



Published in final edited form as:

*J Comp Psychol.* 2010 February ; 124(1): 1–13. doi:10.1037/a0017703.

## Are Birds Smarter Than Mathematicians? Pigeons (*Columba livia*) Perform Optimally on a Version of the Monty Hall Dilemma

Walter T. Herbranson and Julia Schroeder  
Department of Psychology, Whitman College

### Abstract

The “Monty Hall Dilemma” (MHD) is a well known probability puzzle in which a player tries to guess which of three doors conceals a desirable prize. After an initial choice is made, one of the remaining doors is opened, revealing no prize. The player is then given the option of staying with their initial guess or switching to the other unopened door. Most people opt to stay with their initial guess, despite the fact that switching doubles the probability of winning. A series of experiments investigated whether pigeons (*Columba livia*), like most humans, would fail to maximize their expected winnings in a version of the MHD. Birds completed multiple trials of a standard MHD, with the three response keys in an operant chamber serving as the three doors and access to mixed grain as the prize. Across experiments, the probability of gaining reinforcement for switching and staying was manipulated, and birds adjusted their probability of switching and staying to approximate the optimal strategy. Replication of the procedure with human participants showed that humans failed to adopt optimal strategies, even with extensive training.

### Keywords

pigeon; probability; matching; Monty Hall Dilemma

---

Animals are frequently presented with choices where the outcomes associated with each option are ambiguous and probabilistically determined. A classic example is the decision of where to forage for food: a location that has the signs of being food-rich might recently have been picked clean, and a location that appears to be food-poor might contain an unseen cache. Thus, an animal cannot be absolutely certain of the outcome of its choice. A low-probability or “bad” choice might yield a favorable outcome, and a high-probability or “good” choice might produce disastrous results. However, even if the chooser cannot be *absolutely* certain, many of these choices are associated with a particular *probability* of success.

Optimal foraging theory (Stephens & Krebs, 1986) provides specific, quantitative models of how decisions about foraging could be made. It assumes that animals behave in such a way as to maximize potential gains while minimizing expenditures. As such, optimal foraging theory often provides highly accurate descriptions of how a variety of animals forage in natural environments. However, the models are not infallible: there are bound to be cases where animals will deviate from the predictions made by optimal foraging theory. Furthermore, any foraging data gathered in a natural setting will include the effects of numerous uncontrollable variables that are not accounted for by a specific model.

Laboratory investigations on the other hand, contain inherent limitations but can allow for a higher degree of control and precision.

In the laboratory, probability learning experiments have attempted to illuminate how such probabilistic choices are made. In a simple probability learning experiment, reinforcement is available by responding on one of two alternatives, each of which is associated with a specific likelihood of reinforcement. In some cases, responding approximates optimality, in that an animal responds in such a way that no other way of responding would yield a higher payoff (Shimp, 1966, 1973). However, in some situations, the distribution of responses falls short of maximizing the average expected payoff (Bullock & Bitterman, 1962; Mazur, 1981). That is, animals may respond in such a way that adopting a different strategy would yield more desirable outcomes over the long run. The tendency to fall short of optimal choice is not exclusive to pigeons or other animals and frequently includes humans (Fantino & Asfandiari, 2002). Such laboratory investigations of choice are not only methodologically precise but have led to the development of quantitative theories that are useful both in and out of the laboratory. For example, the matching law has long been used to describe choice behavior in controlled laboratory settings (Herrnstein, 1961, 1997) and has also shown promising fit for data collected in the wild (e.g., Houston, 1986).

Several other kinds of behaviors involve the same kinds of uncertainty. Consider the well-researched problem of categorization, where potentially infinite numbers of exemplars must be assigned to a small number of groups or categories. In complex natural situations, category membership may not be at all certain, and similar stimuli can become confused. For example, a camouflaged insect may be seen as a leaf; eyespots may cause an animal to be seen as a different, more threatening species; or a brood-parasitic cowbird's egg may be confused for a nesting host-mother's own. In recent years, mathematical models of categorization have been used to quantify the uncertainty inherent in categorization (Maddox & Bohil, 2004). As with probability learning, responses may be optimal, in that they approximate the highest possible level of accuracy (even if the highest possible level of accuracy is not 100%), or they may be suboptimal and fall short of the highest possible level of accuracy. Furthermore, sub-optimal performances, when found, often follow a consistent pattern, arising in similar kinds of situations. Also, as in probability learning, humans and pigeons do not show exactly the same patterns of optimal and suboptimal responses (Herbranson, Fremouw, & Shimp, 1999, 2004).

It may be initially surprising that humans deviate from optimal choice in systematic ways. However, a long-standing tradition in behavioral economics has called into question the assumption that humans make decisions in ways that are rational and that maximize utility (Tversky & Kahneman, 1981). Thus, both human and nonhuman behaviors show suboptimality, and it is yet uncertain whether the situations in which humans and nonhuman animals fail to maximize will always be the same. Progress toward an answer might come from systematic investigations of circumstances that produce a failure to maximize in one or more species.

One particularly compelling example of suboptimal choice behavior in humans has come to be known by mathematicians and statisticians as the "Monty Hall Dilemma" (MHD). The MHD is anecdotally based on the TV game show program *Let's Make a Deal* (hosted by Monty Hall) and involves a hypothetical contestant who is given the opportunity to select from a set of three closed doors. One of the doors conceals a valuable prize. The two remaining doors have behind them undesirable "zonk" or "goat" prizes. Once the contestant selects a door, Monty Hall opens one of the remaining doors, always revealing one of the two undesirable prizes (and being careful to never uncover the desirable one). With the remaining two doors still closed, he then offers the contestant the choice to either stay with

the initial selection or to switch to the other remaining door. Regardless of which door was initially selected, the great majority opt to stay with their initial choice (Posner, 1991; Tierney, 1991).

The tendency to stay is a powerful example of suboptimal behavior, because a contestant who switches is twice as likely to win as a contestant who stays. One explanation for why switching is better is as follows: Given that contestants do not know the location of the prize, they will choose its location with their first pick approximately one third of the time. On this one third of all trials, switching always loses and staying always wins. However, on the remaining two thirds of all trials, contestants choose doors that contain “goat” prizes. On these two thirds of all trials, switching always wins and staying always loses (recall that the opened door cannot be the one that contains the desirable prize).

Anecdotal examples of people consistently failing to switch and refusing to accept that switching offers an advantage have been empirically supported by several laboratory investigations of the MHD (Granberg, 1999a, 1999b; Granberg & Brown, 1995), all of which show a stubborn tendency to stay or, at best, no preference either way. The reasons for such consistent failures of humans to maximize the probability of winning remain unsettled, though several possibilities have been suggested (see Gilovich, Medvec, & Chen, 1995; Tubau & Alonso, 2003). Given that different species do not always show the same kinds of response biases and optimality, an investigation of the MHD in a nonhuman species might provide some illuminating context and fuel a possible explanation.

## Experiment 1

Pigeons often perform quite impressively on tasks that require estimation of relative probabilities. In some cases, pigeons may even eclipse the performances of human participants (Arkes & Ayton, 1999). For example, Hartl and Fantino’s (1996) pigeons chose more optimally than Goodie and Fantino’s (1996) college students in a parallel situation. Given that the MHD is essentially one that requires a participant to assess probabilities, pigeons might make for a particularly interesting test population. If pigeons, like humans, fail to maximize, then the MHD might be an informative addition to the already rich probability learning literature. Alternatively, if pigeons perform optimally and maximize their expected winnings, one must consider why they do so when the problem is so notoriously difficult for humans. A standard three-key operant chamber easily lends itself as an analog to the three doors in the MHD.

## Method

**Animals**—Six Silver King pigeons (*Columba livia*) were obtained from a local breeder. Each was maintained at approximately 80% of its free feeding weight in a temperature-controlled colony room with a 14:10 light:dark cycle. Each pigeon was housed in an individual home cage with free access to water and grit. All experimental sessions took place at approximately the same time five days per week.

**Apparatus**—Three BRS/LVE operant chambers were used. The front wall of each chamber contained three pecking keys and a feeder through which birds could gain access to mixed grain. Each chamber was interfaced to a personal computer that controlled all experimental events, recorded data, and computed daily statistics.

**Procedure**—All birds were pretrained in sessions consisting successively of habituation, magazine training, and autoshaping (Brown & Jenkins, 1968) until consistent responding was achieved on each of the three response keys. Stimuli during autoshaping were white, red, and green key lights.

A daily experimental session consisted of a series of individual trials (10 trials on the initial day, gradually increased to 100 over the first five days of the experiment). Each trial consisted of an initial choice response from among three options, a second choice response from among two options, food delivery if the second response was correct, and an intertrial interval (ITI).

At the beginning of a trial, the computer randomly selected one of the three keys (each with a probability of one third) as the prize location. Then, all three keys were illuminated with white light. After the bird pecked any one of the lit keys, all three keys were darkened for 1 second. During the 1-second delay, the computer pseudorandomly selected one of the three keys to deactivate for the remainder of the trial, with the constraint that the deactivated key could not be the one selected as the prize location for that trial, nor could it be the one that the bird had pecked on that trial. For the remainder of a trial, the deactivated key remained darkened and any responses on it were ignored. The two remaining keys (the key that the bird had just pecked, and one other key) were then illuminated with green light. If a bird then pecked the key that corresponded to the prize location, it was given approximately 3 seconds access to mixed grain (times varied from bird to bird to maintain individual running weights) and, following a 5-second ITI, moved along to the next trial. If a bird's second response did not correspond to the prize location, no reinforcement was provided, and the procedure simply moved directly to the ITI and the next trial.

For each trial, the result of interest was whether the bird pecked the same key on its second choice (from among two green keys) as it did on its first choice (from among three white keys). If it did, the bird was said to have stayed. If the bird pecked a different key with its second choice than it did with its first choice, it was said to have switched. On average, staying would lead to food delivery one third of the time, and switching would lead to reinforcement two thirds of the time.

## Results

**Probability of switching**—Figure 1 shows the average probabilities of switching and staying on the first day of the experiment (top panel) and on the thirtieth and final day of the experiment (bottom panel). On Day 1, pigeons switched on 36.33% of the trials ( $CI = 25.12$  to  $47.55^1$ ). Note that if a pigeon simply guesses randomly, it would switch on 50% of all trials. This chance performance value lies outside of the confidence interval ( $d = 1.04$ ). By Day 30, pigeons switched on 96.33% of the trials ( $CI = 93.97$  to  $98.70$ ). Again, chance performance lies outside of the confidence interval, ( $d = 16.81$ ). A comparison of Days 1 and 30 shows that pigeons switched on 60.00% more trials on Day 30 ( $CI = 49.44$  to  $70.56$ ,  $d = 5.55$ ).

**Reinforcers earned**—Figure 2 shows the proportion of reinforced trials on the initial and final days of the experiment. On Day 1, pigeons' choices were reinforced on 47.00% of the trials ( $CI = 43.30$  to  $50.69$ ). On Day 30, pigeons were reinforced on 63.83% of the trials, ( $CI = 59.88$  to  $67.78$ ). Thus, pigeons earned reinforcement on 16.83% more trials on Day 30 than on Day 1, an improvement that cannot be attributed to random chance ( $CI = 9.460$  to  $24.207$ ,  $d = 3.82$ ).

Three reference points are relevant to a quantitative analysis of the number of reinforcers earned and are shown on Figure 2 as horizontal lines. These reference points correspond to the expected number of reinforcers earned by a bird that always stays (33%), always

<sup>1</sup>All confidence intervals reported are 95% confidence intervals and refer to the probability of switching. Because all trials required either a stay or a switch response, the probability of staying can always be calculated as 100% minus the probability of switching.

switches (67%), and guesses randomly (50%). The confidence interval for Day 1 indicates that birds were reliably different from the theoretical expectation for staying (33%,  $d = 3.25$ ) and switching (67%,  $d = 4.64$ ) but not guessing (50%,  $d = 0.70$ ). The confidence interval for Day 30 indicates that birds were reliably different from the theoretical expectation for staying (33%,  $d = 6.70$ ) and guessing (50%,  $d = 3.00$ ) but not switching (67%,  $d = 0.70$ ).

**Individual birds' choices**—Table 1 shows data for the first and last session for each of the six birds in Experiment 1. Each column shows one of the nine possible combinations of responses on a trial. Responses on the same key (stay responses) are printed in italics. Notice that on the first day of training there is considerable between-bird variability and that each bird shows a variety of both staying and switching responses. The table also shows that, on the final day of training, all birds had a consistent tendency to switch and that the effect in Figure 1 cannot be attributed to exceptional performances by a subset of the pigeons.

## Discussion

In an iterated MHD, such as the one in Experiment 1, the optimal strategy is to switch on every trial. Doing so will result in wins on two thirds of all trials. While this necessarily means that the remaining one third of trials will be losses, no other strategy can match or exceed it. Most strategies result in reinforcement on 50% of the trials (the same as guessing randomly), and some can yield considerably worse outcomes.

All six pigeons in Experiment 1 learned to switch consistently, and by doing so they earned close to the maximum possible payoff. Considering that switching produces a higher rate of reinforcement than staying, the procedure appears to be one that naturally shapes pigeons to switch. While choices to stay were initially more abundant, choices to switch were twice as likely to be reinforced. The higher probability of reinforcement for switching gradually increased the proportion of switch responses until the suboptimal stay responses were eliminated almost entirely. Once a pigeon was consistently switching, there was no way to further increase the likelihood of reinforcement.

Thus, it might be said that pigeons adopted a strategy<sup>2</sup> of consistently switching. Aside from the optimal strategy of consistently switching, two other possible strategies deserve consideration. The first is one that involves randomly choosing whether to stay or switch. In the absence of any prior history of reinforcement for staying or switching, this is apparently what pigeons did on the initial days of the experiment: not only was there a mixture of stay and switch responses, but the proportion of reinforced trials was consistent with the expected outcome of a random strategy: 50%. By the end of the experiment, however, birds earned many more reinforcers than one could expect if they were responding randomly. The final strategy that should be considered is one that involves consistently staying. While staying results in reinforcement on only one third of trials (and in fact, it is impossible for any strategy to consistently do worse), recall that this is exactly what most humans do when presented with the MHD. Pigeons did show an initial preference for staying, which gradually became a strong tendency to switch. It is unknown whether this initial tendency to stay reflects something specific to the MHD: while it is consistent with what humans do, it was not particularly persistent and could simply be reflective of pigeons' more general behavioral tendency to develop position biases, wherein they preferentially respond on a particular key. Despite the initial preference for staying, it is important to note that the number of reinforcers earned even on the first day exceeded what one could expect via a

<sup>2</sup>The use of the term strategy here is only meant to indicate a consistent pattern in birds' responses. Just as biologists' use of the term "Evolutionary Stable Strategy" (ESS) does not imply planning or forethought on the part of a population, we make no such assumptions about birds here.

consistent staying strategy. Thus, it is critically important that the initial preference to stay was not exclusive, and that even on the first day of training birds frequently switched.

## Experiment 2

A simple and parsimonious interpretation of Experiment 1 might be that pigeons learned to switch rather than stay by trying each possibility and then consistently choosing the one that yielded the most frequent payoffs. If this interpretation is true, then it should be possible to change the likelihood of switching by manipulating the relative likelihoods of reinforcement associated with switching and staying.

By the above interpretation, birds in Experiment 1 chose to switch simply because switching had maximized the probability of gaining reinforcement on previous trials. However, consider a reversed version of the MHD, in which the likelihood of reinforcement is greater if a player stays with their initial choice. This version is interesting because it is conceivable but highly unnatural and could only occur if the location of the prize is determined after the player's initial guess. If pigeons learned whether to stay or switch based on experience with the problem, they ought to be capable of learning to stay just as easily as they learned to switch in Experiment 1. The current interpretation, then, would predict that pigeons will learn to stay, just as they learned to switch in the previous experiment.

## Method

**Animals**—Six Silver King pigeons (*Columba livia*) were obtained from a local breeder. They were maintained under the same conditions as those in Experiment 1.

**Apparatus**—The operant chambers were the same as those used for Experiment 1.

**Procedure**—Birds were pretrained as in Experiment 1 until consistent responding was achieved on each of the three keys.

The structure of a trial was the same as in Experiment 1 with the exception of the method of determining the prize location. In Experiment 2, the prize location could only be determined after a pigeon had made its initial choice (from among the three white keys). Specifically, the prize location was determined to be the key that a pigeon had just pecked with a probability of two thirds. The prize location was linked to the two keys that had not been pecked with a probability of one sixth each (a combined one third for the two together). One of the keys was then deactivated (again, with the constraint that it could not be the key that a pigeon had just pecked nor the prize location), and the remainder of the trial followed the same procedure as Experiment 1.

Again, the critical result for each trial was whether a pigeon stayed (pecked the same key twice) or switched (pecked different keys) on each trial. Note that in Experiment 2 the probabilities of reinforcement associated with staying and switching are the converse of Experiment 1: staying would lead to reinforcement two thirds of the time (rather than one third), and switching would lead to reinforcement one third of the time (rather than two thirds).

## Results

Figure 3 shows the average probability of staying and switching on the first day of the experiment (top panel) and on the fifteenth and final day of the experiment (bottom panel). On Day 1, pigeons switched on 30.17% of the trials ( $CI = 17.60$  to  $42.73$ ). The confidence interval does not encompass the chance performance value of 50% ( $d = 1.35$ ). By Day 15, pigeons switched on 4.33% of the trials ( $CI = -0.20$  to  $8.87$ ). Again the confidence interval



does not contain the chance performance value of 50% ( $d = 8.63$ ). A comparison of Days 1 and 15 shows that pigeons switched on 25.83% fewer trials on Day 15, an improvement that cannot be attributed to chance ( $CI = 11.17$  to  $40.50$ ,  $d = 2.09$ ).

Birds settled into a consistent response pattern more quickly than they did in Experiment 1. In Experiment 1, birds reached a stable baseline, on average, by Day 25 (defined as a period of five consecutive days over which the proportion of switch responses varied less than 10%). In Experiment 2, they reached this point by Day 12, roughly twice as fast.

## Discussion

In Experiment 2, birds learned to stay with their initial choices. Note that they learned to do the opposite (stay) of what pigeons did in Experiment 1 (switch), but they presumably did so for the same reason. In both cases, the adopted response strategy was the one that maximized the potential to earn reinforcement. In Experiment 1 switching provided a higher likelihood of reinforcement than staying, and birds learned to switch. In Experiment 2 staying provided a higher likelihood of reinforcement than switching, and birds learned to stay.

It is interesting that birds apparently learned to stay with their first choice in Experiment 2 more quickly than birds learned to switch in Experiment 1. On the surface, this finding has the potential to bridge the gap between pigeons' optimal performance and humans' suboptimal performance on the MHD. That is, humans have shown a tendency to stay that is extremely difficult to overcome. Pigeons apparently are not quite as rigid, but they did learn to stay more quickly than they learned to switch. However, one should be hesitant to make a claim that staying is easier for pigeons to learn for two reasons. First, pigeons began both experiments with a slight tendency to stay and, thus, had further to go in Experiment 1 to reach a stable optimal strategy of always switching. In Experiment 2, the same initial response distribution was quite close to the optimal strategy of always staying—which would naturally produce a shorter acquisition time given the same learning rate. Second, pigeons likely had the slight tendency to stay in the first place because of the relative distances to the keys required for stay and switch responses. In the case of a stay response, regardless of which key was initially pecked, a pigeon was likely still physically closest to that same key after the 1-second interstimulus interval.

The motivation for Experiment 2 was to test the interpretation that pigeons in Experiment 1 had learned to switch based on their history of reinforcement during previous trials. This interpretation was supported, as birds in Experiment 2 adopted a different strategy appropriate to the different probabilities of reinforcement in effect. While this interpretation seems quite intuitive, it is not the only possibility. Consider that the procedure used in Experiment 2 was highly non-natural, in that staying can only be the optimal strategy if the location of the reinforcer is determined *after* a pigeon has made its initial response. This is quite a bit different from the natural world, in which locations of physical things, such as items of food, do not suddenly change as a result of unrelated behaviors at separate locations. The fact that the correct solution to the original MHD is to switch is dependent on this constraint. Thus, if this real-world constraint were a heavy factor in pigeons' performance, it is entirely possible that pigeons might have responded during Experiment 2 just as they did in Experiment 1, by consistently switching. If this had happened, it would have implied that pigeons' optimal behavior in Experiment 1 was based on something other than (or in addition to) reinforcement history.

What then accounts for the difference between the approximately optimal performance of pigeons and the suboptimal performance of humans in parallel situations? One possibility is that Experiment 1 was an *iterated* MHD, involving repeated trials. The MHD is traditionally

proposed as a single choice. The difference is important because of the result just discussed: that pigeons apparently learned to switch *based on feedback* from previous trials. When the MHD is expressed as a single choice, there is no opportunity to learn from previous trials. Thus, it may seem possible (even likely) that human participants who fail to switch on an isolated MHD trial might nevertheless settle into an optimal switching strategy if given the opportunity to learn over multiple trials. Surprisingly, this is not what happens. Granberg and Brown (1995) gave human participants repeated trials of the MHD, finding that while their participants did become more likely to switch, they did not do so consistently and did not approach the optimal strategy of switching on 100% of trials. Instead they settled into a pattern where they distributed about half of their responses each to staying and switching. Even when an incentive was provided for switching, only a small percentage of their participants reached the optimal solution of switching on every trial.

### Experiment 3

The results from the pigeons in Experiments 1 and 2 are interesting, given that most humans fail spectacularly when faced with the MHD. While it may be that pigeons are simply better suited to the problem, there are some other plausible explanations for the differences between the results presented here and the experiments conducted previously with human participants. In particular, one must consider the possible influences of instruction set and amount of training.

First, investigations of the MHD using human participants usually utilize some sort of anecdote or story, which may or may not involve game shows, doors, and goats. In contrast, the pigeons in Experiments 1 and 2 merely saw illuminated response keys and were given the opportunity to learn via trial and error. This difference may be an important one, given that Burns and Wieth (2004) showed that MHD-style problems having the same underlying mathematical structure, but a different cover story were not equally likely to be solved. Furthermore, people can be made to overestimate the probability of a random event, such as winning a lottery, if the random event includes factors from skill situations, such as choice or competition (Langer, 1975). It is important to notice that these same factors are ones that are often emphasized by the usual MHD game show scenario. So while mathematically, a participant's initial selection does not affect the likelihood of winning in the MHD, it may be the case that the participant believes that it does. If so, then the game show anecdote may contribute to the same "illusion of control" shown by Langer, and result in a suboptimal tendency to stay. Pigeons, therefore, may have performed well in Experiments 1 and 2 because they were not subjected to the MHD cover story. If this were the case, humans presented with a stripped-down version of the problem, parallel to the method used in Experiments 1 and 2 here, and unadorned with a story-line might be more likely to settle into the optimal solution.

Second, because pigeons cannot be given verbal instructions, they must learn the procedure through experience and naturally will complete many trials during the process. For example, pigeons in Experiment 1 completed 30 days of training, with up to 100 trials per day. Pigeons in Experiment 2 learned faster, but still completed 15 days of training. In contrast, human investigations of the MHD have involved far fewer trials—anywhere from 1 (Gilovich et al., 1995) to 50 (Granberg & Brown, 1995). While even Granberg & Brown's participants had reached a stable asymptote that had very little variability by the end of 50 trials, it is still possible that some slower learning process could be at work and that if given more training humans would eventually approach the optimal solution.

Experiment 3 investigates the importance of these two factors by testing human participants using a procedure that matches as closely as possible the methods of Experiments 1 and 2. In



particular, no story was presented, and there was a deliberate attempt to avoid the terminology usually associated with the MHD. In addition, participants completed 200 trials—4 times as many as Granberg and Brown's (1995). If a slower learning process is at work, there should be either a higher level of accuracy or evidence of changing response patterns in the later stages of the experiment.

## Method

**Participants**—Thirteen undergraduates were recruited from introductory level psychology classes. Some participants received course credit for participation. Seven participants participated in Condition 1, and six participated in Condition 2. One participant in Condition 1 was eliminated due to prior familiarity with the Monty Hall Dilemma, leaving 6 participants in each condition.

**Apparatus**—A 15 in. flat panel monitor was used to present stimuli. The monitor was equipped with an infrared touch frame (Carroll Touch), which allowed participants to respond by touching the screen. The monitor and touch frame were both interfaced to a personal computer that controlled all experimental events, recorded data, and computed session statistics.

Three response locations were defined on the monitor (left, center, and right). Each was approximately 3 cm × 3 cm, located in a horizontal row with adjacent locations approximately 8 cm apart. Locations were not permanently marked, but the experimental software could present 3 cm square blocks within each to signal when and where it was appropriate for a participant to respond. The touch frame detected responses within any of the three locations, and ignored responses on other areas of the screen.

**Procedure**—Upon arrival, participants were told that they were participating in an experiment on choice and showed how to use the touch screen monitor. Participants were told, “This is an experiment on choice. You will have the opportunity to earn one point per trial by choosing among locations that appear on the touch screen. On each trial you will first choose from among three white squares, then from among two green squares. The computer will provide feedback after each trial. You should try to earn as many points as you can.” They were not told anything about the Monty Hall Dilemma, and care was taken to avoid terms (such as “stay” or “switch”) frequently associated with the MHD. After the general instructions, participants were allowed to complete three practice trials and ask any questions before the experiment began.

An experimental session consisted of a series of 200 trials organized into four blocks of 50. Between each block, participants were allowed to rest for as long as they wished before continuing to the next block. Each trial consisted of an initial choice response from among three options, a second choice response from among two options, a feedback message, and an intertrial interval (ITI).

**Condition 1**—The structure of a trial in Condition 1 mirrored the structure of a trial in Experiment 1. At the beginning of a trial, the computer randomly selected one of the three screen locations as the prize location. Then a white square appeared in all three screen locations. After a participant touched any one of the locations, all three locations were darkened for 1 second. During the 1-second delay, the computer pseudorandomly selected one of the three locations to deactivate for the remainder of the trial, with the constraint that the deactivated location could not be the one selected as the prize location for that trial, nor could it be the one that the participant had touched on that trial. Green squares then appeared in the two remaining locations. After each response, the green squares were erased, and

feedback (the word “win” or “lose”) appeared in the location of the second response. The feedback remained visible for 2 seconds, and then following a 3-second ITI the program moved on to the next trial. As in Experiment 1, switching on every trial would produce the highest possible number of wins.

**Condition 2**—The structure of a trial in Condition 2 mirrored the structure of a trial in Experiment 2. As in Experiment 2, the prize location could only be determined after a participant made an initial choice from among the three white squares. The prize location was determined to be the location that a participant had just touched with a probability of two thirds. The prize location was linked to the two locations that had not been touched with a probability of one sixth each. The remainder of the trial followed the same procedure as Condition 1. As in Experiment 2, staying on every trial would produce the highest possible number of wins.

## Results

**Condition 1**—Figure 4 summarizes the results of Condition 1, and corresponds to Figure 1’s depiction of pigeons’ choices in Experiment 1. Bars represent the average probability of staying and switching during the first block of 50 trials (top panel) and during the final block of 50 trials (bottom panel). During the first block, participants switched on 56.67% of the trials ( $CI = 34.51$  to  $78.83$ ). Note that the confidence interval indicates that performance was indistinguishable from chance performance (50%,  $d = 0.26$ ). On the final block, participants switched on 65.67% of the trials ( $CI = 45.04$  to  $86.29$ ). Again, the confidence interval indicate that performance was no different from chance ( $d = 0.65$ ). While participants switched on 9.00% more trials on Day 30 than on Day 1, the obtained difference was not statistically distinguishable from no difference ( $CI = -1.18$  to  $19.18$ ,  $d = 0.36$ ).

**Condition 2**—Figure 5 summarizes the results of Condition 2, and corresponds to Figure 3’s depiction of pigeons’ choices in Experiment 2. Bars represent the average probability of staying and switching during the first block of 50 trials (top panel) and on the final block of 50 trials (bottom panel). During the first block, participants switched on 30.00% of the trials ( $CI = 13.21$  to  $46.79$ ). The confidence interval indicates that performance was reliably different from chance ( $d = 1.02$ ). On the final block, participants switched on 27.67% of the trials ( $CI = 11.12$  to  $44.22$ ). Again, performance was reliably different from chance ( $d = 1.16$ ). A comparison of the first and final blocks shows that participants switched on 2.33% fewer trials on the final block, a difference that is indistinguishable from no difference ( $CI = -7.73$  to  $12.39$ ,  $d = 0.12$ ).

**Rate of learning**—Figure 6 shows the probability of switching as a function of training. Each data point represents the average of 10 consecutive trials. Reference lines at .67 and .33 correspond to the probability of reinforcement for switching in Conditions 1 and 2, respectively. A strict matching strategy, where the likelihood of switching is equal to the likelihood of reinforcement for switching, would be expected to approximate these probabilities. In both conditions, the probability of switching quickly approaches and hovers near the matching value. Also notice that the majority of learning appears to happen over the first 50 or so trials, and that there was little change over the last 150 trials.

**Comparison with pigeon data**—In both conditions of the present experiment, participants adjusted their likelihood of switching toward the optimal strategy. Unlike pigeons, their eventual likelihood of switching fell well short of optimality. Human participants on the first block of trials in Condition 1 switched 20.33% more than pigeons on Day 1 of Experiment 1, but the difference was not statistically reliable ( $CI = -1.20$  to  $41.86$ ,  $d = 0.91$ ). In contrast, those human participants were 30.67% less likely to switch on the last

block of trials than pigeons were on the last day of training ( $CI = 12.67$  to  $48.66$ ,  $d = 1.54$ ). That is, despite starting with approximately the same likelihood of switching, pigeons reached a point where their performance was significantly closer to optimal.

In Condition 2, the optimal strategy was to stay on every trial, and pigeons again showed a better ability to shift toward the optimal strategy. Human participants on the first block of trials in Condition 2 showed a probability of switching that was only 0.17% lower than the pigeons on the first day of training in Experiment 2, a difference that is not statistically reliable ( $CI = -18.35$  to  $18.01$ ,  $d = 0.01$ ). In contrast, those human participants were 23.33% more likely to switch (and less likely to stay) on the last block of trials than pigeons were on the last day of training, a difference that is statistically reliable ( $CI = 8.46$  to  $38.21$ ,  $d = 1.45$ ). Again, pigeons and humans began with the same likelihood of switching, but with training pigeons came significantly closer to the optimal strategy.

## Discussion

The results of the present experiment seem to undermine the plausibility of the two alternative explanations advanced earlier. Even without the game show (or any other) story, and with 200 trials of practice, humans failed to adopt the optimal strategies that the pigeons in Experiments 1 and 2 quickly mastered. Despite the lack of improvement over the last 150 trials, we do not necessarily discount the possibility that participants may have been learning in the latter stages of the experiment, or that with further training they could have moved nearer to the optimal solutions. However, such a learning rate, if found, would be slower than in most other probability learning experiments (e.g., Estes, 1961) and even slower than many characteristically slow forms of learning, such as complex multidimensional category learning (e.g., Ashby & Maddox, 1992) and implicit learning (e.g., Reber, 1967).

The fact that the results of this experiment so closely matched the overall patterns of previous human research on the MHD could indicate that the same cognitive processes are at work in the present experiment. Apparently, pigeons' superior performance in Experiments 1 and 2 is not related to the instruction set or number of trials. To what then, might we attribute the success of the pigeons in Experiments 1 and 2? So far, the pigeon results are consistent with maximization. Specifically, pigeons learn to adopt a response strategy that maximizes the likelihood of winning on each trial, as they often do in probability learning experiments. If a maximizing strategy is indeed responsible for pigeons' performance, then one might see a different pattern of responding if there were no optimal response - that is, if neither switching nor staying resulted in a higher likelihood of reinforcement. This possibility is the subject of the next experiment.

## Experiment 4

Recall that Experiments 1 and 2 showed that pigeons can perform optimally on an iterated MHD, and that optimal responding is shaped by previous experience. If this is the case and responding is based on feedback from earlier trials, then ambiguous feedback should eliminate the consistency seen in Experiments 1 and 2. This is easily done by changing the procedure so that determination of the prize location is truly random between the two nonselected options (including the possibility of removing the prize). With this small change, the previously optimal solution no longer holds true. Instead, the probabilities of reinforcement associated with staying and switching are both equal to one third. In such a case, where there is no optimal strategy, pigeons might no longer adopt a consistent staying or switching strategy. Experiment 4 investigates the consequences of this particular variation, where neither switching nor staying offers a relative advantage over the other.

## Method

**Subjects**—Six Silver King pigeons (*Columba livia*) were obtained from a local breeder. They were maintained under the same conditions as those in Experiments 1 and 2.

**Apparatus**—The operant chambers were the same as those used for Experiments 1 and 2.

**Procedure**—Birds were pretrained as in Experiments 1 and 2 until consistent responding was achieved on each of the three keys.

A daily experimental session consisted of a series of individual trials (10 trials on the initial day, gradually increased to 100 over the first five days of the experiment). Each trial consisted of an initial choice response from among three options, a second choice response from among two options, either reinforcement or correction depending on whether the second response was correct, and an intertrial interval.

The procedure was different from Experiment 1 in the way that the deactivated key was determined. In Experiment 4, the prize location was determined before a pigeon made its initial choice (from among the three white keys), as it was in Experiment 1. However, selection of the deactivated key was not subject to the same constraints. The deactivated key could still not be the key that a pigeon initially pecked on that trial, but it could be the key that had been selected as the prize location (in which case, no reinforcement was available, regardless of which key was pecked).

Again, the critical result for each trial was whether a pigeon stayed (pecked the same key twice) or switched (pecked different keys) on each trial. In Experiment 4, the probabilities of reinforcement associated with staying and switching are the same: staying would lead to reinforcement one third of the time, and switching would also lead to reinforcement one third of the time. Thus, there is no statistical advantage to either option.

## Results

Figure 7 shows the average probability of staying and switching on the first day of the experiment (top panel) and on the thirtieth and final day of the experiment (bottom panel). On Day 1, pigeons switched on 33.00% of the trials ( $CI = 19.68$  to  $46.32$ ). This performance was reliably different from chance (50%,  $d = 1.09$ ). On Day 30, pigeons switched on 24.00% of the trials ( $CI = -16.01$  to  $64.01$ ). In this case, performance was not reliably different from chance ( $d = 0.56$ ). A comparison of the first and final days shows that after 30 days, pigeons' tendency to stay was 9.00% lower, but the difference was indistinguishable from no difference ( $CI = -33.86$  to  $51.86$ ,  $d = 0.30$ ).

Despite the similarity in the overall probability of switching on Days 1 and 30, there was considerably more between-bird variability on Day 30 ( $s = 38.13$ ) than there was on Day 1 ( $s = 12.70$ ), and this additional variability explains why the tendency to switch was no longer statistically reliable on Day 30 despite being stronger on average than it was on Day 1.

Figure 8 depicts the source of variability by showing acquisition curves for three of the six birds (birds 10, 11, and 12; the remaining three are very similar to the curve shown for bird 12 and have been omitted to reduce clutter). Notice that each bird shows a markedly different acquisition curve: Bird 10 eventually switched on virtually every trial (100% switching on Day 30), while bird 12 stayed on virtually every trial (2% switching on Day 30), and bird 11 did not settle into any consistent pattern (40% switching on Day 30).

## Discussion

While they require a different sort of analysis, the results of Experiment 4 are still consistent with the interpretations provided for Experiments 1 and 2. Specifically, in all three experiments, tendencies to stay or switch were determined by the outcomes of previous trials. In Experiment 1, switching produced the best rate of reinforcement, and birds consistently switched. In Experiment 2, staying produced the best rate of reinforcement, and birds consistently stayed. In Experiment 4, staying and switching had no effect on the rate of reinforcement, and there was no consistent pattern of staying or switching across (and sometimes within) birds.

The behavior of the pigeons in Experiment 4 is somewhat reminiscent of B.F. Skinner's (1948) well known demonstration of superstition, in which pigeons were reinforced at arbitrary intervals, and each subsequently adopted a different behavior based on what they happened to be doing when reinforcement was delivered. Similarly, in Experiment 4, pigeons were reinforced at the same rate, regardless of whether they chose to switch or stay. If a bird was reinforced after choosing to switch (or stay) on a particular trial, switching (or staying) became more likely and was consequently more likely to be reinforced again. In this way, the outcome of early trials in the experiment may have randomly set a trajectory in motion, resulting in different birds adopting either a switching or staying strategy.

Figure 9 shows data relevant to this explanation. Bars show the percent of all reinforcers that were earned for switching during the earliest experimental sessions (Days 1–5 and 6–10) and during the final session. Notice that the birds (birds 12–16) that settled into a consistent tendency to stay earned the majority of reinforcers for staying during Sessions 1–5, and that proportion became even greater as the experiment continued. Similarly, the one bird (bird 10) that settled into a consistent tendency to switch earned the majority of early reinforcers for switching, and that pattern became stronger in subsequent sessions. Finally, the one bird (bird 11) that did not show a strong tendency either way earned approximately equal numbers of reinforcers for switching and staying early in the experiment.

## General Discussion

Taken together, these experiments show that pigeons can learn to respond optimally in a simulation of the MHD. Furthermore, they also suggest that birds learn to respond optimally based on feedback received from completed trials. In particular, a bird's response strategy can be changed by adjusting the relative likelihoods of reinforcement for switching and staying. The surprising implication is that pigeons seem to solve the puzzle, arriving at the optimal solution while most humans do not.

The MHD seems to be particularly tough for humans, for whom suboptimal choices persist in a wide variety of situations. Granberg (1999b) demonstrated that the suboptimal tendency to stay in the MHD is persistent and appears cross-culturally. Students in Brazil, China, Sweden, and the United States all showed a reliable tendency to stay. Thus, people seem compelled to stay with their initial choice, even when it reduces expected winnings. One possible source of this suboptimality may be the heuristics that are acquired as a part of normal cognitive development. DeNeys (2006) investigated the MHD in students of varying ages, finding that university students almost universally believed that staying and switching were equally likely to win. Younger students showed a progressively weaker belief in equiprobability. Only in the youngest group that was tested (8th graders) did even a small subset of students correctly discern that switching was the most profitable strategy. This finding is consistent with other investigations of suboptimal choice behaviors. For example, Langer and Imber (1979) have shown that overpractice can produce a seemingly paradoxical decrement in performance. They characterize the decrement as a negative consequence of



the otherwise adaptive tendency for overlearning to free up limited attentional resources. Thus, the benefits that come with extensive education may partly be due to the acquisition of heuristics that while efficient, can interfere with some kinds of performance, perhaps including the MHD.

A second potentially important implication relates to the phenomenon of probability matching. When presented repeatedly with two choices, each having a particular likelihood of reinforcement, many species, including humans, tend to match the proportion of their choices to the arranged probabilities (Herrnstein, 1970). In the MHD, the probability of reinforcement for switching is two thirds, while the probability of reinforcement for staying is one third. Thus, some probability matching theories might predict that responses would approximate those proportions: two thirds switching and one third staying. This simple matching theory actually describes the performance of human participants in both conditions of Experiment 3 quite well. Pigeons, however, did something entirely different: instead of matching, they maximized. So while human participants did adopt a consistent pattern of responding that was based on reinforcement probabilities, it was not even close to the optimal one quickly adopted by pigeons.

We propose that the curious difference between pigeon and human behavior described above might parallel the difference between classical and empirical probability. Classical probability is an approach to probability that is based on a complete a priori analysis of the situation. It requires no actual data collection, as when one states the probability of tossing “tails” on a fair coin as one half or the probability of rolling “boxcars” at the craps table as 1 in 36. Empirical probability on the other hand is based on observation and estimates the probability of an event as the relative frequency with which that event has occurred in the past. The MHD can potentially be solved through either approach. A classical probability approach would involve considering every possible outcome and assigning a probability value to each. An empirical approach would involve running through several MHD trials and tracking the outcomes as they accumulate. Humans seem to rely on classical probability theory in attempting to solve the MHD. However, in most cases they do so inappropriately (Franco-Watkins, Derks, & Daugherty, 2003; Tubau & Alonso, 2003), reaching the conclusion that switching and staying should win equally often. Even trained mathematicians seem to have difficulty correctly analyzing the MHD using classical probability. When the MHD appeared in the “Ask Marilyn” column in *Parade* magazine (Vos Savant, 1990), along with an explanation of the solution, columnist Marilyn Vos Savant received 10,000 letters (many from university addresses), 92% of which disagreed with her solution. Even Paul Erdos, perhaps the most prolific mathematician in history, initially did not accept the explanation for why switching is better than staying.

Pigeons, on the other hand, likely use empirical probability to solve the MHD and appear to do so quite successfully. Pigeons might not possess the cognitive framework for a classical probability-based analysis of a complicated problem like the MHD, but it is certainly not far-fetched to suppose that pigeons can accumulate empirical probabilities by observing the outcomes of numerous trials and adjusting their subsequent behavior accordingly. The MHD then seems to be a problem that is quite challenging from a classical approach, but not necessarily from an empirical approach, and the difference may give pigeons something of an advantage over humans. The aforementioned mathematician Paul Erdos demonstrates this proposition nicely. According to his biography (Hoffman, 1998), Erdos refused to accept colleagues’ explanations for the appropriate solution to the MHD that were based on classical probability. He was eventually convinced only after seeing a simple Monte Carlo computer simulation that demonstrated beyond any doubt that switching was the superior strategy. Until he was able to approach the problem like a pigeon—via empirical probability—he was unable to embrace the optimal solution.



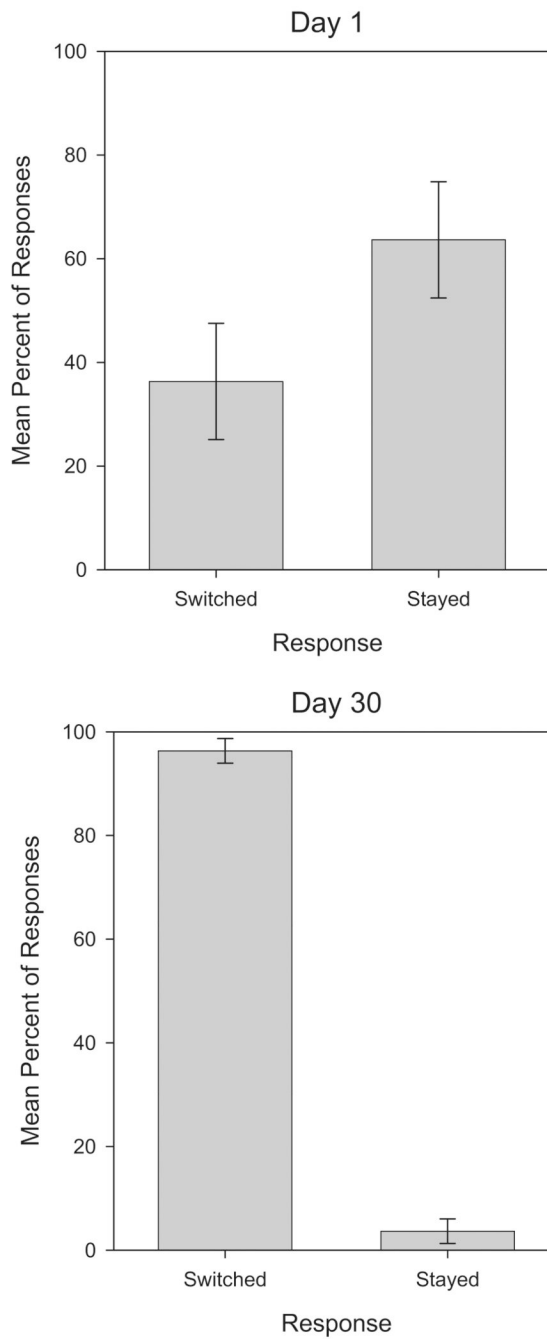
## Acknowledgments

This research was supported by a grant from the National Institutes of Health.

## References

- Arkes HR, Ayton P. The sunk cost and Concorde effects: Are humans less rational than lower animals? *Psychological Bulletin*. 1999; 125:591–600.
- Ashby FG, Maddox WT. Complex decision rules in categorization: Contrasting novice and experienced performance. *Journal of Experimental Psychology: Human Perception and Performance*. 1992; 18(1):50–71.
- Brown PL, Jenkins HJ. Autoshaping of the pigeon's keypeck. *Journal of the Experimental Analysis of Behavior*. 1968; 11:1–8. [PubMed: 5636851]
- Bullock DH, Bitterman ME. Probability-matching in the pigeon. *American Journal of Psychology*. 1962; 75:634–639. [PubMed: 14016794]
- Burns BD, Wieth M. The collider principle in causal reasoning: Why the Monty Hall dilemma is so hard. *Journal of Experimental Psychology: General*. 2004; 133(3):434–449. [PubMed: 15355148]
- DeNeys, W. Developmental trends in decision making: The case of the Monty Hall dilemma. In: Ellsworth, JA., editor. *Psychology of decision making in education*. Hauppauge, NY: Nova Science Publishers; 2006.
- Estes WK. A descriptive approach to the dynamics of choice behavior. *Behavioral Science*. 1961; 6:177–184. [PubMed: 13697447]
- Fantino E, Asafandiar A. Probability matching: Encouraging optimal responding in humans. *Canadian Journal of Experimental Psychology*. 2002; 56(1):58–63. [PubMed: 11901961]
- Franco-Watkins AM, Derks PL, Dougherty MRP. Reasoning in the Monty Hall problem: Examining choice behaviour and probability judgements. *Thinking and Reasoning*. 2003; 9(1):67–90.
- Gilovich T, Medvec VH, Chen S. Commission, omission, and dissonance reduction: Coping with regret in the “Monty Hall” problem. *Personality and Social Psychology Bulletin*. 1995; 21(2):182–190.
- Goodie AS, Fantino E. Learning to commit or avoid the base-rate error. *Nature*. 1996; 380:247–249. [PubMed: 8637572]
- Granberg D. A new version of the Monty Hall dilemma with unequal probabilities. *Behavioural Processes*. 1999a; 48:25–34.
- Granberg D. Cross-cultural comparison of responses to the Monty Hall dilemma. *Social Behavior and Personality*. 1999b; 27(4):431–438.
- Granberg D, Brown TA. The Monty Hall dilemma. *Personality and Social Psychology Bulletin*. 1995; 21(7):711–723.
- Hartl JA, Fantino E. Choice as a function of reinforcement ratios in delayed matching to sample. *Journal of the Experimental Analysis of Behavior*. 1996; 66:11–27. [PubMed: 16812815]
- Herbranson WT, Fremouw T, Shimp CP. The randomization procedure in the study of categorization of multidimensional stimuli by pigeons. *Journal of Experimental Psychology: Animal Behavior Processes*. 1999; 25(1):113–135. [PubMed: 9987862]
- Herbranson WT, Fremouw T, Shimp CP. Categorizing a moving target in terms of its speed, direction, or both. *Journal of the Experimental Analysis of Behavior*. 2002; 78(3):249–270. [PubMed: 12507003]
- Herrnstein RJ. Relative and absolute strength of response as a function of frequency of reinforcement. *Journal of the Experimental Analysis of Behavior*. 1961; 4:267–272. [PubMed: 13713775]
- Herrnstein RJ. On the law of effect. *Journal of the Experimental Analysis of Behavior*. 1970; 13:243–266. [PubMed: 16811440]
- Herrnstein, RJ. *The Matching Law*. Cambridge, MA: Harvard University Press; 1997.
- Hoffman, P. *The man who loved only numbers: The story of Paul Erdos and the search for mathematical truth*. New York, NY: Hyperion; 1998.


- Houston A. The matching law applies to wagtails foraging in the wild. *Journal of the Experimental Analysis of Behavior*. 1986; 45(1):15–18. [PubMed: 16812441]
- Langer EJ. The illusion of control. *Journal of Personality and Social Psychology*. 1975; 33(2):311–328.
- Langer EJ, Imber LG. When practice makes imperfect: Debilitating effects of overlearning. *Journal of Personality and Social Psychology*. 1979; 37(11):2014–2024. [PubMed: 521900]
- Maddox WT, Bohil CJ. Probability matching, accuracy maximizing and a test of the optimal classifier's independence assumption in perceptual categorization. *Perception & Psychophysics*. 2004; 66(1):104–118. [PubMed: 15095944]
- Mazur JE. Optimization theory fails to predict performance of pigeons in a two-response situation. *Science*. 1981; 214:823–825. [PubMed: 7292017]
- Posner G. Nations mathematicians guilty of “innumeracy.”. *Skeptical Inquirer*. 1991; 15:342–345.
- Reber AS. Implicit learning of artificial grammars. *Journal of Verbal Learning and Verbal Behavior*. 1967; 6:855–863.
- Shimp CP. Probabilistically reinforced choice behavior in pigeons. *Journal of the Experimental Analysis of Behavior*. 1966; 9(4):443–455. [PubMed: 5961513]
- Shimp CP. Probabilistic discrimination learning in the pigeon. *Journal of Experimental Psychology*. 1973; 97(3):292–304.
- Skinner BF. “Superstition” in the pigeon. *Journal of Experimental Psychology*. 1948; 38:168–172. [PubMed: 18913665]
- Stephens, DW.; Krebs, JJ. *Foraging theory*. Princeton, NJ: Princeton University Press; 1986.
- Tierney J. Behind Monty Hall's doors: Puzzle, debate and answer? *New York Times*. 1991 July 21.:1–20.
- Tubau E, Alonso D. Overcoming illusory differences in a probabilistic counterintuitive problem: The role of explicit representations. *Memory & Cognition*. 2003; 31(4):596–607.
- Tversky A, Kahneman D. The framing of decisions and the psychology of choice. *Science*. 1981; 211:453–458. [PubMed: 7455683]
- Vos Savant M. Ask Marilyn. *Parade*. 1990 September 9.:15.



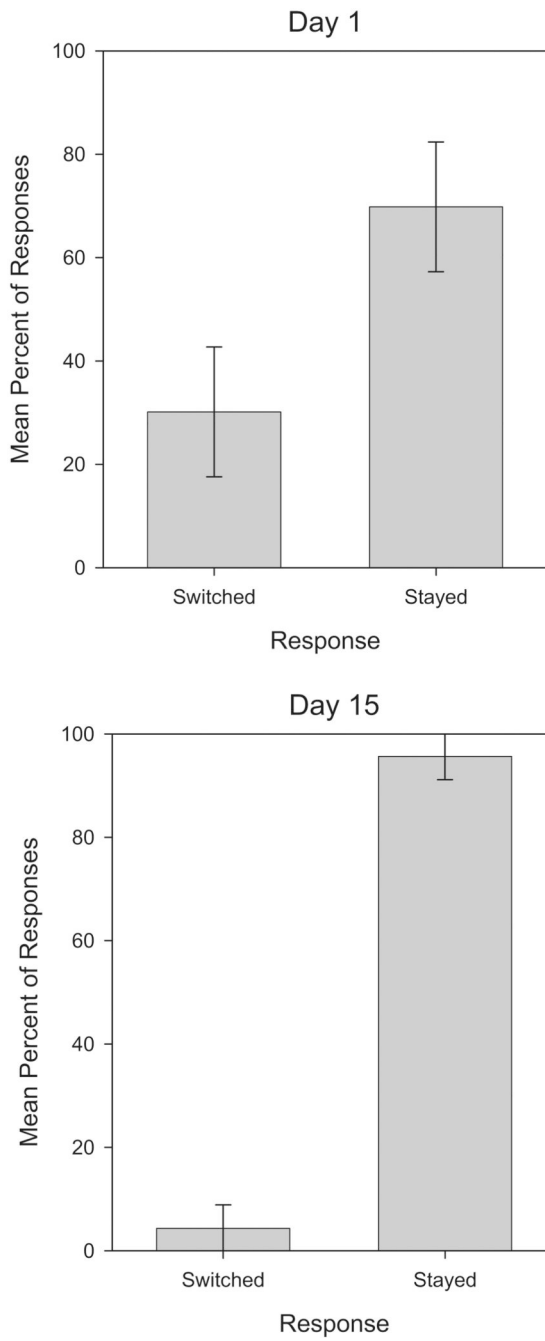
**Figure 1.** Proportion of switching (left bars) and staying (right bars) responses, averaged across all birds in Experiment 1. The top panel represents the first day of training, and the bottom panel represents the 30th day of training. Error bars represent 95% confidence intervals.

## HERBRANSON AND SCHROEDER

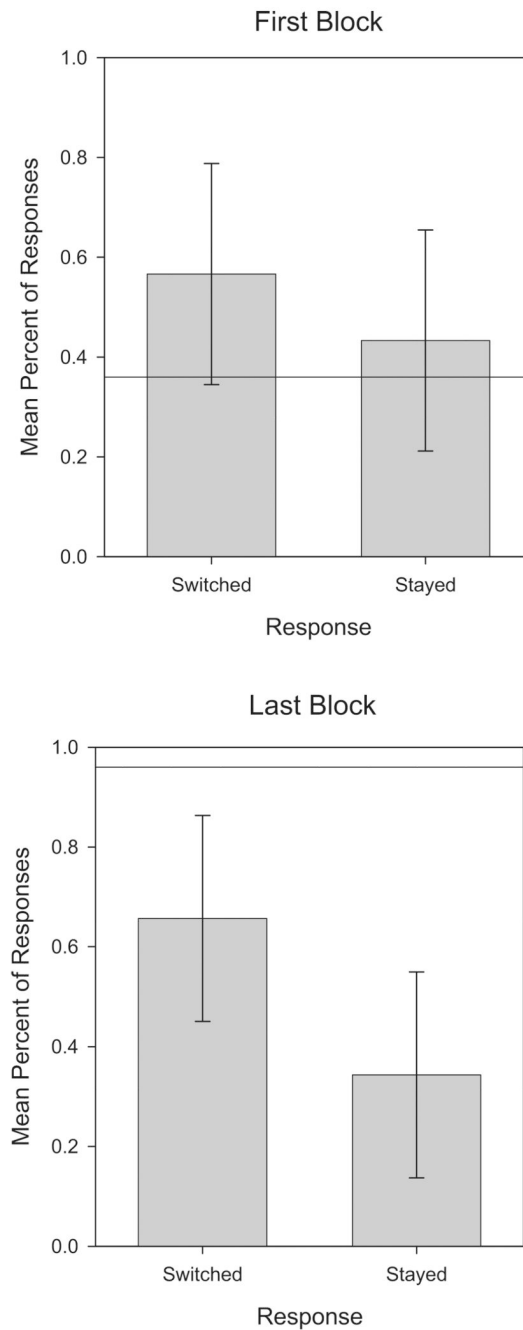
## Reinforcers Earned



**Figure 2.** Percent of reinforced trials on Day 1 (left bar) and Day 30 (right bar) averaged across all birds in Experiment 1. Reference lines show the expected performance for a strategy of always switching (67%), guessing randomly (50%), and always staying (33%). Error bars represent 95% confidence intervals.




**Figure 3.** Proportion of switching (left bars) and staying (right bars) responses, averaged across all birds in Experiment 2. The top panel represents the first day of training, and the bottom panel represents the 15th day of training. Error bars represent 95% confidence intervals.




**Figure 4.** Proportion of switching (left bars) and staying (right bars) responses, averaged across all human participants in Condition 1 of Experiment 3. The top panel represents the first 50 trials, and the bottom panel represents the final 50 trials. Error bars represent 95% confidence intervals. Reference lines show the probability of switching by pigeons during the first (top panel) and final (bottom panel) days of training in Experiment 1.



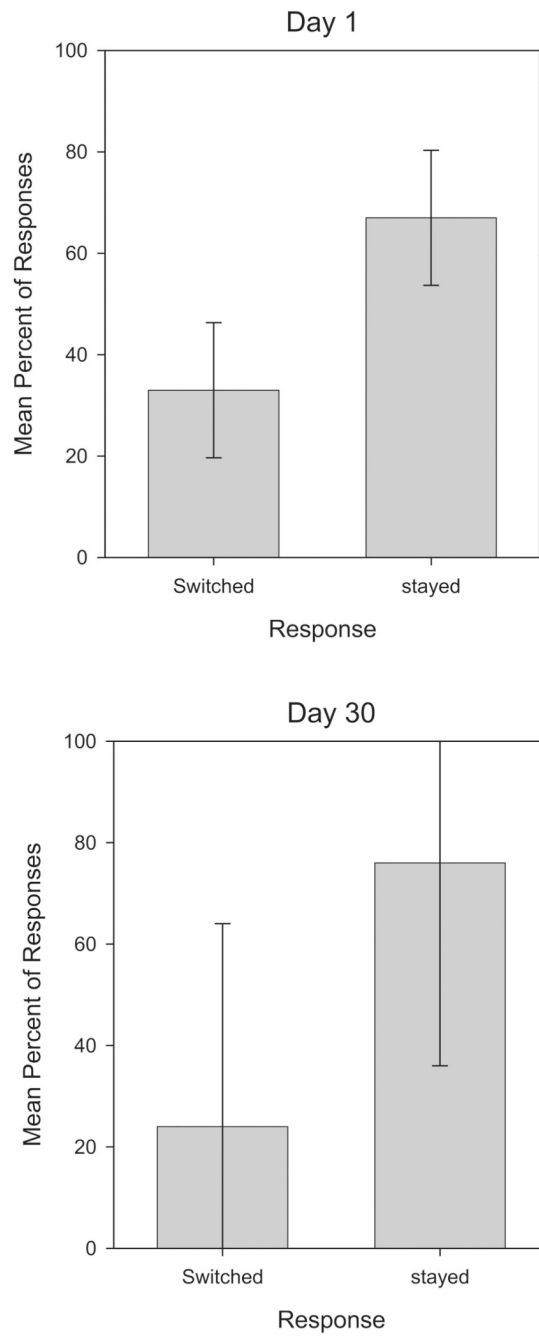


**Figure 5.** Proportion of switching (left bars) and staying (right bars) responses, averaged across all human participants in Condition 2 of Experiment 3. The top panel represents the first 50 trials, and the bottom panel represents the final 50 trials. Error bars represent 95% confidence intervals. Reference lines show the probability of switching by pigeons during the first (top panel) and final (bottom panel) days of training in Experiment 2.

## Probability of Switching Across Trials




**Figure 6.** Probability of switching in both conditions of Experiment 4 across trials. Each data point represents the average of 10 consecutive trials. Error bars represent one standard error.




**Figure 7.** Proportion of switching (left bars) and staying (right bars) responses, averaged across all birds in Experiment 4. The top panel represents the first day of training, and the bottom panel represents the 30th day of training. Error bars represent 95% confidence intervals.

## Probability of Switching Across Days of Training



**Figure 8.** Daily probability of switching for three individual birds (birds 10, 11, and 12) in Experiment 4.



**Figure 9.** Percent of reinforcers earned for staying, for all six birds in Experiment 4. Left (dark) bars represent the first five days of training. Middle (light) bars represent the second five days of training. Right (medium) bars represent the 30th and final day of training.

**Table 1**

Possible Response Paths and Individual Bird Responses for Experiment 1

	<i>L-L</i>	<i>L-C</i>	<i>L-R</i>	<i>C-L</i>	<i>C-C</i>	<i>C-R</i>	<i>R-L</i>	<i>R-C</i>	<i>R-R</i>	<i>Stay</i>	<i>Switch</i>
Bird 1											
First	.34	.10	.02	.00	.28	.04	.02	.06	.14	.76	.24
Last	.01	.51	.47	.00	.00	.01	.00	.00	.00	.01	.99
Bird 2											
First	.20	.10	.14	.04	.16	.16	.06	.02	.12	.48	.52
Last	.00	.52	.48	.00	.00	.00	.00	.00	.00	.00	1.00
Bird 3											
First	.00	.06	.00	.06	.04	.00	.14	.10	.60	.64	.36
Last	.01	.59	.40	.00	.00	.00	.00	.00	.00	.01	.99
Bird 4											
First	.20	.14	.10	.00	.26	.00	.14	.08	.08	.54	.46
Last	.07	.50	.43	.00	.00	.00	.00	.00	.00	.07	.93
Bird 5											
First	.00	.06	.14	.00	.06	.02	.00	.08	.64	.70	.30
Last	.00	.00	.00	.01	.05	.00	.40	.53	.01	.06	.94
Bird 6											
First	.48	.04	.10	.14	.06	.00	.06	.04	.08	.62	.38
Last	.00	.00	.00	.00	.00	.00	.47	.48	.05	.05	.95
Avg.											
First	.20	.08	.08	.04	.14	.04	.07	.06	.28	.62	.38
Last	.02	.35	.30	.00	.01	.00	.15	.17	.01	.03	.97

*Note.* Numbers in italics represent stay responses.