally based off the game show *Let's Make a Deal*. Subjects are presented with three doors to choose from. Behind one door lies a valuable prize, while the other two

offer negligible gains for the participant. After an initial selection, one of the two unselected doors is opened to reveal an undesirable prize and then the contestant must decide whether to stay with the initial selection or to switch. The optimal strategy is to always switch, and doing so doubles the participant's likelihood of winning, yet people more often stay with their initial selection. Evidence shows that even when provided feedback and told the optimal strategy, humans still tend to perform poorly (Granberg & Brown, 1995).

Although humans demonstrate suboptimal choice behavior by not switching, Herbranson and Schroeder (2010) investigated whether pigeons would perform better on a variation of the MHD since relevant response biases and heuristics may differ. In this comparative experiment, pigeons performed a task having the same structure as the MHD, but adapted for an operant chamber. Pigeons performed close to optimally by switching more than 90% of the time. The MHD is one example of a task that requires estimation of relative probabilities, and at which pigeons outperform humans. Therefore, pigeons might make for an interesting test population for other decision-making tasks that produce suboptimal behavior in humans, such as the SP.

In addition to laboratory investigations, there are other reasons to suspect that pigeons may be well suited to the SP. Pigeons must frequently scout in a sequential manner for good areas to forage, usually with much uncertainty about the quality of different locations. Foraging locations vary in quality, so a pigeon must decide if the current site is likely to be the best available or if it would be better to reject the current location and invest additional effort in a continued search for a higher-quality location.

Since this form of decision-making roughly parallels many of the critical elements of the SP, it may be that pigeons possess the cognitive tools needed to perform optimally.

Current Study

This exploratory study focused on the decision-making of pigeons in a variation of the SP. This variation stays true to the mathematical structure of the SP, while modifying the task for an operant chamber. Over the course of this multi-trial experiment, pigeons viewed sequences of colored keylights in an operant chamber. Each of the different colors was associated with a different probability of reinforcement (20%, 40%, 60%, 80%, or 100%), and a pigeon could either peck a colored key for a chance to receive reinforcement (access to grain) with the assigned probability, or peck a different key to view the next candidate (a colored keylight, with its own associated probability of reinforcement) in the sequence. Success on a trial was measured by whether or not the pigeon followed the strategy laid out by optimal stopping theory (viewing n/e candidates and selecting the next option better than those already seen). The greatest deviation from the SP is that not all candidates will be ranked ordinally, meaning that while there will be only one ideal candidate, other, inferior candidates may be of the same quality. For example, if the ideal candidate on a trial corresponds to 40% reinforcement, then all nine of the other candidates will correspond to 20% reinforcement. However, this alteration does not change the procedure and, more importantly, does not alter the optimal stopping strategy.

For this study, there were four main hypotheses about pigeons' behavioral decision-making strategies on this variation of the SP. The first hypothesis was that, over the course of the data collection, the pigeons will reject lower quality candidates

at a higher frequency than higher quality candidates. This means that pigeons will proportionally choose candidates of higher relative rank at a greater frequency than candidates of lower relative rank (i.e. will select the 20% candidate the lowest frequency and the 100% candidate the highest frequency). Second, as the pigeons accumulate experience with the task, the daily average quality of candidate hired will increase. This means that, as the experiment continues, the pigeons will become more selective in the quality of candidate they choose to hire. In line with the second hypothesis, the third hypothesis states that by the end of the experiment the pigeons will consistently view a minimum of four candidates ($10/e$). In conjunction, the fourth hypothesis states that the average number of times the pigeons utilized the optimal solution would increase over time. These predicted results would indicate that pigeons had successfully adopted the solution laid out by optimal stopping theory.

Method

Animals

Seven white carneau pigeons (*Columba livia*) were used for this experiment. All were kept at approximately 80% free feeding weight in a temperature-controlled colony room with a 14:10 hour light:dark cycle. Each pigeon was housed in an individual cage with free access to water and grit. All experimental sessions took place five days a week at approximately the same time.

Apparatus

Seven BRS/LVE operant chambers were used during the experimental session. Each chamber contained three pecking keys and a feeder through which pigeons could

access grain. Each chamber was connected to a computer that controlled the daily trials and record data.

Stimuli

Stimuli presented during both the pretraining the experimental phase consisted of the same elements: color (red, green, blue, cyan, or magenta) and a geometric shape (a cross, consisting of two perpendicular white lines). Each trial included a simple two-element sequence in which a white-cross appeared on each of the three keys. Once a key was pecked, a solid colored light followed, replacing the cross. This solid colored light represented a candidate that the bird could either select (by pecking the colored keylight) or reject (by pecking a different key). Pecking the colored keylight cleared all three keys and resulted in reinforcement with a designated probability depending on the color (.2, .4, .6, .8, and 1.0; probabilities were assigned to different colors for each bird). Reinforcement consisted of approximately 3 seconds of access to grain. A peck to one of the remaining two keys displaying a cross led to the cross being replaced by a new colored keylight. The former colored keylight was replaced by a cross.

Procedure

Pretraining. During the pretraining phase, all sessions focused on habituating and autoshaping the pigeons for the experiment. At first, the pigeons were reinforced with $p = 1.0$ for following the sequence of pecking the white cross and then the following color. After 5-10 days, five different rates of reinforcement were implemented for each of the five colors (.2 - 1.0 in increments of .2). Each reinforcement probability was matched with a different color for each bird in order to

mitigate the effects of color preferences. Each of these colors represented the quality of a “candidate,” as specified in the SP.

Manipulation check. Before beginning the experimental phase, a five-day manipulation check was run. Pigeons were tested for successful association of color with the designated reinforcement probabilities. On each of 100 daily trials, two different colors were presented on adjacent keys. The first peck to either key ended a trial and resulted in reinforcement with the designated probability. It was expected that pigeons would peck the color associated with the higher rate of reinforcement. A successful check would indicate that birds acquired the relevant associations between colors and probabilities of reinforcement, and used them to display a preference for the “higher quality candidates.” A second two-day manipulation check was run in the experimental phase when switching from forced to unforced trial to ensure the learned association was still present.

Experiment. A daily experimental session consisted of a series of 80 individual trials. Each trial consisted of a sequence of up to ten candidates, each represented by one of the five colors used in pretraining. An individual trial began with a white-cross presented on each of the three response keys. When a white cross was pecked, it was replaced with a colored light (representing the first available candidate). The pigeon could then select either the current candidate present (by pecking the colored key) or opt to reveal an alternative candidate (by pecking the white cross displayed on either of the other two keys). If the pigeon opted to reveal the next candidate, the current candidate disappeared (the colored light was replaced by a cross) and the color representing the next available candidate was displayed on the pecked key. This

continued until a pigeon selected a candidate or all ten possibilities had been presented. In the latter case, pecks to crosses had no consequences and a bird could only choose the final candidate to terminate a trial. The likelihood of the pigeons being rewarded was dependent on the probability of reinforcement associated with the selected candidate. Reward consisted of approximately 3-s access to mixed grain (times varied from bird to bird to maintain 80% free feeding weight). If no reward was given, there was a 3-s delay before moving on to the next trial.

The ten candidates used in each trial were pseudo-randomly selected with the following constraints. On each trial, a “best candidate” was randomly chosen from 4 equally likely possibilities (.4, .6, .8, or 1.0). Then, 9 other candidates were randomly generated from 4 possibilities (.2, .4, .6, or .8) with the constraint that all must be less than the previously determined “best candidate” for that trial. The order of the ten candidates was then randomized for presentation. To offset possible impulsivity of the pigeons, they were initially forced to view at least three candidates on each trial. This means pigeons could not select any of the first three candidates within the ten-candidate sequence. This constraint was removed after 50 days and subsequently, pigeons were tested for an additional ten days. The data collected from each trial included the number of candidates viewed (1 - 10), the quality of candidate selected (.2 - 1.0), if the best candidate was selected (true or false), and if the optimal solution was used (true or false).

Results

Manipulation Check

Pigeons successfully learned associations between key color and quality of candidate during the pretraining stage of the experiment. Table 1 shows the average percentage of correct selections during the six day manipulation check for each bird. A series of one-sample t-tests against a null hypothesis mean of .50 indicates that each individual bird chose the better option on more than 50% of trials. If the pigeons were choosing between the two candidates randomly, the average percentage would verge on 50%, chance probability. The 95% confidence intervals for the mean difference, also included in Table 1 confirm this conclusion. Since the mean difference between the pigeons performance and the test value lies within the confidence interval, it can be concluded that the pigeons were preferentially selecting the higher quality candidates at a statistically significant rate.

Table 1. Results of One Sample t-test Compared to the Test Value = 0.5

Bird	<i>M</i>	<i>SD</i>	<i>n</i>	Mean Difference	95% <i>CI</i> of the Difference	<i>t</i>	<i>df</i>	<i>p</i>
1	.77	.11	6	.27	.15, .38	5.96	5	.002
2	.77	.07	6	.26	.19, .33	9.57	5	<.001
3	.80	.04	6	.30	.26, .35	17.59	5	<.001
4	.76	.05	6	.26	.21, .31	13.43	5	<.001
5	.82	.06	6	.32	.24, .40	10.20	5	<.001
6	.75	.03	6	.25	.22, .29	18.65	5	<.001
7	.72	.04	6	.22	.18, .25	14.68	5	<.001

Experimental Phase

To address the four main hypotheses of the study, three blocks, each containing ten days, were analyzed from the 60-day experiment. Block 1 was the first ten days that consisted of a full 80 forced trials (days 15 - 24). Block 2 was the last ten days of the 80 forced trials (days 41 - 50). Block 3 was the first ten days of the 80 unforced trials (days 51 - 60). Each analysis looked at differences between the blocks, the differences

among the ten days within each block, and the interaction between days (within each block) and blocks.

Acceptance rates of the five candidate ranks. A 5 (candidate quality: 20%, 40%, 60%, 80%, 100%) x 3 (block: 1, 2, 3) repeated measures ANOVA¹ was used to test the first hypothesis, that pigeons will select higher quality candidates a higher percentage of the time than lower quality candidates (i.e. pigeons will select 100% candidates more often than 80% candidates, 80% candidates more often than 60% candidates, and so on). Results from the analysis indicated that the hire rate differed statistically between the different qualities of candidates, $F(4, 36) = 246.12, p < .001$. Table 2 shows the results of post hoc comparisons using a Bonferroni correction. This post hoc test reveals significant differences between all qualities of candidates, except between the 60% candidate and the 80% candidate ($p = 1.0$).

Table 2. Mean Acceptance Rates and Standard Error for Candidates

Candidate	<i>M</i>	<i>SEM</i>	Mean Difference				
			20%	40%	60%	80%	100%
20%	.57	.01	-	-.24*	-.31*	-.31*	-.37*
40%	.81	.01	-	-	-.06*	-.07*	-.13*
60%	.87	.01	-	-	-	-.01	-.07*
80%	.88	.01	-	-	-	-	-.06*
100%	.94	.01	-	-	-	-	-

* $p < .05$

In addition, the ANOVA also revealed a significant main effect of blocks, $F(2, 18) = 36.27, p < .001$. Post hoc comparisons using a Bonferroni correction revealed that pigeons' performance on all combinations of the three blocks differed

¹ Note that this repeated measures ANOVA does not include days within block as mentioned in the previous paragraph. However, since the intention was to compare the quality of candidate selected over time, it was sufficient to collapse the days within each block for the analysis.

significantly, as seen in Table 3. Block 3 had significantly higher overall acceptance rates than block 1, $p < .001$, and block 2, $p < .001$. Block 2 was also significantly higher than block 1, $p = .04$.

Table 3. Mean Performance and Standard Error Between Blocks

Blocks	<i>M</i>	<i>SEM</i>	Mean Difference		
			1	2	3
1	.76	.01	-	-.04*	-.13*
2	.80	.01	-	-	-.09*
3	.89	.02	-	-	-

* $p < .05$

The two-way repeated measures ANOVA also revealed a significant interaction between candidate quality and block, $F(8, 72) = 7.16, p < .001$. Figure 1 shows the interaction between blocks and candidates, which shows different rates of increased hiring probability across blocks. This interaction shows that in all three block pigeons' display a preference for better candidates, though the increase is steepest during block 3 for lower-quality candidates.

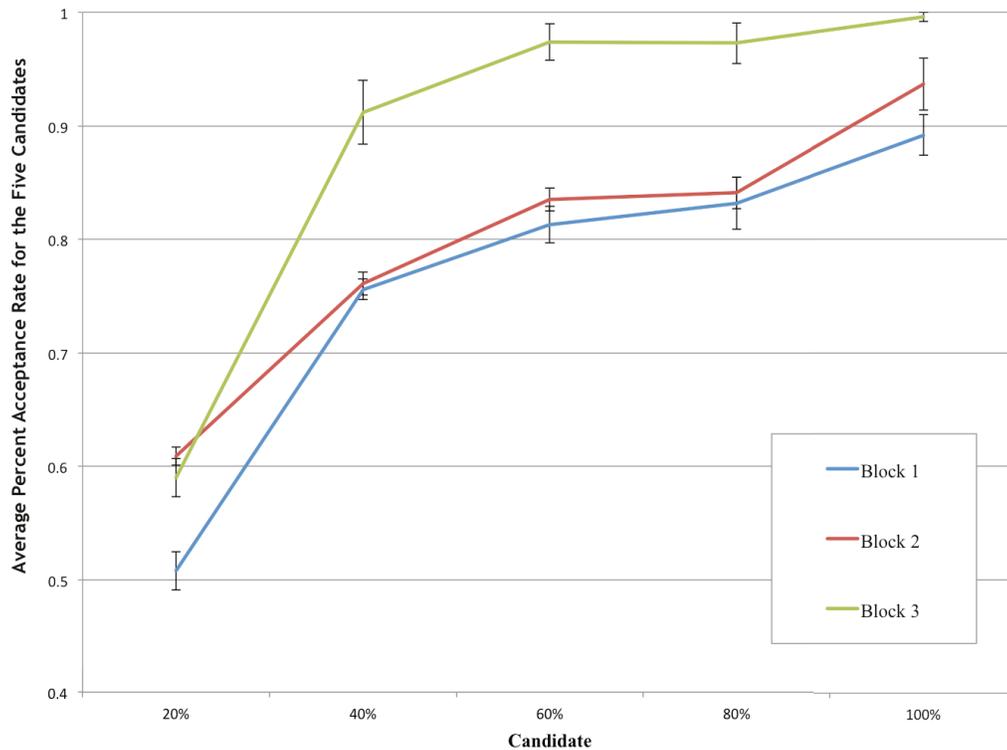


Figure 1. Interaction Between Blocks and Candidates for Acceptance Rate for the Five Quality of Candidates. Error bars represent standard error.

Average quality of candidate hired. A 3 (block: 1, 2, 3) x 10 (day within block: 1-10) repeated measures ANOVA was used to test the second hypothesis, that the average quality of the candidates that the pigeons select will increase over time (i.e. over the three blocks). Results from the analysis indicated that the mean candidate hired did not differ statistically between the three blocks, $F(2, 12) = 0.42, p = .67$, suggesting that the quality of the pigeons' candidate selection did not increase over the three blocks. Furthermore, the mean candidate hire did not differ statistically between days in each block, $F(9, 54) = 1.04, p = .42$, suggesting that the quality of the pigeons' candidate selection was stable within each block. Additionally, no interaction was observed between days and blocks, $F(18, 108) = 0.96, p = .52$. This indicates that pigeons performance was stable overtime.

Average number of candidates seen. A 3 (blocks: 1, 2, 3) x 10 (days within block: 1-10) repeated measures ANOVA was used to test the third hypothesis, that the average number of candidates the pigeons view will increase over the three blocks, eventually viewing a minimum of four candidates (which is inline with the optimal solution of the SP). The results from the analysis indicated that there was a significant difference between blocks, $F(2,10) = 2163.20, p < .001$, as well as a statistically significant interaction between day and block, $F(18,108) = 2.04, p = .013$. Figure 2 shows this interaction.

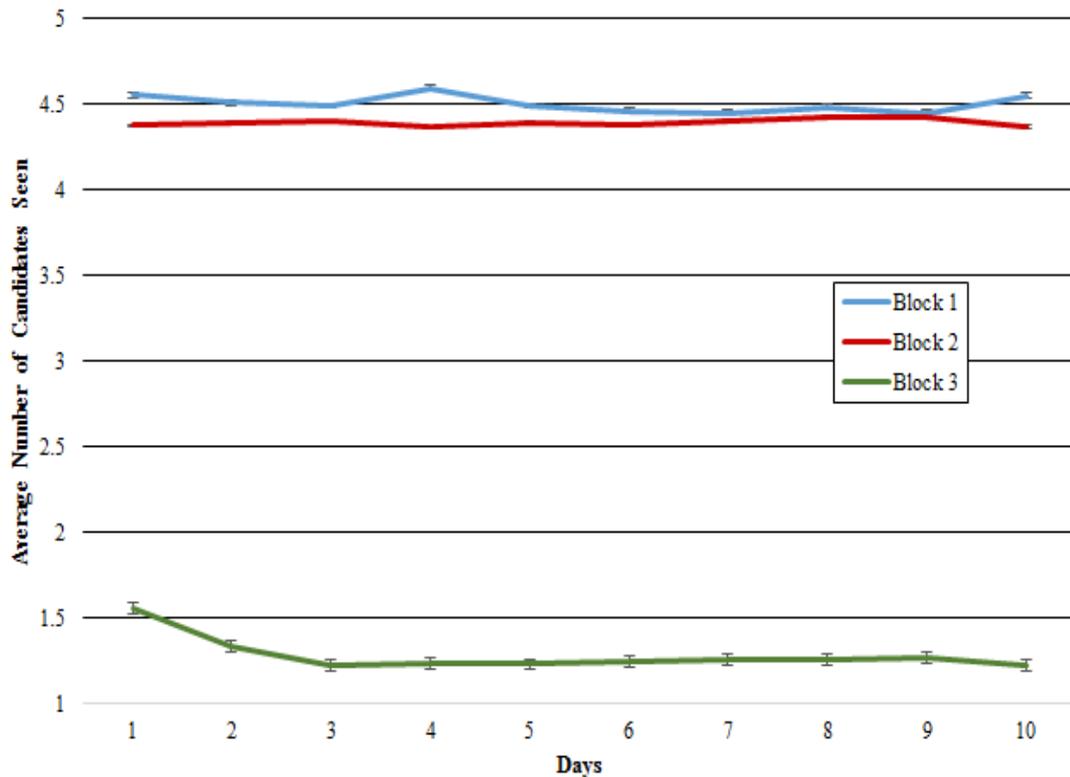


Figure 2. Interaction Between Blocks and Days for the Average Number of Candidates Seen. Error bars represent standard error.

Post hoc comparisons using Bonferroni correction, seen in Table 4, revealed that there was a significant difference between block three and blocks one and two. No

significant difference was seen between block one and block two. In addition, the results determined that the mean number of candidates seen differed significantly between the ten days within each block $F(9,54) = 2.74, p = .01$. However, post hoc comparisons revealed that there was no significant difference between any pair of two days within the blocks.

Table 4. Mean Number of Candidates Seen and Standard Error between Each Block

Blocks	<i>M</i>	<i>SEM</i>	Mean Difference		
			1	2	3
1	4.50	.07	-	.11	3.22*
2	4.40	.05	-	-	3.11*
3	1.28	.06	-	-	-

* $p < .05$

Average percentage of times the optimal solution was used. A 3 (blocks: 1, 2, 3) x 10 (days within block: 1-10) repeated measures ANOVA was used to test the fourth hypothesis, that the times the optimal solution was used would increase over time. There was a significant main effect of blocks, $F(2, 12) = 268.81, p < .001$. There was no significant main effect of days, suggesting that the pigeons' application of the optimal solution did not differ over the ten days within a block, $F(9, 54) = 1.15, p = .34$. Furthermore, there was no significant interaction between blocks and days, $F(18, 108) = 1.43, p = .13$. Post hoc tests using Bonferroni correction revealed that pigeons performed optimally at a significantly lower frequency in block 3 compared to block 1 (see Table 5), $p < .001$, and block 2, $p < .001$. However, no difference was found between block 1 and block 2, $p = .15$. This suggests that the removal of the forced trials in the third block reduced the number of times the optimal solution was used.

Table 5. Mean and Standard Deviations for Application of the Optimal Solution

Blocks	<i>M</i>	<i>SEM</i>	Mean Difference		
			1	2	3
1	4.50	.07	-	.11	3.22*
2	4.40	.05	-	-	3.11*
3	1.28	.06	-	-	-

* $p < .05$

Average percentage of wins. A 3 (blocks: 1, 2, 3) x 10 (days within block: 1-10) repeated measures ANOVA was used to look at the number of wins throughout the study. A win was defined as a pigeon selecting the best candidate out of the ten available. There was no significant main effect of blocks, suggesting that the removal of forced trials during block three did not alter the win rate, $F(2, 12) = 2.41, p = .13$. There was no significant main effect of days within each block, $F(9, 54) = 1.69, p = .12$. Furthermore, there was no significant interaction between blocks and days within block, $F(18, 108) = .97, p = .50$. Figure 3 shows the average daily percentage of wins during all three blocks in comparison to the percentage of average wins based on random chance (10%) and the expected average wins for the optimal solution (37%), which are represented by the black horizontal lines. Note that the average percentage of wins lies above 10% for all days excluding days six, seven, and nine in block 2 and day five in block 3.

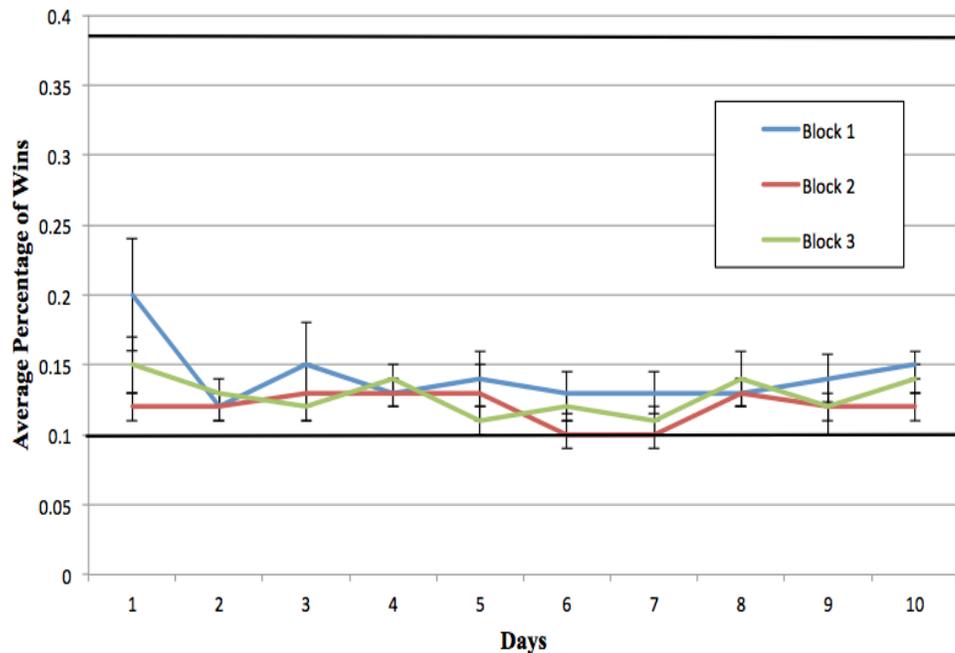


Figure 3. Average Percent of Wins During the Ten Days Constituting Blocks 1, 2, and 3. Error bars represent one standard error. Black horizontal bars represent average expected wins based on random chance (10%) and based on the optimal solution (37%).

Discussion

In this experimental version of the SP, the optimal strategy is to view at least four of the ten candidates and then to pick the next candidate having a higher relative rank than all previously viewed candidates. Doing so will result in a win (selection of the overall best candidate) approximately 37% of the time. While this optimal solution is seemingly simple and gives the highest probability of success when compared to any other strategy, humans do not seem to apply this strategy to the SP (Seale & Rapoport, 1997). Results from this experiment suggest that pigeons, similarly to humans, do not apply this strategy and therefore fail to perform optimally by making selections too early.

Since all of the pigeons performed comparably, with no significant deviance or outliers, their data was collapsed to a single group. None of the seven pigeons learned to consistently apply the optimal strategy over the course of the experiment. In fact, their performance stayed relatively stable when it came to applying the optimal strategy. This goes against the four hypotheses which overall predicted that the pigeons would improve with experience. To be more specific, the number of times that the optimal solution was used did not increase over time (and in fact decreased after the removal of forced trials). This was likely due to the fact that the average number of candidates seen dropped significantly after the removal of the three-candidate minimum within the forced trial phase. On both forced and unforced trials, pigeons often selected the first candidate they were permitted to choose. This means that during the forced trials the pigeons had a much higher likelihood of performing optimally since the constraint mandated a fundamental aspect of the optimal solution: viewing a minimum of four candidates. Additionally, unlike predicted, the average quality of candidate hired stayed stable over time. These findings suggest that pigeons were not altering their decision-making behaviors towards the optimal solution.

However, some of our results indicated that the pigeons were applying a consistent, meaningful strategy, even if it was non-optimal. Throughout the experiment, the pigeons selected higher quality candidates at a higher rate than lower quality candidates. This finding means that even though the pigeons were often impulsive by picking the first candidate they could (i.e. the fourth candidate during the forced trials phase and the first candidate during the unforced trials phase) they still implemented some selectivity in the hiring process by passing up lower quality

candidates more often. This is illustrated by the result that the 20% candidate had the lowest acceptance rate while the 100% candidate had the highest acceptance rate. Thus it could be said that, although the pigeons failed to adopt the optimal strategy, they still engaged in a kind of decision-making process that was more complex than just randomly choosing to hire or reject, or always choosing the first available option.

The fact that the pigeons often won at a rate that was better than chance also shows that the pigeons did something meaningful. While the average win percentage was still significantly below the success rate of the optimal solution, it was often significantly higher than 10%, thus indicating that the pigeons' behavior typically resulted in more wins than would be expected by chance. So despite the apparent impulsivity in the pigeons' decision-making behavior, there was also some form of a decision-making strategy which resulted in a higher than chance success rate on the SP.

Explanations

As mentioned previously, optimal foraging theory posits that animals will develop cognitive strategies that push them to maximize their gains while minimizing costs (Stephens & Krebs, 1986). While it was initially theorized, under optimal stopping theory, that pigeons' cognitive processes would lead to optimal performances on the SP, the foraging patterns that pigeons have evolved to use in the wild may not be compatible with this variation of the SP. Given the inevitable tradeoff between speed and accuracy, it is possible that the pigeons chose to emphasize speed over accuracy, by selecting a candidate and skipping immediately to the end of a trial. While this did not always result in reinforcement, it brought the onset of the next trial sooner, with its potential for reinforcement. Thus, the pigeons could have been applying a

strategy that, in real world foraging conditions, might have resulted in greater efficiency, particularly in a food-rich environment. Unlike in the lab, real-world foraging sessions are not defined by a set number of trials. Rather, they are defined by time intervals in which they can forage. These time intervals are influenced by both external factors (i.e, the end of daylight) and internal factors (i.e. satiation). Although nonhuman animals do not have fixed time constructs (i.e, hours, days), they do rely on temporal cues and adjust behaviors accordingly. Therefore, moving more quickly through the candidates rather than spending the extra time to select candidates of higher quality might reflect a more successful foraging strategy in real-world conditions. Given a limited time frame during which to forage, it might be more efficient for pigeons to opt for a higher number of low-quality feeding opportunities than to expend time looking for higher-quality feeding opportunities. This strategy also minimizes their cognitive expenditure and energy used while sorting through candidates.

In this way, pigeons' performance on the SP mirrors humans' performance in Lee's (2006) research into an optimal stopping problem that resembles the SP. Results from the study suggests that one common strategy decision-makers implemented was a simple threshold-based model, which has also been observed in other studies of the SP (Bearden et al., 2005). In this strategy, humans chose the first option that exceeded a fixed threshold for that list of numbers. This decision-making strategy might well be described as a "satisficing" model. In behavioral decision-making, those who implement a satisficing model seek options that are "good enough", even if they are not the best option (Simon, 1956). These satisficers stand contrary to maximizers who

spend extensive time seeking out all possible alternatives to make the absolute best choice (Cheek & Schwartz, 2016). Based off optimal foraging theory, it was hypothesized that pigeons would behave as maximizers and thus adopt the optimal solution to the SP (which also assumes a goal that corresponds to maximizing). However, the results of this study suggest that, in the case of the SP, pigeons behave more similarly to a satisficing model of behavior. By the end of the experiment, pigeons seemed to have been satisfied with selecting candidates of 40% and higher quality, accepting them at a 90% rate or higher, whereas 20% candidates were accepted at a much lower rate of around 60%. Furthermore, much like the decision-making patterns of the humans in Lee's (2006) study, the pigeons did not appear to learn or adapt to the SP with experience. This fixed pattern of behavior in the pigeons mirrors humans failure to learn the optimal stopping strategy. Therefore, many of the ways in which the pigeons display sub-optimal performance on this specific variation of the SP align with some of the common suboptimal behavioral choices seen in humans on optimal stopping problems.

Another behavioral aspect to consider is that pigeons often exhibit impulsive tendencies. Previous studies have shown that pigeons regularly chose a shorter period of reinforcement paired with a shorter delay over a longer period of reinforcement paired with a longer delay (Ainslie, 1974; Logue, Chavarro, Rachlin, & Reeder, 1988). In these studies, pigeons displayed impulsivity by consistently picking the key with a shorter delay even though it yields a smaller payoff. In line with these studies' findings, our results suggest that pigeons may be too impulsive to pass up the potential for immediate gratification. In the case of the current study, the tendency for

the pigeons to select the first option regardless of quality, rather than viewing more candidates, may be a sign of similar impulsivity. All of these studies suggest that time plays a large role in impulsivity, in that pigeons opted for immediate potential gratification even though it was a suboptimal choice. Once again, this ties into the speed-accuracy tradeoff in foraging, in which pigeons prioritize speed over accuracy. Overall, the current study indicates that pigeons' decision-making strategies may not lead to optimal performances, even though these strategies may be optimal in real-world foraging conditions.

Limitations & Future Directions

In the current study, several adaptations were made to the classical SP in order to investigate the decision-making processes of pigeons. The major deviation from the traditional SP was that the pigeons in this study always had a chance of being reinforced regardless of whether or not the selected candidate was the best in the set. The use of variable reinforcement allowed us to establish relative ranks of candidates but could also be a potential limitation. More specifically, the classic SP defines a "success" or a "win" as the selection of the single best candidate. Selection of any other candidate is considered a loss. In the current study, however, the pigeons always had at least a possibility of reinforcement regardless of which candidate they selected. It is possible that this deviation affected pigeons' behavior on the SP because selecting a suboptimal candidate was not costly enough, and could result in the same desirable outcome as a "win". Future studies with pigeons could work to make a different SP setup that is more reflective of the original SP to see if this variation would lead to a maximizing model of behavior.

Because this study deviated in some potentially important ways from the original SP, its results should be compared to other formulations with caution. Although adaptations needed for this experimental investigation limited the internal construct validity of the study by violating certain rules of the original SP, some of those changes made it potentially more reflective of real-world decision-making situations. When foraging in the wild, animals encounter multiple feeding opportunities that vary in quality. In these situations, they must decide whether to invest time and energy in foraging at a potentially fruitful site, or pass it up and invest additional search time in the hope of finding a higher quality site at which to forage. Similarly, hiring managers interview candidates of varying quality when seeking out a new hire. Neither of these scenarios are ‘all-or-none’ given that even a candidate that is third in overall relative rank can still be a good hire. Importantly, these scenarios are different from the classical formulation of the SP, which defines a win only as hiring the absolute best candidate from the available pool. An interesting future direction would be to incorporate this more realistic definition of success in human studies of the SP. Although there are already a number of studies that have examined human performance on the classic SP (Bearden et al., 2005; Seale & Rapoport, 2000), no one has investigated how human beings perform on versions of the SP that do not define success in this ‘all-or-none’ fashion. Future researchers should be mindful of these considerations when designing their studies. Not only could a comparative investigation allow for additional explanations as to why the pigeons in the current study behaved as they did, but it might also provide some additional insight into the

differences that exist between the decision-making processes of human beings and pigeons.

This study also illuminates the potential problem with studying decision-making comparatively in the unnatural environmental setting of the lab. Since it can be challenging to develop procedures that accurately reflect the conditions present in the real-world, it can be hard to know how applicable the results of controlled laboratory studies are to situations that animals might encounter in the wild. The external validity is also limited to animals raised within laboratory settings and may not extend to animals in the wild. Therefore, the lab may not be an ideal setting to study animals' behavioral decision-making within the framework of optimal foraging theory. Optimal foraging theory emerged from observations of animals in the wild, and it was meant to predict the behavior of animals who adapt to their natural environments (Pyke, 1984). Therefore, it may be an overextension of the theory to study it in laboratory settings.

For example, ethical considerations limit researchers' ability to simulate certain characteristics of the wild such as the constant threat of starvation or predation. Optimal foraging theory incorporates these real-world factors by assuming that animals' cognitive processes evolved in contexts where these factors are present. This theory assumes that these foraging behaviors are optimal and deeply ingrained. However, just because they are optimal in real-world foraging situations, one should not assume that these behaviors translate to a laboratory decision-making task. Therefore, optimal foraging theory may not be relevant to researchers who are attempting to predict the foraging behaviors of laboratory animals. Future studies

should consider this when studying the behavior of animals removed from their natural environments.

Note that these limitations do not imply that meaningful comparative studies cannot be done in the laboratory. Comparative studies deepen our understanding of cognitive processes in both human and nonhuman animals and are valuable to the field. Examining the behavior of animals in a controlled setting has the potential to provide insight into the cognitive processes used by wild animals, even if these behaviors are not optimal in a laboratory task. For instance, while the pigeons' impulsiveness was detrimental to their application of the optimal solution in this study, this behavioral tendency may benefit them in natural environments. In this task, the cost of passing up a candidate in order to view another one did not include a significant energy expenditure (i.e. flying until they reach another foraging site) nor an appreciable loss of time; however, in the wild these are both considerable factors as pigeons may benefit from choosing the first foraging opportunity rather than risking the uncertainty of distance and time to the next one. Future studies on pigeons' performance on the SP could work to minimize the behavioral tendency of impulsivity in pigeons. A study conducted by Zentall and Laude (2013), comparing pigeons' behavior to human gambling, found that housing pigeons in enriched environments reduced their impulsiveness. Enriched environments include group housing, and providing opportunities for animals to interact with their environment (i.e, puzzles, objects). Therefore, a future direction on the SP could look at how enriched housing environments may influence performance. Additionally, future studies may benefit from incorporating delays between the stimulus and reinforcer, which may subvert the

pigeon's impulsive tendencies. These variations on housing environments and methodology might improve performance on the SP by reducing impulsive behaviors.

Conclusion

Broadly speaking, cross-species comparisons have the potential to enrich our understanding of human cognition as well as how our cognitive processes differ from other species. This study suggests that pigeons' performance as well as humans' performance on the SP may stem from a satisficing model of decision-making by consistently selecting a candidate too soon. However, under some circumstances, humans act as maximizers (Parker, De Bruin, & Fischhoff, 2007), as do pigeons (Herbranson & Schroeder, 2010). A complete understanding of decision-making should include an understanding of the situations that would lead to each pattern, as well as account for the cross-species differences. In particular, researchers should be more mindful of applying optimal foraging theory to laboratory studies since that may be an overextension of the theory. Ultimately, this study sheds light on the complexity of behavioral decision-making strategies, in both human and non-human animals, in relation to optimal stopping theory and the SP.

References

- Ainslie, G. W. (1974). Impulse control in pigeons. *Journal of the Experimental Analysis of Behavior*, *21*(3), 485–489.
- Barnard, C. J., & Brown, C. A. J. (1981). Prey size selection and competition in the common shrew (*Sorex araneus* L.). *Behavioral Ecology and Sociobiology*, *8*(3), 239-243.
- Barnard, C. J., & Brown, C. A. J. (1985). Risk-sensitive foraging in common shrews (*Sorex araneus* L.). *Behavioral Ecology and Sociobiology*, *16*(2), 161-164.
- Bearden, J. N., Murphy, R. O., & Rapoport, A. (2005). A multi-attribute extension of the secretary problem: Theory and experiments. *Journal of Mathematical Psychology*, *49*(5), 410-422.
- Cheek, N. N., & Schwartz, B. (2016). On the meaning and measurement of maximization. *Judgment and Decision Making*, *11*(2), 126-146.
- Ferguson, T. S. (1989). Who solved the secretary problem? *Statistical Science*, *4*(3) 282-289.
- Granberg, D. & Brown, T. A. (1995). The monty hall dilemma. *Personality and Social Psychology Bulletin*, *21*, 711–723.
- 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.
- Laude, J. R., Pattison, K. F., & Zentall, T. R. (2012). Hungry pigeons make suboptimal choices, less hungry pigeons do not. *Psychonomic Bulletin & Review*, *19*(5), 884-891.

- Lee, M. D. (2006). A hierarchical bayesian model of human decision-making on an optimal stopping problem. *Cognitive Science*, *30*(3), 1-26.
- Logue, A. W., Chavarro, A., Rachlin, H., & Reeder, R. W. (1988). Impulsiveness in pigeons living in the experimental chamber. *Animal Learning & Behavior*, *16*(1), 31-39.
- Parker, A. M., De Bruin, W. B., & Fischhoff, B. (2007). Maximizers versus satisficers: Decision-making styles, competence, and outcomes. *Judgment and Decision Making*, *2*(6), 342.
- Pattison, K. F., Laude, J. R., & Zentall, T. R. (2013). Environmental enrichment affects suboptimal, risky, gambling-like choice by pigeons. *Animal Cognition*, *16*(3), 429-434.
- Pyke, G. H. (1984). Optimal foraging theory: A critical review. *Annual Review of Ecology and Systematics*, *15*, 523-575.
- 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.
- Seale, D. A., & Rapoport, A. (2000). Optimal stopping behavior with relative ranks: The secretary problem with unknown population size. *Journal of Behavioral Decision Making*, *13*(4), 391-411.
- Simon, H. A. (1956). Rational choice and the structure of the environment. *Psychological Review*, *63*(2), 129-138.
- Stephens, D. W., & Krebs, J. R. (1986). *Foraging theory*. Princeton University Press.

Tversky, A., & Kahneman, D. (1973). Availability: A heuristic for judging frequency and probability. *Cognitive Psychology*, 5(2), 207-232.

Zentall, T. R. (2014). Suboptimal choice by pigeons: An analog of human gambling behavior. *Behavioural Processes*, 103, 156-164.

Zentall, T. R., & Laude, J. R. (2013). Do pigeons gamble? I wouldn't bet against it. *Current Directions in Psychological Science*, 22(4), 271-277.