Variable and sub-optimal responses to a choice problem are a persistent default mode

Amelia R. Hunt¹, Warren R.G. James¹, Josephine Reuther¹, Melissa Spilioti², Eleanor K. Mackay¹,
and Alasdair D.F. Clarke³

1. School of Psychology, University of Aberdeen
2. Division of Surgery and Interventional Science, University College London
3. School of Psychology, University of Essex

Abstract

Here we report persistent choice variability in the presence of a simple decision rule. Two analogous choice problems are presented, both of which involve making decisions about how to prioritize goals. In one version, participants choose a place to stand to throw a beanbag into one of two hoops. In the other, they must choose a place to fixate to detect a target that could appear in one of two boxes. In both cases, participants do not know which of the locations will be the target when they make their choice. The optimal solution to both problems follows the same, simple logic: when targets are close together, standing at/fixating the midpoint is the best choice. When the targets are far apart, accuracy from the midpoint falls, and standing/fixating close to one potential target achieves better accuracy. People do not follow, or even approach, this optimal strategy, despite substantial potential benefits for performance. Two interventions were introduced to try and shift participants from sub-optimal, variable responses to following a fixed, rational rule. First, we put participants into circumstances in which the solution was obvious. After participants correctly solved the problem there, we immediately presented the slightly-less-obvious context. Second, we guided participants to make choices that followed an optimal strategy, and then removed the guidance and let them freely choose. Following both of these interventions, participants immediately returned to a variable, sub-optimal pattern of responding. The results show that while constructing and implementing rational decision rules is possible, making variable responses to choice problems is a strong and persistent default mode. Borrowing concepts from classic animal learning studies, we suggest this default may persist because choice variability can provide opportunities for reinforcement learning.

Keywords: problem-solving, decision, eye movements, variability, insight, training

We are often faced with decisions about how to allocate limited resources to achieve more than one goal. A rational strategy to apply to these dilemmas is to compare the resources required to achieve our goals to the resources we have available. If the combined requirements of two goals exceeds the available resources, we should focus on accomplishing just one of them. If we have enough resources, though, we can attempt to accomplish more than one goal. There are many examples of these kinds of dilemmas in daily life, from deciding which home improvement projects to embark on with a limited budget, to deciding to switch off the car radio when navigating through an unfamiliar neighborhood. Under many circumstances, the decision can be complicated by uncertainty about the task requirements, failures to appreciate the limitations in one's own resources, or the perceived importance or relative urgency of the goals themselves. However, Clarke and Hunt (2016) demonstrated that even when these decisions dilemmas are highly simplified and people have all the information they need, they fail to approach an optimal solution. We propose an explanation for this failure in the current study.

Clarke and Hunt (2016) posed the same decision dilemma in a range of different task contexts. In all contexts, participants first performed a single task across a range of difficulty levels, establishing the level of skill of each individual. Participants were then given two targets and a decision to make - to try and accomplish both, or to focus on accomplishing one. For example, in the "throwing" version of the task, participants first threw beanbags into hoops of varying distances away, allowing both the participants and the experimenters to observe their throwing skill. In the second part of the experiment, there were two hoops, and participants were not told which one was their target until they chose a place to stand. The optimal choice of standing position to maximize potential accuracy depends on the distance between the hoops. If the hoops are close together, it makes sense to stand between them, because the beanbag has a good chance of landing in the hoop irrespective of which is designated as the target. But if the hoops are far enough apart that the expected odds of success from a central position is less than 50%, standing next to one of the two hoop makes more sense: in the long term, they will be 100% accurate if the chosen hoop is the target, and 0% accurate at the other hoop, averaging out to 50%. Although this strategy is simple to understand and implement, the results of Clarke and Hunt (2016) demonstrate that in this and two other tasks (visual detection and memorizing strings of digits), naive participants not only fail to maximize their potential accuracy on the task, most of them do not modify their behaviour in a way that varies systematically with manipulations of task difficulty at all (see also Morvan and Maloney, 2012). In a follow-up study, James, Clarke and Hunt (2017) ruled out an explanation based on missing or inaccurate information about throwing ability. They asked half of the participants, just before they executed each throw on the 45 trials in the choice phase of the experiment, to report their expected odds of success to hit each of the two targets from their chosen standing position. They were able to do so quite well, with a

correlation of $r=.89$ between actual and estimated probability of success. Despite this accurate self-awareness of ability, James and colleagues also replicated the unsystematic and sub-optimal decisions about standing position found in the original study. Moreover, the standing position choices of the group who were asked to estimate their odds of success was indistinguishable from the control group, who were not asked this question. This result demonstrates that participants have all the information they need to make better choices, and even calling attention to this information at the time of the choice does not provoke choices that would improve performance.

Why do people fail to maximize success despite having a simple solution and all the necessary information to do so? Various cognitive biases could play a role, such as loss aversion (Kahnemann & Tversky, 1984) or an achievement motive (McClelland, 1953). These classic studies show that people's perceptions of concepts that influence choice, such as value, risk, effort, gain and loss, are not always aligned with external benchmarks, and are not always evaluated in a way that maximises what could objectively be defined as success. The achievement motive, for example, has been demonstrated in a clever experiment in which participants choose a place to stand to toss a ring onto a post (Atkinson and Litwin, 1960). Although people could choose to stand very close and always hit the target, most people instead choose a location from which they can expect to be successful on about half of the throws. This choice of standing position has been suggested to reflect a tradeoff between the frequency of successful tosses and a sense of satisfaction that comes from knowing those successes were not a *fait accompli*. Could the achievement motive explain the sub-optimal decisions observed by Clarke and Hunt (2016)? It may explain *some* participants' choices on *some* trials, particularly in the beanbag throwing task which shares some features with the ring-toss task. An achievement motive might provoke participants to remain in the center when the hoops are close together, and when the hoops are separated far enough that they have <50% chance of hitting them from the center, they might drift towards one of the hoops, staying far enough away that they will be assured some satisfaction for getting it in. This pattern of behaviour would approximate an optimal solution, but not quite achieve it. The results as a whole do not follow this pattern, in both Clarke and Hunt (2016) and James et al., (2017). Looking at individual choices shown in the results of both of these papers, one can pick out one or two participants who are consistent with this pattern, but underlying even these minority of cases is a widely varying set of choices *within* each individual. It seems even more doubtful that achievement motivation could account for choices in the visual detection and memory versions of the task, as choices in these tasks show a similar range of variability within the same condition, and a need to achieve would have a far less straightforward effect on choices in these tasks. Variability is indeed the most striking feature of the results observed by Clarke and Hunt (2016) and James et al. (2017). Even given exactly the same conditions, most individual

participants displayed a wide range of different choices from one trial to the next. When there is a consistent rule to follow that would maximise accuracy, any variability in choices will lower overall success. This variability also makes any explanation for the results based on a single, simple cognitive bias unlikely. More likely, a wide range of different cognitive biases and strategies can influence choices on each trial, and some may dominate in some choices and not others.

Choice variation is consistent with a lack of *insight* into the problem's solution. When the solution is explained during debriefing, it has been our experience that participants understand it easily, and often express some consternation for not having figured it out for themselves. Taking this further, in the current study, we demonstrate that the optimal solution can be easily implemented when it is known: as a starting point, we report decision behaviour for non-naïve participants in both the throwing and detection versions of this task. Four of the participants are authors on this paper, and our experience of having completed these tasks was that, with the benefit of knowing what the optimal strategy is, implementing it is trivial. The only challenge came from estimating at exactly what difficulty level to switch between trying to accomplish both and focusing on one. Naïve participants, in contrast, do not show any consistent change in strategy with task difficulty even at the extremes.

There is a large literature on insight problem solving, but the question of how exactly to define insight is still not widely agreed. Most people can agree that there is a distinctive experiential difference between solving problems that require following complex cognitive routines (like a Rubik's cube) relative to the "aha" moment that comes with solving some specific types of problems. In the Gestalt school, this step-change was called "restructuring", and was suggested to reflect a re-organisation that allowed the same information to be represented in a different way. Consistent with this sudden change in information, Metcalfe and Wiebe (1987) asked participants to report "warmth" ratings, that is, how close they felt to arriving at a solution, every ten seconds until they solved the problem. Warmth ratings predicted when people would solve the control set of problems (algebra), but no such relationship was evident for the insight problem set. The authors propose that the inability to predict whether a solution is imminent could provide an operational definition of an insight problem. Weisberg (2006) later emphasized that for a problem to be defined as an insight problem, this discontinuity in the solution process needed to be caused by restructuring, and this restructuring had to be necessary for a solution to be achieved. Under this definition, the decision dilemmas we present in these experiments are not typical insight problems: Participants repeatedly make choices and see the consequences of each decision play out. They can complete each decision with or without insight about the optimal strategy, but using this strategy will lead to significant improvements in their accuracy in the long run.

A definition and approach to insight that is more germane to our paradigm may therefore come from classic animal learning studies¹. Maier (1940), who was also a Gestaltist, examined the extent to which animals have the capacity to spontaneously adjust how they respond to stimuli based on knowledge. He argued for three different categories for solving problems: variability, equivalence reactions, and reasoning. *Variability* in choice responses is presented as a functional tendency, assumed to support trial-and-error learning, and stands in contrast with repetitive, inflexible behaviour that is a hallmark of cortical damage (e.g. Gelb and Goldstein, 1920). Krechevsky (1937) argues that varying choices allows the organism to explore “means-end readiness”, or in other words, the set of available options and their consequences. These options can be refined into hypotheses, and then a solution, through a gradual process of reinforcement learning. In *Equivalence Reactions*, hypothesis learning can be transferred to other problems that share features with the original, though usually with some hesitancy. This process permits acceleration of the learning process when circumstances are in some way familiar, though if wrongly applied, it can interfere with learning. *Reasoning* differs from learning altogether, and refers to the ability to spontaneously integrate information from outside the experimental context to take responses in a new direction. More generally, reasoning refers to the ability to use abstract information to guide behaviour, a concept later referred to as *teleological control* by Dickinson (1985). Dickinson similarly drew a stark distinction between *habits*, which are consistent responses elicited by the immediate environment, and *actions*, which are controlled by the animal’s knowledge about consequences. Habits will always follow from the same conditions, but actions allow an organism to change the choices it makes given what the organism knows. This dichotomy has held up well in the literature, with accruing evidence that actions and habits are subserved by distinct neural systems (e.g. O’Doherty, Cockburn & Pauli, 2017).

In bringing this distinction between habits and actions back to simple human choices, we can use eye movements as a starting point. We move our eyes several times each second, and each of these movements can be thought of as a choice -- a resolution to a mini-problem of where the most useful and interesting visual information is coming from at the moment. Solving this problem optimally, although possible, requires complex computations even in a highly simplified and predictable environment (Najemnik and Geisler, 2008). In a complex and unpredictable environment, a more efficient approach to

¹ Although the 1930’s is typically thought of as having been dominated by behaviourism, there existed a minority of researchers who were investigating and developing theories of animal problem-solving. Dewsbury (2000) suggests this research has been forgotten because it was marginalized by the behaviorists at the time, and subsequently swept aside by the cognitive revolution in the 1950’s, which tended to somewhat indiscriminately characterize all the research of this era as behaviourist.

these kinds of “small” problems may be to solve them with variability (Krechevsky, 1937), allowing the constraints of the immediate environment to shape the set of viable choices and randomly varying within that set to allow for flexibility and learning to occur. Consistent with this notion, a stochastic model of fixation selection during visual search, which selects fixations at random from a population of common saccade vectors, describes human search behaviour reasonably well (Clarke, Green, Chantler & Hunt, 2016; Clarke, Stainer, Tatler & Hunt, 2017). A similar process of random selecting from a population of possible responses could guide many decisions, preventing stereotyped behaviour while avoiding overthinking of minor choices. Indeed, the choice experiment we replicate here (Clarke and Hunt, 2016) was originally designed as a method for assessing optimal eye movements (Morvan and Maloney, 2012). The similar degree of suboptimal variability in the context of choosing locations to fixate, choosing standing positions, and choosing material to memorize is striking. Stochastic selection from a set of plausible options may be a reasonable model for many simple choices under conditions of uncertainty, not just eye movements.

Variability in choice leads to arbitrary outcomes on each trial, but over time, choices associated with success would be reinforced and repeated. If there is a pattern to be discerned, then responses could eventually become more optimal, even in the absence of explicit knowledge of a decision rule. Variability itself has been previously argued to be an operant (Neuringer, 2002), providing the basis for learning and exploration. As already noted, the previous results from this decision paradigm are highly variable, suggestive of a trial-and-error approach. In contrast, the responses of the expert participants reported in the first experiment below are, we can safely assume, based on reason. Compared with naive participants, participants who were deliberately using the optimal strategy were, of course, much closer to optimal. However, it is important to note that these “reasoned” choices are also far less variable. These results demonstrate two important points. First, although some learning is required to understand the constraints, a separate process of reasoning allows a logical decision rule that optimises accuracy to be applied to a wide range of different circumstances. Second, applying this decision rule restricts variability, with the potential consequence of restricting future learning.

With these two extremes as points of comparison, our question in this study was initially about whether and how we can transition people from variability to reasoning (or in Dickinson’s terminology, from habit to action). One possibility is that equivalence reactions could provide a transition into spontaneous integration; that is, applying the correct solution in one circumstance could lead to its application in circumstances with similar properties. This generalization may prompt participants to restructure the successful choices as an abstract decision rule. On the other hand, applying a decision rule carries some

risks. If applied to the wrong circumstances, the rule would not only lead to sub-optimal outcomes, it would also restrict the variability of responses that is the basis for learning. Given these risks, taking a conservative approach to equivalence reactions and spontaneous integrations may be adaptive in the long run, leading to a high threshold of similarity needed for a hypothesis, or a strategy, to transfer across contexts. In this series of experiments, we explore this threshold. The first experiment simply establishes that the decision rule can be easily applied when it is known. The remaining experiments create circumstances in which participants carry out the optimal strategy in one context, and then provide a similar (Experiment 2) and the same (Experiment 3) context to observe whether the strategy persists.

Experiment 1: Optimal Choices in Non-Naive Participants

Here we present the results from a group of non-naive participants who completed replications of both the throwing and detection versions of the experiments reported in Clarke and Hunt (2016). In the throwing task, participants make a choice about where to stand to throw a beanbag into one of two possible hoops. They are only told which hoop is their target after they have made their choice. The detection task is similar in structure, except participants make a choice about where to look to detect a probe that can appear in one of two boxes (see also Morvan and Maloney, 2012). In both tasks a range of distances between the possible target locations (hoops/boxes) are presented, and the logic is the same: stand/look in a central location, equidistant between the two potential target locations when they are close together, and stand/look close to one potential target location when they are too far apart to throw to/see reliably from a central position. The results of this first experiment demonstrate that the optimal strategy can be easily implemented when it is explicitly known, as a basis for comparison with Experiments 2 and 3.

Methods

Experiment 1a: Throwing

Participants. Four participants completed the throwing task. Three are authors on this paper (AH, WJ, JR). The fourth participant was a fellow lab member, familiar with the paradigm. All participants had knowledge of the optimal strategy but had not completed the experiment previously.

Materials and Procedure. The procedure was similar to that used in Clarke and Hunt (2016, Experiment 2), with minor modifications. An outdoor, sheltered area of concrete slabs just outside the building (Figure 1) was used because the slabs (measuring 0.46 x 0.61m) were useful for marking and recording hoop placement and standing positions, respectively. In the first phase, flat hoops with a diameter of 0.40m were placed at a range of distances away (1.38m, 3.22m, 4.14m, 5.06m, 6.9m & 9.2m).

Participants threw 12 bean bags into a hoop at each distance in each of two directions and throwing performance (out of a maximum of 24) was used to determine the point at which each participant's accuracy dropped below 50%.

Figure 1: Illustration of the methods for the two tasks in Experiment 1. Throwing task: the goal is to get a bean-bag into one of two hoops of a specified colour (left side). Participants must first choose a place to stand, and are then told which one of the two hoops is their target. Detection task: the goal is to report whether the light grey dot is in the upper or lower part of the box (right side). The dot can appear in the left or right box. The participant must first fixate a box, and then the dot will appear. The separation between the boxes varies.

In the second phase, participants were again asked to throw bean-bags into hoops, but this time there were two potential targets for each throw and participants needed to choose where to stand before finding out which of the hoops was the actual target. In the original experiment (Clarke and Hunt, 2016), the distance between the hoops in the second phase was based around the participants' performance in the first phase. In the current version, six distances were randomly selected from a range of distances based on the performance of participants in previous experiments (0.46m to 11.5m). The random selection of distances within this range ensured that participants would be unable to rely on their knowledge of the structure of previous experiments in deciding whether they should stand in the middle or next to one hoop. Instead they had to base their decision on knowledge about their own ability. To begin the session, three pairs of hoops matched in color were placed at each of three separations (red was the closest separation, yellow the middle, blue the farthest). After 45 choice trials, the hoops were shifted to three new separations for a further 45 trials. The colour of the bean-bag that the participant drew at random from a bag before each throw determined which pair of hoops were potential targets on that trial (e.g., if a red beanbag was drawn, one of the two red hoops would be the target on that trial). Participants then took the beanbag and chose a place to stand. After they chose their standing position, they were told which hoop was the target, and they then attempted to get the beanbag into that target hoop. The bean bags would only be placed back into the bag once all nine had been thrown, to ensure that participants made an equal number of decisions for each distance. There were 15 trials for each distance (90 in total). Which of the two hoops

would be designated as the target on each trial was determined by a pre-generated random sequence. On each trial, the experimenter recorded the color of the beanbag, the standing position (based on numbers chalked on the wall over each row of paving slabs), and throwing accuracy (0 or 1).

Experiment 1b: Detection

Participants. Five participants (4 female, all right-handed) completed the detection task. Four are authors on this paper (AH, WJ, EM, JR) and the fifth participant was a fellow lab member, familiar with the paradigm. All participants were aware of the optimal strategy, had normal or corrected to normal vision, and like in the throwing task (Experiment 1a), were unaware of their personal switch-point. All participants provided informed consent.

Setup. The stimuli and task replicated those used by Clarke and Hunt (2016, Experiment 1) with a few minor changes. All stimuli were presented against a uniform grey background (50% white). The square boxes had an edge length of 1.0° visual angle and were lighter than the background (75% white). The target was a small, light grey dot (80% white) that appeared for 500ms either in the top (dot-up configuration) or the bottom (dot-down configuration) of one of the boxes. The “page-up” and “page-down” buttons of standard keyboard were used to collect responses, matching dot-up to the page-up key, and dot-down to the page-down key. If the participant blinked or if the eyes moved more than 1° while the target was on, the trial was terminated.

The experiment was executed in Matlab 7.9.0 (R2009b) on a PowerMac running Apple Macintosh Pro OS X (version 10.6.4) software using Psychtoolbox (Brainard, 1997) and EyeLinkToolbox functions (Cornelissen, Peters & Palmer, 2002). Stimuli were presented on a 24-inch Sony Trimaster EL OLED monitor with a resolution of 1080p and a refresh rate of 60 Hz. A chin rest, with forehead bar, assured a viewing distance of 54cm. The right eye was tracked using a desktop-mounted EyeLink 1000 (version 4.594) (SR Research Ltd, Mississauga, Ontario, Canada) to record eye position at 1000 Hz. A 9-point calibration sequence was used. The calibration and validation process was repeated every time the participant moved away from the chinrest. The process was also triggered if the participant broke fixation during the stimulus presentation five times in a row, or ten times cumulatively for each block of trials. The experiment took place in a dark, quiet room.

Experimental Design and Procedure. The experiment lasted for approximately 1 hour 30 minutes and consisted of two phases: a ‘Sensitivity Mapping Phase’, directly followed by a ‘Decision Phase’. Both phases are outlined below.

Sensitivity mapping phase. The sensitivity mapping phase consisted of four blocks. During each block, boxes were presented in 8 different distances: 2.7°, 3.9°, 5.2°, 6.8°, 8.4°, 10.1°, 11.4° or 12.5° between the centrally presented fixation mark and the centre of each box. Each distance was presented 12 times per block (96 trials per block, total 384 trials). The order of distances was randomized. All trials per distance were presented in succession. Each trial started with the presentation of a black fixation cross in the centre of the screen. Participants were asked to initiate the trial sequence by key press while fixating the cross intersection (drift check). After a stable fixation of 700 ms (within a 1° radius circle around the cross) two equidistant, grey boxes were presented, one on the left and one on the right of the fixation cross, with the target present in one of them. After 500 ms of stimulus presentation a blank (grey) screen appeared, indicating that a response was expected. A red screen was displayed and the trial terminated whenever participants looked more than 1° from fixation during the trial interval (i.e. starting from the key press at the start of the trial, ending with the onset of the response screen).

Decision phase. The Decision phase consisted of four blocks. During each block, boxes were presented in 9 different distances (measured from the centre of the screen to centre of the box). Two of the distances were constant among participants (8°, 18°). The participants' individual switch-point (E_0) was the basis for the remaining 7 distances (E_0 [-3°, -2°, -1°, $\pm 0^\circ$, +1°, +2°, +3°]). The switch-point was estimated from the sensitivity mapping data as described by Morvan and Maloney (2012). Each distance was presented 10 times per block (90 trials per block, total 360 trials). Trial order was randomized. Participants initiated each trial with a key press while fixating a black fixation cross to initiate a drift check. The cross was positioned 4° above the horizontal meridian with an offset of half the (to-be-presented) box distance either to the left or to the right. After 700 ms of stable fixation, three boxes were presented along the horizontal meridian. The middle box always remained centred horizontally, and the left and right boxes were equidistant from the central box on either side. The fixation cross was positioned equally often between either the centre and left box or the centre and right box. As we were not interested in the choice between the right and the left box, but between the centre and side-boxes, this equates the distance from the fixation to each of these two locations. Participants were instructed to make a saccade towards the box of their choice. Once fixation was stable inside one of the boxes for 50ms, the target configuration was presented for 500ms in either the left or the right box (never the centre). Figure 1B illustrates the target presentation of a typical trial of the decision phase.

Results

Choices. We first visually compared each participant's choice behaviour to an individualized estimate of their optimal strategy (that is, choices that would have achieved optimal performance). The optimal choices for each participant in each experiment are determined based on performance in the first phase of

the experiments (see supplementary material for performance curves). For the throwing experiment, the participant should choose to stand in the center for distances where accuracy from a central position is expected to be greater than 50%, and should stand near one hoop or the other for distances where accuracy from center is less than 50%. This is shown in Figure 2, with a dark blue line showing optimal performance and black circles showing actual standing position. It is clear from this figure that these expert participants made standing position choices that were close to optimal.

Figure 2. Results for each expert participant in the throwing experiment (1a) on the top row. Each dot is light grey and represents standing position on a single trial; the darker dots represent up to 15 overlaid trials. The dark blue line shows the optimal standing position for each participant. The standing position has been normalized to the distance from the center to the hoop. This means deviations from the line for small values of delta represent relatively small changes in standing position. The second row shows the results for each expert participant in the detection experiment (1b). Each dot represents the proportion of fixations made towards one of the side boxes at each value of delta. The dark blue line represents the optimal fixation strategy for each participant.

Similarly, for the detection experiment, the participant should choose to look at the center box when the separation between boxes is small enough that expected accuracy from the center is $>75\%$. For expected accuracy from the center of less than 75% , participants should instead choose the left or right box, because if the target appears here they will be 100% correct and if it appears in the other box they will be 50% correct, giving an expected overall accuracy of 75% . This is shown as the blue line in Figure 2.

Again, it is clear that expert participants' choices of where to fixate in this task (the black dots) were close to optimal.

Performance. For each participant in each experiment, we calculated the accuracy they would have achieved under an optimal strategy (i.e., had their choices followed the blue lines in Figures 2 and 3). The size of the difference between actual proportion correct and each participant's optimal performance ranges from underperforming by 0.089 to out-performing by 0.081. Overall, the mean difference from optimal accuracy was close to 0: -0.017 for the throwing task, and 0.036 for the detection task.

Figure 3. Comparison of actual proportion correct to optimal (an estimate of the proportion correct the participant would have achieved given optimal choices) for the throwing and detection experiments. Each line represents an expert participant. Performance is close to optimal.

Although participants' choices were broadly consistent with a near-optimal strategy, it is clear from the data that not all participants were accurate in switching at precisely the right point. These small inconsistencies had a trivial effect on accuracy, however, and all participants achieved a level of accuracy in the task that nearly equaled the accuracy expected under an optimal strategy. The results demonstrate that implementing the optimal strategy when it is known is straightforward. With this in mind, in the following experiments we attempt to nudge naive participants to implement an optimal strategy.

Experiment 2: From reaching to throwing

Clarke and Hunt (2016) demonstrated that participants fail to modify their behaviour in response to task difficulty across three tasks: the two tasks employed above (visual detection and throwing) as well as a third task, in which participants were given two strings of digits to memorise and did not know which one they would be asked to report. The logic of the digit string experiment was similar to the others: for

short digit strings that are easy to remember, it makes sense to memorize both, but when the two strings together are too long to remember accurately, participants should focus on just one. The participants in this experiment sometimes focused on just one string and sometimes looked at both, but their behaviour in this respect did not vary systematically with the length of the strings. In a fourth experiment of Clarke and Hunt (2016), the participants' task was simplified to a trivial level: participants were told they would need to pick up one of two beanbags, which were placed on a long table. There were three chairs - one in the middle and one at either end. The participant selected a chair to sit on. After they chose a chair, they were told which of the two beanbags they would need to pick up. Clearly, if the beanbags are within arm's reach from the central chair, the participant should sit there. If they are too far apart to reach from the center, the participant should select a chair close to one of the beanbags, ensuring they can at least reach one, if it is designated as the target. Clarke and Hunt (2016) found that all participants followed this pattern, choosing the center chair for central beanbags, and one of the side chairs when the beanbags were at either end of the table. Although participants perform the reaching version of the decision task optimally, they do not necessarily need to "solve" the task in this context. Participants have been asked to choose between two goals and they simply sit in the chair that is in front of one of them. The direct interaction between the choice (which chair to sit in) and the goal (pick up one of the beanbags from the table) means the "choice" is less about allocating resources, and more about finding the best way to execute a direct command. The choice of chair essentially accomplishes the task. In the throwing, detection and memorizing versions of the task, in contrast, there is an intervening action that is made easier or harder by the preceding choice. It is therefore not especially surprising that participants perform the table/reaching task optimally: they would have to not understand what the goals are to select a chair from which they can reach none of the beanbags they have been instructed to pick up. Indeed, Clarke and Hunt (2016)'s reason for running this experiment was to make sure participants did understand what it was they were supposed to be accomplishing.

Although the relationship between the choice and the goal is far more direct, the choice of chair in the table task is directly analogous to the standing position choices that would optimize throwing accuracy. In the next experiment, we primed one group of participants by asking them to perform the table task immediately before doing the throwing task. We expected that performing optimal choices in the table task would bring the decision between locations into a framework where at least some of the participants would recognize the strategy they should apply in more complex versions of this task.

Methods

Participants. Thirty-two students from the University of Aberdeen were recruited. Participants had no prior knowledge of the hypothesis of the experiment. They were randomly assigned to either the primed group (those who did the reaching task first and then the throwing task) or the control group (those who did only the throwing task), with 16 in each group. There was a £5 remuneration or appropriate academic credits awarded for participating. In this and all subsequent experiments, the experimental protocol was reviewed and approved by the Aberdeen Psychology Ethics Committee, and participants all gave signed, informed consent to participate.

Materials and Procedure. Two experimenters conducted this study. The first experimenter carried out the reaching task (detailed below) with half of the participants. The control group participants remained with the first experimenter for five to ten minutes to ensure that the second experimenter would not be able to tell if the participant had carried out the reaching task based on their time of arrival. The second experimenter carried out the throwing task with both groups, and was blind to their condition assignment.

Reaching Task. Six PVC beanbags of three colours (two red, two yellow and two blue) were placed on a long table equidistant from each other (Figure 5). At the centre of the table were the red beanbags, halfway over each side of the table were the yellow beanbags (one on each side), and at the far end were the blue beanbags (one on each end). Three chairs were positioned alongside the table: one on each end (opposite the blue beanbags) and one in the centre (opposite the red beanbags). In order to make sure participants were aware of their own reaching span, they were first asked to sit on the chair positioned at the centre of the table. With their back always touching the chair they were asked whether they could reach and touch the red, green and blue beanbags individually. They were then asked to stand up and move away from the table before starting the main part of the experiment. They were told which colour beanbag they would be asked to pick up, and that they had to choose a chair to sit on, at which point they would be told which of the two beanbags of that colour was the target. This was repeated for each colour. The order of the colour-location combination was the same for each participant, starting from the middle beanbag and working towards the sides. Once all three colours had been tested, they were taken to the other experimenter for the second part of the experiment.

Figure 4. The reaching task setup. The participant selects a chair from which to pick up one of two beanbags of a specified colour. After choosing a chair and sitting in it, they are told which of the two beanbags to pick up.

Throwing task. The experiment was an abbreviated version of the throwing task used in Experiment 1a. Six hoops with a diameter of 0.40m were taped down in a row. The red hoops were 4.60m apart, the yellow hoops were 8.28m apart, and the blue hoops were 11.96m apart. These hoop separations were selected on the basis of previous experiments; participants should stand in the center for the closest distance, and next to one of the hoops for the farthest distance, to achieve optimal performance. As in Experiment 1a, the colour of the beanbag the participant was handed on each trial specified which pair of hoops were potential targets on that trial. Participants took the beanbag and chose a place to stand. After choosing their standing position, they were told which hoop was the target (determined based on a pre-generated random list), and they then attempted to get the beanbag into that target hoop. The experimenter stood on the grass to the side and handed beanbags to the participant in a pre-specified but random order of colours, and kept a record of standing position and throwing success or failure. There were 15 trials for each distance (45 total).

Results

All participants in the reaching group correctly carried out the reaching task. That is, when they were asked to reach for one of the two red beanbags (the closest separation), they sat in the central chair. When the distance between the beanbags was larger, they moved to either the left or right side chairs, resulting in a 50% success rate. This replicated the behaviour observed by Clarke & Hunt (2016).

Standing Position Choices. The optimal solution in the throwing task would give standing position = 0 for the closest hoop distance, and standing position = 1 for the furthest (standing position has been normalized over hoop separations, with 1 being the distance from the center to the hoop). It is clear from Figure 5 that participants fail to adopt this strategy, replicating the striking findings of Clarke & Hunt (2016). We were interested in whether participants would be able to learn from their optimal performance in the reaching task, to adopt a more optimal strategy in the throwing task, but it is clear from the results from the primed group that participants fail to do so to a similar extent as the control group.

Figure 5. Standing positions chosen by each individual in Experiment 2. Each dot is a trial and each facet is a single participant. The top two rows are the control group, and the bottom two rows were primed by making optimal decisions the reaching/table task before completing the throwing task. Optimal decisions would stand close to 0 at the smallest value of delta and close to 1 at the largest. By this definition, no individual, in either group, chose optimal standing positions.

We can compare performance between groups by examining the change in mean normalized standing position for the closest versus the furthest hoop separation, which should be close to 1 if participants are optimal, and close to 0 if they are random. A comparison against 1 confirms that neither group could be described as optimal: the range for the control group is -0.18 to 0.21, and for the primed group it is -0.49 to 0.60. The difference in range between groups here captures the wider variation in standing position choices for the primed group, and suggests a comparison between groups on this measure should be approached with caution. Nonetheless, this comparison is consistent with a modest effect of priming, with the group who had the experience of making optimal decisions in the table task moderating their choices more with distance than the group that did not have this experience: a Welch two-sample t-test is just over the threshold for significance $t(18.75)=2.14$, $p=.046$). Because this result is right on the threshold, and the groups were unmatched in variance, we included two further experiments in the supplementary material that prime participants using the reaching task and measure the effect on the throwing task. Inconsistent with the current results, both of these experiments find no evidence that the optimal solution transfers from the table to the throwing task. This suggests that if priming with the table task does improve choice behaviour, the effect is very small. The more striking result from this experiment is that participants are highly variable in their choices, and very far from optimal, even after priming with the table task.

Throwing Performance. Accuracy for the three hoop separations in Experiment 2 are shown in Figure 6. Because we did not run a separate session measuring each person’s throwing accuracy over distance like we did in Experiment 1, we cannot simulate optimal accuracy for each individual. However, we know participants could achieve an average accuracy of 50% if they stood at one hoop in the far condition, and it is clear that both the primed and control groups fall short of this standard.

Figure 6: Standard boxplots depicting the accuracy at each hoop separation in Experiment 2. The primed group completed the reaching task before the throwing decision and the control group did not. There is no difference in accuracy between groups, and both under-perform relative to optimal, which should be at least 50% in the medium and far hoop conditions. The “optimal” group data is a simulation of 32 participants doing 15 trials each with a probability of success of .5.

Experiment 3: Guided optimal eye movements

Based on the results of Experiment 2 (and the two similar experiments reported in the supplementary material), we could not conclude that making optimal choices on the table task nudged participants to make better choices in the subsequent throwing task. One plausible explanation is the change in context. In everyday circumstances, one would not expect the same logic to apply to choosing a seat at a table as to choosing where to stand to make an accurate throw. The change in circumstances (inside to outside, reaching to throwing) may have been too different for people to relate the two choice situations, so they returned to solving the problem using variability.

In the next experiment, we therefore tested transfer of decision strategies without a change in the task. In the first session, we directed eye movement choices. In the second session, we let participants freely choose. Our question is whether repeatedly executing optimal eye movement behaviour in the first session would lead to improvements in performance in the same choice task once the guidance has been

removed. We employed the visual detection task used in Experiment 1b, but participants completed a training block where they were cued to fixate a particular box. The target would only appear after they had fixated that box. The cue always directed the participants' fixations to the location that would optimize their detection accuracy, given their own visual acuity as measured in the preceding sensitivity mapping phase. After completing this session, participants completed a session where they freely chose which box to fixate, just as in the original experiment. The key question is whether guidance in the form of the cues directing the eyes to the optimal location will result in more subsequent optimal decisions in the free-choice task relative to a control group who was given an equivalent amount of practice, but no guidance.

There is evidence that solutions to insight problems can be provoked by guiding the eyes to relevant locations (Grant and Spivey, 2003) or in a pattern that is consistent with the solution (Thomas and Lleras, 2008). This suggests executing task-appropriate responses might lead participants to restructure information in a way that is consistent with the responses they are executing. In our experiment, being required to repeatedly perform optimal eye movements could cause participants to persist in these eye movements when given the opportunity to freely choose. An interesting additional question we were interested in addressing here was whether guided optimal behaviour would lead not only to subsequent improvements in a free-choice setting, but would also be accompanied by *reason*, that is, whether the guided participants also would have more explicit insight about *why* they were making these choices. This experiment's methods and analysis plan were pre-registered on the Open Science Framework: <https://osf.io/yan5k/>.

Methods

Participants. Twenty-four participants (mean age 21, range 18-26, 19 female, 23 right-handed) were recruited from the University of Aberdeen community. All participants were unaware of the purpose of the experiment and had normal or corrected to-normal vision. Participants were offered either 5 credits (for undergraduate psychology students) or a £15 reimbursement for their participation. This study was reviewed and approved by the School of Psychology ethics committee. All participants provided informed consent.

Stimuli and Materials. The same stimuli and materials as in Experiment 1b were used.

Experimental Design and Procedure. The experiment took place over two sessions. Session 1 lasted approximately 1 hour 30 minutes, and Session 2 lasted approximately 1 hour. The experiment consisted of three phases; a *Sensitivity Mapping Phase*, a *Directed Eye Movement Phase* and a *Decision Phase*. The

Sensitivity Mapping Phase and the Decision Phase followed the same procedure as in Experiment 1b. The Directed Eye Movement Phase is described below. Participants were randomly assigned to one of two groups: the control group completed the Sensitivity Mapping Phase and the Decision Phase in Session 1, and the primed group completed the Sensitivity Mapping Phase and the Directed Eye Movement Phase in Session 1. In Session 2, both groups completed the Decision Phase. All participants were instructed to respond to the target as accurately as possible. No feedback was provided. To ensure that participants understood the task and that a psychometric curve could be fitted to the performance data successfully, participants needed to score 80% or above in the closest distance in the sensitivity mapping phase to take part in both sessions of the experiment and to be included in data analysis. Two participants were excluded on this basis and replaced (to retain a total of 12 in each group).

Directed Eye Movement Phase. This phase was similar to the decision phase, with the difference that participants were not given a choice over which of the three presented boxes they would fixate. Instead, the three boxes were presented alongside the fixation cross at the beginning of the trial, and after 700ms of stable fixation of the fixation cross, one of the boxes ‘blinked’ (disappeared for 400 ms and then reappeared) and participants were instructed to execute a saccade to this box and remain fixated there until the target appeared. This phase consisted of 4 blocks of 10 trials for each of the 9 distances used (90 trials per block, total 360 trials). The same distances as in the decision phase were used. Which of the boxes blinked was dependent on the box distance and the participants’ individual switch-point. For distances equal or smaller than the switch-point distance, the central box blinked. For all other distances, randomly either the left or the right side-box blinked.

Results

Figure 8 shows the fixation choices (center or side box) for the two groups in the second session, in which they all made free choices about which box to fixate (see Supplementary Material for a more detailed figure, showing each participant relative to their own switch point for both sessions). As a reminder, the optimal strategy would be to fixate the center box for small values of delta, and one of the side boxes when delta is large, as observed in the behaviour of the non-naïve participants in Experiment 1b. This strategy is not evidence in the results from either the primed or the control group. None of the 24 participants in this experiment could be described as optimal when allowed to freely choose which box to fixate, even in the primed group, who completed 360 trials in which they were guided to execute fixations using this optimal strategy.

Figure 8. The proportion of fixations made to the side boxes in the second session, separately for the control and primed groups. Each line is a participant. Full plots for the first and second session, with switch points for each participant, are included in the Supplementary Material.

The distribution of accuracy relative to optimal accuracy for each participant can be seen in the scatterplots in Figure 7. In Session 1, the primed group accuracy even slightly exceeds optimal accuracy, likely because the estimate of optimal is based on the sensitivity mapping phase, which came first in Session 1; accuracy may have improved with practice, leading to better detection in the decision phase. The control group falls below the line. In the second session, where both groups freely choose which box to fixate, there is no discernable difference between the control and primed groups.

Figure 7. Actual and optimal accuracy for each participant in the first and second session of Experiment 3. The primed group were guided to make optimal fixations in the first session, and both groups made free decisions about where to fixate in the second session. Actual accuracy slightly out-performs optimal accuracy, likely because discrimination performance improved from the initial visual sensitivity session (on which optimal estimates are based) to the first session of the main experiment that immediately followed it. There is no difference between groups in the second session.

To analyze these detection performance results, we calculated a difference score between each participant's expected target discrimination accuracy (under an optimal model) and their actual observed accuracy, and compared the difference scores for the decision phase. This gives a measure of how close to optimal each participant is (in terms of Figure 7, this score describes how far from the line each point is). During Session 1, participants who were guided to make optimal choices were, of course, optimal (i.e. they had a difference score close to 0). This score was lower for the instructed participants in Session 1 than for the control group, who received no guidance (between-group $t(22) = 4.89$; $p < .001$). In other words, being guided to optimal choices about which box to fixate does indeed significantly improve detection accuracy relative to when these choices are freely made (see the Verification Phase at the end of the experiment reported by Morvan and Maloney, 2012, for similar evidence that guided optimal fixations can significantly improve detection performance relative to when participants freely choose). For the control group, accuracy is nominally closer to optimal accuracy in Session 2 compared to Session 1, presumably due to practice, but this is not significant ($t(22) = .866$; $p = .395$). To address the main question of whether training improves performance, we look at Session 2 performance. Participants who initially were guided to make eye movements that maximize their chance of correct target discrimination were hypothesized to continue to perform optimally when freely choosing where to fixate in Session 2, and should therefore differ from optimal less than the control group. However, a t-test comparing Session 2 performance (difference from optimal) between the two groups was not significant ($t(22) = .581$; $p = .566$).

To model the data, we used a Bayesian generalised linear mixed effect model to compare the differences in accuracy between the two groups, and to compare the effects of practice and instruction on accuracy over difference values of delta. The advantage of using this approach is that it allows us to estimate the effect size, and the associated uncertainty. Modelling was done using rStan v2.17.3 (Carpenter et al., 2017) and the rethinking package v1.59 (McElreath, 2016) for R 3.4.0. A logit link was used and we accounted for differences between individuals by allowing intercepts to vary. We included three fixed effects: group (control/instructed); session (in the first session, one group was instructed where to saccade to, while in the second session, neither group were given instruction); and delta (the separation between the centre and side boxes). The model that allowed for all two-way interactions and a three-way interaction between these predictors outperformed the other models we considered (using WAIC) so we report results from that model. Full details of the modelling procedure are given in the R script [add link to OSF after publication].

To visualise the model, we computed the predicted accuracy for the average participant, and recoded our variables to give: baseline (the control group in the first session), instructed (the instructed group in the

first session), practice (the control group in the second session) and transfer (the instructed group in the second session). The posterior distributions are shown in Figure 9. We can see there is some evidence for an effect of learning for the small separations (the transfer condition outperforms the effect of practice), but this does not hold over for the large separations. A plausible interpretation of this result is that instructions biased people to fixate the center more often than the sides; this bias promotes near-optimal performance for the close separations, to the detriment of performance at the far separations. Alternatively, there may have simply been more participants with a central bias in our primed group to begin with.

Figure 9. Posterior distributions showing the predicted accuracy for the average participant across three box separations. Baseline corresponds to the control group’s first block, and Practice their second block. Instructed corresponds to the instructed group’s first block, and Transfer to their second block. Instruction improves accuracy across all separations. Transfer is slightly better than Practice at the smaller separations, but this difference is not observed at the farthest separations. There are also small practice effects at the close and medium separations, but at the farthest separation this distribution almost perfectly overlaps with the baseline.

General Discussion

In these experiments, optimal performance hinges on understanding the domain-general principle that should guide the decision. Even though this principal is intuitive once it is known, and could be implemented easily by non-naïve participants (as demonstrated in Experiment 1), people do not appear to spontaneously construct or exploit this principle. They do not apply the correct response more broadly even when they have just implemented the principle in another context (in Experiment 2), nor do they generalize from having experienced the benefits of using it in the same context (in Experiment 3). Instead, participants appear to default to a variable, sub-optimal patterns of responding.

The results may initially seem to be at odds with previous research finding that guiding eye movements can guide thought (Grant and Spivey, 2003; Thomas and Lleras, 2009). Similar conclusions have been drawn from studies showing that gesturing can influence strategies in problem-solving (e.g. Alibali, Spencer, Knox and Kita, 2011). These results all suggest that motor behaviour can translate to abstract thought to promote solution-finding in difficult problems, consistent with an “embodied cognition” framework (e.g. Wilson, 2002). Based on this framework, one might have expected that performing responses consistent with the optimal solution would have prompted at least some of our participants to recognize the solution and apply it in future situations, but this did not occur. As noted in the introduction, contemporary research on problem-solving has tended to focus on relatively complex, abstract problems with a single correct solution that is easily recognized when achieved. These kinds of problems usually require reason to be solved; trial-and-error learning of the solution is generally not an efficient approach. Our results do not contradict the notion that actions can guide people towards solving these more complex and abstract kinds of problems, because the problem we have presented in these experiments, like most of the problems we encounter in our daily lives, does not lead to an impasse if it is not solved, but instead leads to less efficient behaviour if it is not solved *optimally*. Indeed, the solutions to the routine dilemmas of a typical day (what to eat, what to wear, what to say) are highly context-dependent and difficult to predict, and most people would agree it is best not to “over-think” these.

The majority of participants in the current experiment had some degree of bias towards one choice option or another, but also a large amount of variability from trial to trial under the same conditions. However, not all participants were variable, with a minority responding the same way on nearly every trial; for example, four out of the 32 participants in Experiment 2 always stood in the center to throw. Effort minimization (Irons and Leber, 2016), stress (Shors and Dryver, 1992), and distraction (e.g. Wolford, Newman, Miller & Wig, 2004) could contribute to rigid or stereotyped behaviour among a subset of participants. Factors that can contribute to differences in variability are discussed by Neuringer (2002), and include age, training, and the perceived constraints of the environment. The throwing experiment takes place outdoors, and while it is sheltered from rain and most wind directions, the temperature conditions do vary, and anecdotally, our experimenters have observed that participants are more likely to stand in the center when the conditions are especially cold. This is consistent with the notion that the conditions can influence the choices of the participants, shifting them towards or away from variability (Maier, 1940). In Krechevsky’s (1937) experiments, the variability in choices in the brain-damaged rats was restricted relative to the healthy rats. All these findings suggest choice variability as a kind of “luxury”, expressed under comfortable but uncertain conditions. The sub-optimal but invariant behaviour of a minority of participants in our sample is also consistent with the concept of *habits* as defined by Dickinson (1985): In rats, a response that has been consistently rewarded can, over time, lead to

perseverative responses after the behaviour ceases to be reinforced. An intriguing recent study by DeWit, Kindt, Knot et al (2018), however, failed to induce habits in humans under similar protocols to those used previously in rats. In other words, people were highly sensitive to changes in the contingencies between responses and their reward or punishment, even after over-training with the original contingencies. This finding, together with the general tendency towards variability observed in most of our participants, suggests intriguing relationships between variability in behaviour and flexibility in responding to environmental change that remain to be explored in humans.

As noted in the introduction, the variability observed both between and within individuals in these experiments precludes an explanation for the results based on a single cognitive bias or motivation. However, idiosyncratic behaviour could be due to participants just not caring particularly about their own accuracy. Arguing against this, participants do express intrinsic satisfaction in getting the beanbags in the hoops (and frustration at missing), and Experiment 2 required only 45 trials; participants did not appear bored or disengaged. Experiment 3 was longer and accuracy in the detection task is, admittedly, less motivating than landing the beanbag inside the hoops. However, in their original experiment exploring optimal eye movements upon which this experiment is based, Morvan and Maloney (2012) gave substantial financial incentives for detection accuracy and found the same, sub-optimal pattern. One cannot rule out that even higher stakes could elicit more optimal choices, however, and this is a question we are exploring in further experiments.

An intriguing trade-off is implied by our results: applying a consistent decision rule based on knowledge will restrict choice variability. Restricting this variability stunts the potential for learning to shape behaviour in a way that flexibly adapts to dynamic and unpredictable environments. If applying reason comes at a cost to potential learning, a conservative use of reason might be warranted, especially in solving the simple, repeated problems of daily life. An important reminder here is that we are using Maier's (1940) definitions of both variability and learning in this context, where variability is an approach to solving a problem (such as a rat choosing a direction in a maze) that allows the animal to define and refine the array of choice options and their consequences, and learning is the tendency to repeat choices that lead to positive outcomes. In our experiments, the variable but sub-optimal choices of our participants over the series of trials is consistent with a trial-and-error approach, and the failure to recognize or maintain what is, in retrospect, an obvious strategy suggests this trial-and-error approach is a strong and persistent default mode. Although this explanation for our results is speculative, a similar tradeoff between learning and logic has been previously observed in the context of probability matching. When asked to repeatedly guess which of two events will occur on each trial, participants tend to be sensitive to the probability of these events. That is, an event that occurs with a probability of .8 will be

guessed more often than one with a probability of .2. In fact, participants tend to match the ratio of their guesses to the probability of the events. Of course, this is a sub-optimal strategy for maximizing guessing accuracy, resulting in an overall success rate (in this example) of .68, compared to the .8 participants could attain if they guessed the event with the higher probability on every trial. This tendency towards probability matching has been interpreted as a cognitive limitation (e.g. West and Stanovich, 2003). However, Gaissmaier and Schooler (2008) noted that maximizing by only ever guessing the more likely event can become sub-optimal if the sequence is not actually random, that is, if there are patterns in the sequence that can be discovered and exploited. Indeed, participants who probability-match (in a random sequence block) were more likely to detect and exploit a pattern in the sequence when it was introduced in a separate block. This is consistent with the notion that applying a constant rule (e.g., always guessing the more likely event) can optimize accuracy as long as the critical conditions persist (i.e. that the sequence stays random, and the likely event continues to be more likely). In a dynamic natural environment, random sequences are rare, and patterns can emerge and change. The cost of applying a consistent response rule is that the behaviours that could be reinforced by the current conditions will not be executed, and thus will never be reinforced and repeated.

The variable responses we observed in the choice behaviour of the participants in our experiments could in part be inherent to the participant; under conditions of choice uncertainty, there may be an element of random selection between options to avoid becoming fixed in a particular stimulus-response pattern and missing opportunities for learning, as described above. Some variability may also come from trial-to-trial variations in the immediate context. Summerfield and Tsetsos (2015) argue that inefficient economic decisions can be accounted for in the context of *efficient coding*: neural representations of different choice options emphasize the features that are most diagnostic in the immediate context. In other words, the relative appeal of different choices will vary according to the local context, as a result of a mechanism that facilitates their differentiation across a wide range of potential circumstances. Being sensitive to idiosyncratic changes in the environment can lead to behaviour that is sub-optimal in terms of maximizing potential gains and minimizing risk and energy expenditure in an environment that, over the long term, is stable. Many so-called “sub-optimal” choices, like probability matching, could result from over-sensitivity to short-term, idiosyncratic variations in local context. Similarly, opting for larger reward later over a smaller one now is only optimal if the environment is stable. However, environments are rarely stable, and incorrectly assuming they are could lead to disastrous outcomes, such as giving up a short-term sure reward in favour of a longer-term uncertain reward that never comes to fruition. Kolling, Wittman and Rushworth (2014), for example, showed that participants can strategically shift from being risk-averse to risk-prone depending on how large the payoffs are and how many chances they have

remaining to achieve a goal. In this case, as in probability matching, being sensitive to changes in the local context yields better outcomes than a consistently applied principle that incorrectly assumes a stable world. The trade-off between reinforcement learning on the one hand, and following logical decision rules on the other, is also reminiscent of the classic distinction between exploration and exploitation in organizational psychology (March, 1991).

In conclusion, we have found a simple and intuitive decision rule is not recognized or adopted by most participants. Choices are governed by a complex set of individual and contextual factors and their interaction, even for seemingly simple decisions such as a rat deciding which way to turn at a junction point in a maze (e.g. Tolman, 1938). Nonetheless, based on the overall pattern of behaviour of the naive participants who completed the choice task in the current set of experiments, we can conclude that few, if any, settled on an optimal strategy even after being guided to make optimal choices in a similar, or the same, context. After training participants to make optimal choices, the majority immediately returned to making variable and idiosyncratic decisions. Given how strong and persistent this behaviour is, the benefits of this variability may, in many circumstances, outweigh the costs of finding and applying consistent rules.

References

- Alibali, M.W., Spencer, R.C., Knox, L. and Kita, S. (2011). Spontaneous gestures influence strategy choices in problem solving. *Psychological Science*, 22, 1138-1144.
- Atkinson, J. W., & Litwin, G. H. (1960). Achievement motive and test anxiety conceived as motive to approach success and motive to avoid failure. *The Journal of Abnormal and Social Psychology*, 60, 52-63.
- Brainard, D.H. (1997). The psychophysics toolbox. *Spatial Vision*, 10, 433-436.
- Carpenter, B., Gelman, A., Hoffman, M. D., Lee, D., Goodrich, B., Betancourt, M., Brubaker, M., Guo, J., Li, P., & Riddell, A. (2017). Stan: A probabilistic programming language. *Journal of statistical software*, 76(1).
- Clarke, A.D.F., Greene, P., Chantler, M.J. & Hunt, A.R. (2016). Human search for a target on a textured background is consistent with a stochastic model. *Journal of Vision*, 16(7):4
- Clarke, A.D.F., Stainer, M., Tatler, B. & Hunt, A.R. (2017). The saccadic flow baseline: Accounting for image-independent biases in saccadic behaviour. *Journal of Vision*, 17(11):12.
- Clarke, A.D.F. & Hunt, A.R. (2016). Failure of intuition when choosing whether to invest in a single goal or split resources between two goals. *Psychological Science*, 27, 64–74.

- Cornelissen, F.W., Peters, E.M. & Palmer, J. (2002). *Behavior Research Methods, Instruments, & Computers*, 34, 613.
- Dewsbury, D.A. (2000). Comparative cognition in the 1930's. *Psychonomic Bulletin and Review*, 7, 267-283.
- Dickinson, A. (1985). Actions and habits: the development of behavioural autonomy. *Philosophical Transactions of the Royal Society B*, 308, 67-78.
- Gaissmeier, W. & Schooler, L.J. (2008). The smart potential behind probability matching. *Cognition*, 109, 416-422.
- Grant, E.R. & Spivey, M.J. (2003). Eye movements and problem solving: guiding attention guides thought. *Psychological Science*, 14, 462-466.
- James, W., Clarke, A.D.F. & Hunt, AR. (2017). Failure to use probability of success when deciding whether to pursue one goal or two. in *Proceedings of the 39th annual meeting of the Cognitive Science Society*. Cogsci, pp. 2279-2283.
- Kahnemann, D. & Tversky, A. (1984). Choices, values, and frames. *American Psychologist*, 39, 341-350.
- Kolling, N., Wittman, M. and Rushworth, M.F.S. (2014). Multiple neural mechanisms of decision making and their competition under changing risk pressure. *Neuron*, 81, 1190-1202.
- Krechevsky, I. (1937). Brain mechanisms and variability I: Variability within a means-end-readiness. *Journal of Comparative Psychology*, 23, 121-138.
- Krechevsky, I. (1937). Brain mechanisms and variability II: Variability where no learning is involved. *Journal of Comparative Psychology*, 23, 139-163.
- Maier, N.R.F. (1940). The behavior mechanisms concerned with problem solving. *Psychological Review*, 47, 43-58.
- March, J.G. (1991). Exploration and exploitation in organizational learning. *Organizational Science*, 2, 1-147.
- McClelland, D.C., Atkinson, J.W., Clark, R.A. & Lowell, E.L. (1953). *The Achievement Motive*. New York: Appleton-Century-Crofts.
- Metcalfe, J. and Wiebe, D. (1987). Intuition in insight and noninsight problem solving. *Memory and Cognition*, 15, 238-246.
- Morvan, C. and Maloney, L. T. (2012). Human visual search does not maximize the post-saccadic probability of identifying targets. *PLoS Computational Biology*, 8(2):e1002342.
- Najemnik, J., & Geisler, W. S. (2008). Eye movement statistics in humans are consistent with an optimal search strategy. *Journal of Vision*, 8 (3):4.
- Neuringer, A. (2002). Operant variability: Evidence, functions, and theory. *Psychonomic Bulletin and Review*, 9, 672-705.

- Nowakowska, A., Clarke, A.D.F. & Hunt, A.R. (2017). Human visual search is far from ideal. *Proceedings of the Royal Society B: Biological Science*, 284, 2016.2676.
- O'Doherty, J.P., Cockburn, J. & Pauli, W.M. (2017). Learning, reward, and decision making. *Annual Review of Psychology*, 68, 73-100.
- Shors, T.J. & Dryver, E. (1992). Stress impedes exploration and the acquisition of spatial information in the right-arm radial maze. *Psychobiology*, 20, 247-253.
- Summerfield, C. & Tsetos, K. (2015). Do humans make good decisions? *Trends in Cognitive Sciences*, 19, 27-34.
- Tatler, B.W. (2007). The central fixation bias in scene viewing: Selecting an optimal viewing position independently of motor biases and image feature distributions. *Journal of Vision*, 7(14):4.
- Thomas, L. & Lleras, A. (2009). Swinging into thought: directed movement guides insight problem solving. *Psychonomic Bulletin and Review*, 16, 719-723.
- Tolman, E.C. (1938). The determiners of behavior at a choice point. *Psychological Review*, 45, 1-41.
- Weisberg, R., DiCamillo, M. & Phillips, D. (1978). Transferring old associations to new situations: a non-automatic process. *Journal of Verbal Learning and Verbal Behaviour* 17, 219-228.
- Weisberg, R.W. (2006). *Creativity: Understanding innovation in problem solving, science, invention, and the arts*. John Wiley & Sons.
- West, R.F. & Stanovich, K.E. (2003). Is probability matching smart? Associations between probabilistic choices and cognitive ability. *Memory & Cognition*, 31, 243-251.
- Wilson, M. (2002). Six views of embodied cognition. *Psychonomic Bulletin & Review*, 9, 625-636.
- Wolford, G., Newman, S., Miller, M. B., & Wig, G. (2004). Searching for patterns in random sequences. *Canadian Journal of Experimental Psychology*, 58, 221–228.

Acknowledgements. The authors would like to thank Alex Irvine, Ellen Angus, and Anna Nowakowska for help with aspects of data collection, and to the Perception group at the University of Aberdeen for helpful comments on this work. This research was made possible through support from the James S. McDonnell Foundation (Scholar Award to ARH).