

BRIEF REPORT

Hot-Hand Bias in Rhesus Monkeys

Tommy C. Blanchard
University of Rochester

Andreas Wilke
Clarkson University

Benjamin Y. Hayden
University of Rochester

Human decision-makers often exhibit the *hot-hand phenomenon*, a tendency to perceive positive serial autocorrelations in independent sequential events. The term is named after the observation that basketball fans and players tend to perceive streaks of high accuracy shooting when they are demonstrably absent. That is, both observing fans and participating players tend to hold the belief that a player's chance of hitting a shot are greater following a hit than following a miss. We hypothesize that this bias reflects a strong and stable tendency among primates (including humans) to perceive positive autocorrelations in temporal sequences, that this bias is an adaptation to clumpy foraging environments, and that it may even be ecologically rational. Several studies support this idea in humans, but a stronger test would be to determine whether nonhuman primates also exhibit a hot-hand bias. Here we report behavior of 3 monkeys performing a novel gambling task in which correlation between sequential gambles (i.e., temporal clumpiness) is systematically manipulated. We find that monkeys have better performance (meaning, more optimal behavior) for clumped (positively correlated) than for dispersed (negatively correlated) distributions. These results identify and quantify a new bias in monkeys' risky decisions, support accounts that specifically incorporate cognitive biases into risky choice, and support the suggestion that the hot-hand phenomenon is an evolutionary ancient bias.

Keywords: cognitive bias, ecological rationality, foraging, hot hand, rhesus macaque

Humans are surprisingly poor at dealing with randomness. We have trouble distinguishing random sequences from ordered ones, are poor at generating random choices, and have a strong tendency to perceive patterns in random data (e.g., Falk & Konold, 1997; Nickerson, 2002). One of the best known examples of our limitations in the face of randomness is the *hot-hand phenomenon*—the tendency to perceive illusory streaks in sequential events whose probabilities are in reality independent.

This phenomenon was first noticed in basketball. Both players and fans judge a player's chance of hitting a shot to be greater following a successful shot than a miss, even though shots are

almost always statistically independent (Gilovich, Vallone & Tversky, 1985; cf. Bar-Eli, Avugos, & Raab, 2006). In recent years, psychologists have further delved into the *proximate mechanisms* of the hot-hand phenomenon and its opposite the gambler's fallacy (e.g., Ayton & Fischer, 2004; Oskarsson, van Boven, McClelland, & Hastie, 2009) and explored the role of the hot-hand bias in other sports (e.g., Avugos, Köppen, Czienskowski, Raab, & Bar-Eli, 2013; Raab, Gula, & Gigerenzer, 2012), its role in gambling behavior (e.g., Croson & Sundali, 2005), and investigated it in different age (Castel, Drolet Rossi, & McGillivray, 2012), and clinical populations (Wilke, Scheibehenne, Gaissmaier, McCanney, & Barrett, in press).

Recent research on the *ultimate function* of the hot-hand phenomenon, however, suggests that it is an adaptive human universal, tied to an evolutionary history of foraging for clumpy resources rather than an erroneous cognitive fallacy that only occurs in sports or financial settings (Wilke & Barrett, 2009; Wilke & Mata, 2012; Wilke & Todd, 2012; see also Reifman, 2011). In a cross-cultural study, Wilke and Barrett (2009) found that the hot-hand phenomenon occurs in both Western cultures and traditional foraging cultures, and that it seems to be a kind of psychological default which is only partly erased by experience with true randomizing mechanisms like coin tosses. In their computerized tasks, Wilke and Barrett (2009) had participants predict hits and misses when foraging for various kinds of natural resources (e.g., fruits) and man-made artificial ones (e.g., parking spots). Although

Tommy C. Blanchard, Department of Brain and Cognitive Sciences, University of Rochester; Andreas Wilke, Department of Psychology, Clarkson University; Benjamin Y. Hayden, Department of Brain and Cognitive Sciences, University of Rochester.

We are grateful to Marc Mancarella, Krina Patel, Alberto Sepulveda Rodriguez, Mariah Meyer, and Meghan Castagno for help in data collection. This research was supported by a NSF CAREER award (1253576) and a Brain and Behavior Research Foundation Award to Benjamin Y. Hayden.

Correspondence concerning this article should be addressed to Tommy C. Blanchard, Department of Brain and Cognitive Sciences, University of Rochester, Meliora Hall 360, Box 270268, Rochester, NY 14627. E-mail: tblanchard@bcs.rochester.edu

all resource distributions were generated randomly, participants exhibited the hot-hand phenomenon across all the resource types with the strongest effects occurring for resource types that relate to natural kinds.

This finding is less puzzling once one considers that aggregation in space and time, rather than a random distribution, is likely to have been the norm for most of the natural resources humans encountered over evolutionary time. Natural resources that primates forage for, such as specific plants and animals, rarely distribute themselves in a purely random manner in their natural environment because individual organisms are not independent of one another (Taylor, 1961; Taylor, Woiwod, & Perry, 1978; cf. Hutchinson, Wilke, & Todd, 2008). Although these deviations can be in the direction of greater dispersal, most often, these deviations from randomness are in the direction of aggregation, because aggregation offers considerable benefits such as a common habitat, mating and parenting, or the benefits of group foraging (Bell, 1991; Krause & Ruxton, 2002). Because humans have been hunters and gatherers for a very long part of our history, it could well be that our evolved psychology is adapted to assume such aggregated resource distributions as the default (Tooby & DeVore, 1987).

Indeed, the hot-hand bias might be adaptive in contexts where clumps sometimes exist but dispersal is rare, as in foraging (cf. Haselton et al., 2009). When trying to predict the best foraging site, using a strong prior expectation for clumped resources is likely to provide better guesses than a random prior. At the same time, when faced with sequences of independent and equiprobable events, the hot-hand bias does not decrease accuracy, because all strategies produce chance-level performance (see Scheibehenne, Wilke, & Todd, 2011). Thus, what has been seen as a systematic error in our decision-making apparatus may actually be a design feature of our cognitive system to help us estimate the locations of forageable resources in physical environments. This explanation highlights the role of ecological (and evolutionary) rationality—the principle that there is a match between the statistical structure of objects and information of current (and past) environments and the judgment and decision-making strategies of humans and other organisms (e.g., Fawcett et al., 2014; Todd & Gigerenzer, 2012).

One prediction of this idea is that nonhuman species with similar foraging histories might share our proclivity for positive recency in independent sequential events. Although there has been research on decision-making biases in nonhuman primates (e.g., Lakshminarayanan, Chen, & Santos, 2008; Stevens, 2010) and the match between decision environments and animal cognition (e.g., Houston, McNamara, & Steer, 2007; Hutchinson & Gigerenzer, 2005), we know of no direct test of the hot-hand phenomenon in other animals. In uncertain environments, for instance, monkeys do exhibit a bias toward choosing an option again after it has resulted in a reward, and shifting away from an option again after it has failed to result in a reward—a win–stay, lose–shift strategy, the equivalent of a hot-hand bias in these tasks (Barraclough, Conroy, & Lee, 2004; Hayden, Heilbronner, Nair, & Platt, 2008; Hayden, Nair, McCoy, & Platt, 2008; Lau & Glimcher, 2005). In most tasks, this bias is costless, so it may simply reflect a tendency to overweight the recent past in choices, rather than an assumption that the environment is positively clumped. In the current article, however, we directly test for hot-hand biases in monkeys by setting up a task in which rewards are in some conditions posi-

tively correlated and in others negatively correlated. Thus, in this task there should rationally be no expectations of a positively correlated reward structure. We hypothesized that if monkeys have an innate expectation for resources to be clumpy we should see behavioral asymmetries in the adoption of the optimal behavior in positive versus negative correlation conditions.

We examined behavior of three monkeys on a novel two-option risky choice task. On each trial one option offered a reward and the other did not. The chance the rewarded option would switch was set to one of 10 values and varied across days. Monkeys' patterns of choices were consistent with the hypothesis that they see more positive autocorrelation than there is, across all conditions. These results are concordant with previous human results and suggest that monkeys' risky choices are determined by an inherent bias toward expectations of environmental clumpiness.

Method

Subjects and General Method

All procedures were approved by the University of Rochester Institutional Animal Care and Use Committee and were designed and conducted in compliance with the Public Health Service's Guide for the Care and Use of Animals. Three male rhesus monkeys (*Macaca mulatta*) served as subjects. The subjects used were juvenile rhesus macaques (all 3- to 5-years-old). Initially, each animal was provided with a small mount to facilitate head positioning using standard techniques (see Blanchard, Pearson, & Hayden, 2013). Subjects were habituated to laboratory conditions and trained to perform decision tasks for liquid reward. Prior to the experiments we report here, the subjects had been trained to perform oculomotor tasks for liquid (juice or water) rewards through standard reinforcement training. They all had extensive experience with standard gambling tasks, although not with this specific task. Gaze position was collected with an Eyelink 1000 Camera System (SR Research, Osgoode, Canada). A standard solenoid valve controlled the delivery of liquid reward.

Correlated Outcomes Task

To investigate the hot-hand phenomenon in monkeys, we created a novel *correlated outcomes task* (see Figure 1). On each trial of the task, a fixation spot first appeared in the center of the monitor. Once the subject acquired fixation and maintained it for 150 ms, two targets appeared, 275 pixels to the left and right of the central spot. Both stimuli were vertically oriented rectangles (80 pixels wide, 300 pixels tall) containing easily distinguishable photographs. Within a condition, we used two different emotionally neutral nature scene photographs. To help the animals differentiate between the conditions (see below for an explanation of the different conditions), a unique pair of pictures was used for each condition. The two stimuli always appeared in the same locations. Following their presentation, the subject selected a target by shifting gaze toward it. On each trial, only one of the two targets was rewarded. If the subject chose the rewarded target, they would receive an aliquot of water as a reward (0.15 mL). Regardless of what was chosen, a green circle would appear over the stimulus that was rewarded that trial, and a red circle would appear over the stimulus that was unrewarded; these would remain visible for 1 s.

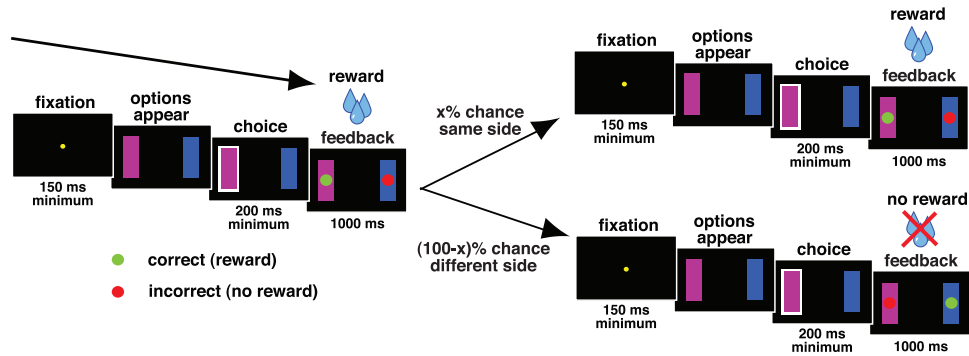


Figure 1. Task schematic of the correlated outcomes task. See text for details. The color version of this figure appears in the online article only.

A 100 ms intertrial interval followed, during which the screen was blank (see Figure 1 for task schematic). We used Psychtoolbox to control visual display of the behavioral task (Brainard, 1997) and EyeLink Toolbox for processing of eye position data (Cornelissen, Peters, & Palmer, 2002).

If the animal chose the rewarded target on the previous trial, the same side would be rewarded with a probability of 10% to 90%, depending on the condition being tested. Nine conditions were tested, four levels of negative autocorrelation (-4 , -3 , -2 , -1 , equivalent to a 10%, 20%, 30%, or 40% chance, respectively, of the same side being rewarded on the subsequent trial), no autocorrelation (0, equivalent to a 50% chance of the same side being rewarded on the subsequent trial), and four levels of positive autocorrelation ($+1$, $+2$, $+3$, $+4$, equivalent to a 60%, 70%, 80%, 90% chance, respectively, of the same side being rewarded on the subsequent trial). The optimal strategy in the negatively correlated condition is to always switch after a win. The optimal strategy in the positive correlation condition is to always stay after a win. There is no optimal strategy in the no autocorrelation condition. Only one condition was run within a daily session. The order of conditions was randomly varied within and between subjects. We added one feature to the task to reduce perseveration and random guessing. If the animal chose the unrewarded target on one trial, then we would repeat the location of the rewarded target on the following trial; the trial would thus be identical to the previous, to force the subject to switch to progress and prevent the animal from continuing to perseverate on a single target. All of our analyses were only of behavior following a rewarded trial.

We chose a different random ordering of the nine conditions for the three subjects. Only one day of data was collected for each condition, and only one condition was run per day. Subjects performed a mean of 1,244 trials per condition (mean of 1,271 for Subject B; 1,106 for Subject H; 1,354 for Subject C). Sessions typically lasted 2–3 hr, and ended when either the subject had received the maximum amount of liquid reward they would work for, or when they had been in the lab for the maximum scheduled time.

Training

All three subjects were previously trained to perform simple decision-making tasks involving making saccades to targets for

rewards. During training (as well as during testing), subjects typically did a minimum of 1,000 trials in a session, and one session per day. Subjects were first trained on the strongest negatively and strongest positively correlated conditions, performing only one of these conditions per session. This was done to give subjects familiarity with the task, and to ensure that the subjects were able to successfully employ the optimal strategies in both of these cases—switching to the other after a reward in the negatively correlated condition, and staying with the same target after a reward in the positively correlated condition. We required the subject to reach two training benchmarks before we began data collection. First, a subject was required to reach 75% optimal performance on both training conditions (that is, 75% of choices were stays in the positively correlated case, and 75% of choices were switches in the negatively correlated case) during sessions of at least 1,000 trials each. This typically took a week of training. We then tested the subject's ability to easily change their strategy between the two conditions. To do this, we had animals perform 500 trials of the positive correlation condition and 500 trials of the negative correlation condition back-to-back on the same day, requiring that they reached the same 75% optimal performance threshold within these 500 trial blocks. All animals made at least 75% optimal choices on each of the two conditions within these 500 trials. After this final day of training and testing, we would begin data collection on the nine conditions. None of the data analyzed came from the training period.

Results

The key question in our analysis is whether monkeys are more predisposed to pursue the optimal course of action in clumpy (i.e., positively correlated) environments than in dispersed (i.e., negatively correlated) ones. By using matched levels of clumpiness and dispersion, we could compare environments in which temporal aggregation was systematically varied but all other variables were matched.

The likelihood of monkeys switching following a win for each of the different correlation conditions is plotted in Figure 2. Let us first consider monkeys' likelihood of switching following wins in the most weakly clumped environment (autocorrelation of $+1$; see Method for a description of correlation levels). The optimal strategy in this environment (as well as in any of the four clumped

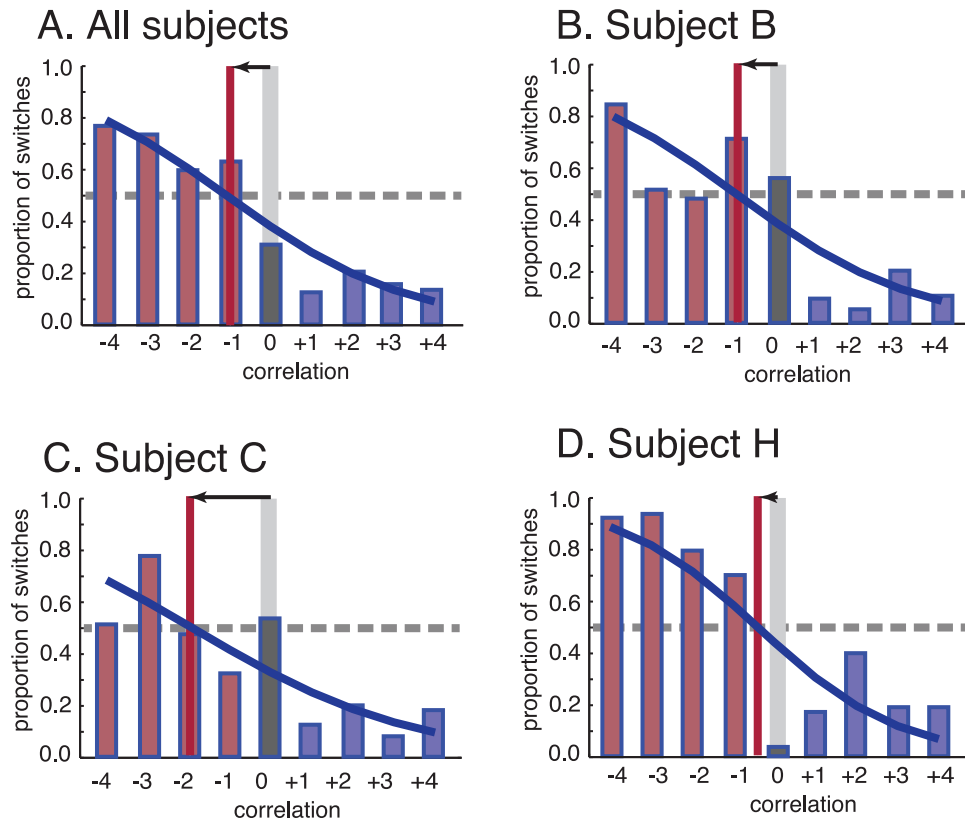


Figure 2. Animals show a bias toward staying following a win. The proportion of switches following a win in each correlation condition. Sigmoidal curve is best fit from a logistic regression. Horizontal dashed line indicates point of indifference. The red vertical line indicates crossover point, the value where the subject is indifferent between staying or switching following a win (see Results for further information). The arrow pointing to the red vertical line indicates the offset of the indifference point from 0. Plots shown for A all subjects combined, and B–D each subject individually. The color version of this figure appears in the online article only.

environments) is to never switch following a win (Nowak & Sigmund, 1993). We found that Monkey B nonetheless switched 10.03% (Binomial 95% CI [8.7,11.5]) of the time, a value that is quite close to the optimal value of zero. We found similar low switching likelihoods in Monkeys C (13.12% of trials; 95% CI [10.6, 15.8]) and H (17.90% of trials; 95% CI [15.1, 20.9]).

In the weakly dispersed environment (autocorrelation of -1), the optimal strategy is to switch on 100% of trials (see Bicca-Marques, 2005; Scheibehenne et al., 2011). If behavior in the clumpy and dispersed environments is symmetric, then we should expect switching coefficients of 80% to 90%. However, Subject B showed a switching likelihood of 71.64%, much closer to random guessing (i.e., 50%) than 10.03%. To test whether this difference is statistically significant, we performed a two-proportion z test on the likelihood of choosing the optimal strategy in the two cases; that is, we tested the raw likelihoods of switching (for the clumpy environment) and of staying (for the dispersed environment). Thus, for Subject B, we tested an 89.97% chance of choosing the optimal strategy in the clumpy environment against a 71.64% chance of choosing the optimal strategy in the dispersed environment, giving a difference of 18.33%. We found that this difference is significant (one-tailed two-proportion z test, $z = 12.56$, $p < .0001$, 95% CI on the difference [15.56, 21.10]). This pattern

suggests that this monkey had more difficulty adapting its choice strategy and exploiting the pattern when the environment was dispersed than when it was clumped. Similarly, in the dispersed environment, Subject H switched 70.39% of the time, 11.71% less than their proportion of optimal choices in the clumpy environment ($z = 7.99$, $p < .0001$, 95% CI [6.72, 16.71]). Finally, Subject C switched at only 32.89% of trials, a highly suboptimal strategy, suggesting a particular difficulty in learning to exploit the pattern of this weakly dispersed environment. This was 53.99% less than their optimal behavior in the clumped environment, and this difference was large and significant ($z = 4.37$, $p < .0001$, 95% CI [49.91, 58.06]).

To assess performance for our full dataset, we computed a standard psychometric curve with a sigmoidal shape on the likelihood of switching for all subjects as a function of correlation (logistic regression; Figure 2). We defined the indifference point as the best fit crossing of the sigmoid. This number gives the value where the subject is indifferent between staying or switching following a win. An animal that treated positively and negatively correlated environments symmetrically to have a crossover point at a correlation of 0—that is, they should only be indifferent between staying and switching when there is no correlation between trials. A crossover value below 0.5 would indicate the animal is biased

toward staying, as it would mean the subject prefers to stay for some negative (below 0.5) correlations. We found that for Subject B, the crossover was 0.403 (95% CI [.393, .413], different than .5 with $p < .0001$, logistic regression; Figure 2B). That is, we would expect that we would need to set the correlation to .403 to make Subject B indifferent between staying and switching. For Subject C, the crossover was 0.314 (95% CI [.302, .325], different from .5 with $p < .0001$; Figure 2C); for Subject H, the crossover was 0.462 (95% CI [.454, .471], different than .5 with $p < .0001$; Figure 2D). Taking the data of the three subjects together the crossover point was 0.394 (95% CI [.388, .399], different from .5 with $p < .0001$; Figure 2A). Together these data demonstrate that these three monkeys do not treat positively and negatively correlated environments in the same way. Because all of the values are below .5, subjects are treating the environment's correlation as if it is more positively correlated than it actually is.

It is possible that, despite the training program, monkeys only slowly learned the statistics of the environment. If this were the case, one possible interpretation of our results is that they do not reflect a stable behavioral bias but instead reflect a learning bias. If this were the case, we would expect the pattern of results we report above to change when we focus our analyses on trials late within a condition, when the animal had adequate time to adjust to the environment of the condition. We thus repeated the same analyses as above on just the last 500 trials of each session (roughly the last third). We found the same patterns again. Specifically, for Subject B, the intercept was 0.387 (95% CI [.369, .405], different from .5 with $p < .0001$). For Subject C, the intercept was 0.368 (95% CI [.355, .381], different from .5 with $p < .0001$). For Subject H, the intercept was 0.454 (95% CI [.443, .466], different from .5 with $p < .0001$). This suggests that our results reflect stable behavioral biases that are hard to overcome.

Having established that our observed affects are not due to learning within a session, one possible confound remains: Subjects may have not fully understood the structure of the task but gradually learned it over time, across sessions. If this were the case, we would expect a significant trend over time for the subjects—specifically, we would expect that their behavior would gradually approach the optimal strategy (always switch with negative correlations; never switch with positive correlations) over days of training. However, we did not see this. We examined the correlation between session number and likelihood of switching in each session on a session by session basis. For Subject B, the correlation was $-.24$ ($p = .504$, not significant, 95% CI $[-.65, .17]$). For Subject C, the correlation was $.36$ ($p = .168$, not significant, 95% CI $[-.11, .73]$). For Subject H, the correlation between likelihood of switching and day of recording was $-.33$ ($p = .384$, not significant, 95% CI $[-.70, .12]$).

Discussion

We have shown that, in a very simple gambling task that manipulates the degree of autocorrelation between trials, monkeys show more optimal performance in clumped than in dispersed resource environments. One of the major factors that drove the evolution of our minds and brains was the need to forage (e.g., Hills, 2006; Passingham & Wise, 2012) and to correctly assess environmental contingencies (e.g., Alloy & Tabachnik, 1984; Genovesio, Wise, & Passingham, in press). Foraging environments

are often characterized by clumpiness—the aggregation of resources in space and time (e.g., Wilke, Hutchinson, Todd, & Czienskowski, 2009). We hypothesized that this clumpiness has influenced decision-making strategies in evolutionary terms, and has led to a propensity to see positive recency in independent events (i.e., the hot-hand bias).

Across multiple domains, animals show biases in their ability to learn outcome-based strategies known as win–stay/lose–shift (hot-hand bias) and win–shift/lose–stay (Barraclough et al., 2004; Lau & Glimcher, 2005; MacDonald, Pang, & Gibeault, 1994; Randall & Zentall, 1997). It appears that win–stay strategies appear to be more natural for monkeys than win–shift ones (Barraclough et al., 2004; Hayden et al., 2008; Hayden et al., 2011; Heilbronner & Hayden, 2013). In the present study, we sought to reduce the potential impact of these biases by including an extensive training period and only examining data collected after the conclusion of training. Our finding that preferences do not change during the later periods of the task is consistent with the idea that learning has concluded and that learning rate differences no longer matter. These previous research studies have also not generally penalized animals for exhibiting such hot-hand biases, meaning they may have adopted such a strategy as a convenient heuristic and simply had no reason to alter their behavior. Our results here show that, even with hundreds of trials, animals will continue to behave suboptimally in negatively correlated environments, suggesting that the hot-hand bias is strong and difficult to overcome.

Previous studies of animal behavior have found what is known as spontaneous alternation behavior: the tendency for animals to switch back and forth between two equally valued options. Although this has primarily been observed in rodents (Dember & Richman, 1989; Hughes, 2004), a tendency to alternate has also been found in monkeys (Lau & Glimcher, 2005). Although this bias may also have an ecological function (facilitating exploration in a frequently changing environment), this tendency has a weaker effect on behavior than their reinforcement history (Lau & Glimcher, 2005), a result consistent with our own. Thus, although animals may tend to alternate more than would be predicted by their reinforcement history, this tendency seems to be secondary to a tendency to choose recently rewarded options, and thus secondary to the hot-hand bias we report here.

One potential explanation for the hot-hand bias could be that animals have limited memory spans and thus overextrapolate from recent outcomes. Such a theory would predict that animals would show a bias for seeing negatively correlated sequences as being more anticorrelated than they are—the opposite from what we observe here. Instead, our results are consistent with the hypothesis that monkeys see more positive correlation than there actually is, and may possess a strong propensity to assume positive recency as a default expectation which is difficult to learn out of—as has been shown in humans (see Wilke & Barrett, 2009). Careful human studies indicate that this bias is robust, is not due to Western culture, and is strongest in foraging contexts, but reduced in contexts where people have learned that outcomes are independent (cf. Wilke & Barrett, 2009).

One interpretation of the animals' behavior is that they perceived a higher causal correlation between their chosen option and the reward outcome than there truly was. That is, even when there was no contingency between their chosen action and the outcome, they still perceived a positive contingency between their choice

and the reward, leading them to be more likely to “stay.” One factor that increases perception of contingency is frequent outcomes (sometimes referred to as the outcome-density effect; Blanco, Matute, & Vadillo, 2013; Vallée-Tourangeau, Murphy, & Baker, 2005); another is frequent causal candidates (the cue-density effect; Blanco et al., 2013; Vadillo et al., 2011). People also overattribute contingency when the causal candidate is the subject’s own actions (the illusion of control; Langer, 1975; Thompson, 1999), possibly because self-involvement usually increases the frequency of the causal candidate, causing a cue-density effect (Yarritu, Matute, & Vadillo, 2014). These explanations, among others, are all consistent with our data here—the causal candidate was the monkey’s own action, and in the no-contingency condition the causal candidate (choice) and outcome (reward) occurred on approximately 50% of trials, levels at which one would expect to find outcome-density and cue-density effects. Indeed, these biases may serve as the proximate mechanisms by which the hot-hand bias is realized. Future research could focus on elucidating the connection between these biases.

We acknowledge two limitations of the present study. First, our results were from animals that are acting, not simply observing. Although the hot-hand bias has been found in both actors and observers (Gilovich, Vallone, & Tversky, 1985), there is also a tendency to overattribute contingency to one’s own actions (Langer, 1975). Our study specifically looked at actors—the monkeys were making judgments of contingency based on their own actions—and although we suspect they would apply to monkeys making contingency judgments based on observed events, our data are silent on that issue. Second, it is possible that repeating an eye movement is easier than alternating between eye movements. We find this unlikely to play a large role in our results, as the oculomotor system is carefully engineered so that all eye movements are almost effortless to perform, and both humans and monkeys perform saccadic eye movements 3–4 times per second constantly during waking hours.

Future research should further explore these findings in nonhuman primates, but we believe that our results emphasize the importance of the often-neglected notion of ecological (and evolutionary) rationality in the evolution of the primate mind.

References

- Alloy, L. B., & Tabachnik, N. (1984). Assessment of covariation by humans and animals: The joint influence of prior expectations and current situational information. *Psychological Review*, *91*, 112–149. doi:10.1037/0033-295X.91.1.112
- Avugos, S., Köppen, J., Czienskowski, U., Raab, M., & Bar-Eli, M. (2013). The “hot hand” reconsidered: A meta-analytic approach. *Psychology of Sport and Exercise*, *14*, 21–27. doi:10.1016/j.psychsport.2012.07.005
- Ayton, P., & Fischer, I. (2004). The hot hand fallacy and the gambler’s fallacy: Two faces of subjective randomness? *Memory & Cognition*, *32*, 1369–1378. doi:10.3758/BF03206327
- Bar-Eli, M., Avugos, S., & Raab, M. (2006). Twenty years of “hot hand” research: Review and critique. *Psychology of Sport and Exercise*, *7*, 525–553. doi:10.1016/j.psychsport.2006.03.001
- Barracough, D. J., Conroy, M. L., & Lee, D. (2004). Prefrontal cortex and decision making in a mixed-strategy game. *Nature Neuroscience*, *7*, 404–410. doi:10.1038/nn1209
- Bell, W. J. (1991). *Searching behaviour: The behavioural ecology of finding resources*. New York, NY: Chapman and Hall.
- Bicca-Marques, J. C. (2005). The win-stay rule in foraging decisions by free-ranging titi monkeys (*Callicebus cupreus cupreus*) and tamarins (*Saguinus imperator imperator* and *Saguinus fuscicollis weddelli*). *Journal of Comparative Psychology*, *119*, 343–351. doi:10.1037/0735-7036.119.3.343
- Blanchard, T. C., Pearson, J. M., & Hayden, B. Y. (2013). Postreward delays and systematic biases in measures of animal temporal discounting. *Proceedings of the National Academy of Sciences of the United States of America*, *110*, 15491–15496. doi:10.1073/pnas.1310446110
- Blanco, F., Matute, H., & Vadillo, M. A. (2013). Interactive effects of the probability of the cue and the probability of the outcome on the overestimation of null contingency. *Learning & Behavior*, *41*, 333–340. doi:10.3758/s13420-013-0108-8
- Brainard, D. H. (1997). The psychophysics toolbox. *Spatial Vision*, *10*, 433–436. doi:10.1163/156856897X00357
- Castel, A. D., Drolet Rossi, A., & McGillivray, S. (2012). Beliefs about the “hot hand” in basketball across adult life span. *Psychology and Aging*, *27*, 601–605. doi:10.1037/a0026991
- Cornelissen, F. W., Peters, E., & Palmer, J. (2002). The eyelink toolbox: Eye tracking with MATLAB and the Psychophysics Toolbox. *Behavior Research Methods, Instruments & Computers*, *34*, 613–617. doi:10.3758/BF03195489
- Croson, R., & Sundali, J. (2005). The gambler’s fallacy and the hot hand: Empirical data from Casinos. *Journal of Risk and Uncertainty*, *30*, 195–209. doi:10.1007/s11166-005-1153-2
- Dember, W., & Richman, C. (Eds.). (1989). *Spontaneous alternation behavior* (Vol. viii). New York, NY: Springer-Verlag Publishing. doi:10.1007/978-1-4613-8879-1
- Falk, R., & Konold, C. (1997). Making sense of randomness: Implicit encoding as a basis for judgment. *Psychological Review*, *104*, 301–318. doi:10.1037/0033-295X.104.2.301
- Fawcett, T. W., Fallenstein, B., Higginson, A. D., Houston, A. I., Mallpress, D. E. W., Trimmer, P. C., & McNamara, J. M. (2014). The evolution of decision rules in complex environments. *Trends in Cognitive Sciences*, *18*, 153–161. doi:10.1016/j.tics.2013.12.012
- Genovesio, A., Wise, S. P., Passingham, R. E. (in press). Prefrontal-parietal function: From foraging to foresight. *Trends in Cognitive Sciences*.
- Gilovich, T., Vallone, R., & Tversky, A. (1985). The hot hand in basketball: On the misperception of random sequences. *Cognitive Psychology*, *17*, 295–314. doi:10.1016/0010-0285(85)90010-6
- Haselton, M. G., Bryant, G. A., Wilke, A., Frederick, D. A., Galperin, A., Frankenhuis, W. E., & Moore, T. (2009). Adaptive rationality: An evolutionary perspective on cognitive bias. *Social Cognition*, *27*, 733–763. doi:10.1521/soco.2009.27.5.733
- Hayden, B. Y., Heilbronner, S. R., Nair, A. C., & Platt, M. L. (2008). Cognitive influences on risk-seeking by rhesus macaques. *Judgment and Decision Making*, *3*, 389–395.
- Hayden, B. Y., Heilbronner, S. R., Pearson, J. M., & Platt, M. L. (2011). Surprise signals in anterior cingulate cortex: Neuronal encoding of unsigned reward prediction errors driving adjustment in behavior. *The Journal of Neuroscience*, *31*, 4178–4187. doi:10.1523/JNEUROSCI.4652-10.2011
- Hayden, B. Y., Nair, A. C., McCoy, A. N., & Platt, M. L. (2008). Posterior cingulate cortex mediates outcome-contingent allocation of behavior. *Neuron*, *60*, 19–25. doi:10.1016/j.neuron.2008.09.012
- Heilbronner, S. R., & Hayden, B. Y. (2013). Contextual factors explain risk-seeking preferences in rhesus monkeys. *Frontiers in Neuroscience*, *7*, 7. doi:10.3389/fnins.2013.00007
- Hills, T. T. (2006). Animal foraging and the evolution of goal-directed cognition. *Cognitive Science*, *30*, 3–41. doi:10.1207/s15516709cog0000_50
- Houston, A. I., McNamara, J. M., & Steer, M. D. (2007). Do we expect natural selection to produce rational behaviour? *Philosophical Transactions*

- tions of the Royal Society B, 362, 1531–1543. doi:10.1098/rstb.2007.2051
- Hughes, R. N. (2004). The value of spontaneous alternation behavior (SAB) as a test of retention in pharmacological investigations of memory. *Neuroscience & Biobehavioral Reviews*, 28, 497–505. doi:10.1016/j.neubiorev.2004.06.006
- Hutchinson, J. M. C., & Gigerenzer, G. (2005). Simple heuristics and rules of thumb: Where psychologists and behavioural biologists might meet. *Behavioural Processes*, 69, 97–124. doi:10.1016/j.beproc.2005.02.019
- Hutchinson, J. M. C., Wilke, A., & Todd, P. M. (2008). Patch leaving in humans: Can a generalist adapt its rules to dispersal of items across patches? *Animal Behaviour*, 75, 1331–1349. doi:10.1016/j.anbehav.2007.09.006
- Krause, J., & Ruxton, G. D. (2002). *Living in groups: Oxford series in ecology and evolution*. Oxford, UK: Oxford University Press.
- Lakshminaryanan, V., Chen, M. K., & Santos, L. S. (2008). Endowment effect in capuchin monkeys. *Philosophical Transactions of the Royal Society B*, 363, 3837–3844. doi:10.1098/rstb.2008.0149
- Langer, E. J. (1975). The illusion of control. *Journal of Personality and Social Psychology*, 32, 311–328. doi:10.1037/0022-3514.32.2.311
- Lau, B., & Glimcher, P. W. (2005). Dynamic response-by-response models of matching behavior in rhesus monkeys. *Journal of the Experimental Analysis of Behavior*, 84, 555–579. doi:10.1901/jeab.2005.110-04
- MacDonald, S. E., Pang, J. C., & Gibeault, S. (1994). Marmoset (*Callithrix jacchus jacchus*) spatial memory in a foraging task: Win-stay versus win-shift strategies. *Journal of Comparative Psychology*, 108, 328–334. doi:10.1037/0735-7036.108.4.328
- Nickerson, R. S. (2002). The production and perception of randomness. *Psychological Review*, 109, 330–357. doi:10.1037/0033-295X.109.2.330
- Nowak, M., & Sigmund, K. (1993). A strategy of win-stay, lose-shift that outperforms tit-for-tat in the Prisoner's Dilemma game. *Nature*, 364, 56–58. doi:10.1038/364056a0
- Oskarsson, A. T., van Boven, L., McClelland, G. H., & Hastie, R. (2009). What's next? Judging sequences of binary events? *Psychological Bulletin*, 135, 262–285. doi:10.1037/a0014821
- Passingham, R., & Wise, S. (2012). *The neurobiology of the prefrontal cortex: Anatomy, evolution, and the origin of insight*. Oxford, UK: Oxford University Press.
- Raab, M., Gula, B., & Gigerenzer, G. (2012). The hot hand exists in volleyball and is used for allocation decisions. *Journal of Experimental Psychology: Applied*, 18, 81–94. doi:10.1037/a0025951
- Randall, C. K., & Zentall, T. R. (1997). Win-stay/lose-shift and win-shift/lose-stay learning by pigeons in the absence of overt response mediation. *Behavioural Processes*, 41, 227–236. doi:10.1016/S0376-6357(97)00048-X
- Reifman, A. (2011). *Hot hand: The statistics behind sports' greatest streaks*. Dulles, VA: Potomac Books.
- Scheibehenne, B., Wilke, A., & Todd, P. M. (2011). Expectations of clumpy resources influence predictions of sequential events. *Evolution and Human Behavior*, 32, 326–333. doi:10.1016/j.evolhumbehav.2010.11.003
- Stevens, J. R. (2010). Rational decision making in primates: The bounded and the ecological. In M. L. Platt & A. A. Ghazanfar (Eds.), *Primate neuroethology* (pp. 96–116). Oxford, UK: Oxford University Press. doi:10.1093/acprof:oso/9780195326598.003.0006
- Taylor, L. R. (1961). Aggregation, variance and the mean. *Nature*, 189, 732–735. doi:10.1038/189732a0
- Taylor, L. R., Woiwod, I. P., & Perry, J. N. (1978). The density-dependence of spatial behaviour and the rarity of randomness. *Journal of Animal Ecology*, 47, 383–406. doi:10.2307/3790
- Thompson, S. C. (1999). Illusions of control: How we overestimate our personal influence. *Current Directions in Psychological Science*, 8, 187–190. doi:10.1111/1467-8721.00044
- Todd, P. M., & Gigerenzer, G. (2012). What is ecological rationality? In P. M. Todd, G. Gigerenzer, & the ABC Research Group. (Eds.), *Ecological rationality: Intelligence in the world* (pp. 3–32). New York, NY: Oxford University Press. doi:10.1093/acprof:oso/9780195315448.003.0011
- Tooby, J., & DeVore, I. (1987). The reconstruction of hominid behavioral evolution through strategic modeling. In W. G. Kinsey (Ed.), *The evolution of primate behavior: Primate models* (pp. 183–237). New York, NY: New York SUNY University Press.
- Vadillo, M. A., Musca, S. C., Blanco, F., & Matute, H. (2011). Contrasting cue-density effects in causal and prediction judgments. *Psychonomic Bulletin & Review*, 18, 110–115. doi:10.3758/s13423-010-0032-2
- Vallée-Tourangeau, F., Murphy, R. A., & Baker, A. G. (2005). Contiguity and the outcome density bias in action–outcome contingency judgments. *The Quarterly Journal of Experimental Psychology Section B*, 58, 177–192.
- Wilke, A., & Barrett, H. C. (2009). The hot hand phenomenon as a cognitive adaptation to clumped resources. *Evolution and Human Behavior*, 30, 161–169. doi:10.1016/j.evolhumbehav.2008.11.004
- Wilke, A., Hutchinson, J. M. C., Todd, P. M., & Czienskowski, U. (2009). Fishing for the right words: Decision rules for human foraging behavior in internal search tasks. *Cognitive Science*, 33, 497–529. doi:10.1111/j.1551-6709.2009.01020.x
- Wilke, A., & Mata, R. (2012). Cognitive bias. In V. S. Ramachandran (Ed.), *Encyclopedia of human behavior* (2nd edition). Maryland Heights, MD: Elsevier.
- Wilke, A., Scheibehenne, B., Gaissmaier, W., McCanney, P., & Barrett, H. C. (in press). Illusionary pattern detection in habitual gamblers. *Evolution and Human Behavior*.
- Wilke, A., & Todd, P. M. (2012). The evolved foundations of decision-making. In M. K. Dhami, A. Schlottmann, & M. Waldmann (Eds.), *Judgment and decision making as a skill: Learning, development and evolution* (pp. 3–27). Cambridge, UK: Cambridge University Press.
- Yarritu, I., Matute, H., & Vadillo, M. A. (2014). Illusion of control: The role of personal involvement. *Experimental Psychology*, 61, 38–47. doi:10.1027/1618-3169/a000225

Received January 15, 2014

Revision received April 30, 2014

Accepted May 1, 2014 ■