



Variable and Sub-Optimal Responses to a Choice Problem Persist Despite Experience With a Simple and Effective Decision Rule

Amelia R. Hunt¹, Warren R. G. James¹, Josephine Reuther²,
Melissa Spilioti¹, and Alasdair D. F. Clarke³

Abstract

Choices in a focus-divide dilemma tend to be variable and suboptimal. How readily will people adopt a simple, rational choice rule to improve performance? We asked participants to choose where to stand to throw a beanbag into one of two possible hoops (Experiments 1A, 1B, 1C), and to choose where to fixate to discriminate a target that could appear in one of two possible boxes (Experiment 2). The rational choice follows a simple logic: when hoops/boxes are close together, standing/fixating near their midpoint is optimal. When hoops/boxes are far apart, standing at/fixating one hoop/box will optimise success. We attempted to nudge participants towards optimal choices by providing circumstances in which they carried out the correct solution, and then returning them to the original context to freely choose. In both experiments, participants executed optimal choices during the intervention, and then immediately returned to a sub-optimal pattern of responding when presented with the original context. Experiments 3A and 3B demonstrate that the rational decision rule can be easily applied and improves performance. We conclude that variable and idiosyncratic responses are a persistent choice strategy. Borrowing concepts from classic animal learning studies, we suggest rule-following risks missing opportunities for reinforcement learning afforded by choice variability.

Keywords

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

Received: 6 May 2025; revised: 12 August 2026; accepted: 16 August 2026

Introduction

Imagine deciding which home improvement project to embark on with a limited budget. You may desire to refurbish both your kitchen and your bathroom, but attempting to do both at once is expensive. Having one of these projects fully completed is surely more desirable than running out of money halfway through and being left with only partially useable rooms. However, if you have a large enough budget, it is faster and more convenient to get all the work done at once. This is an example of a *focus-divide* dilemma (James et al., 2023): How should we allocate limited resources to achieve competing goals? A rational principle to apply when faced with such decisions would be to use an estimate of the resources required to achieve the goals to decide how many goals to try and accomplish. Given two

relatively easy goals, it makes sense to try and accomplish both. But if the combined difficulty of achieving the two goals exceeds the available resources, focusing on accomplishing one is usually the better strategy. Although this seems like simple logic, people generally fail to apply it in experimental adaptations of the focus-divide dilemma (e.g. Clarke & Hunt, 2016; Hesse et al., 2020; James et al.,

¹School of Psychology, University of Aberdeen, UK

²Department of Experimental Psychology, University of Göttingen, Germany

³School of Psychology, University of Essex, UK

Corresponding Author:

Amelia R. Hunt, School of Psychology, University of Aberdeen, Aberdeen AB24 3FX, UK.

Email: a.hunt@abdn.ac.uk

2017). In the series of experiments presented in this article, we attempt to train or ‘nudge’ participants towards the logical solution. Whether or not participants can apply the decision rule when we make it available sheds light on the cause of the failure in the first place. Moreover, being able to improve choice behaviours with simple interventions opens avenues for minimising or preventing the potential real-world consequences of these common choice errors.

Clarke and Hunt (2016) posed focus-divide decision dilemmas in a range of different contexts: Throwing beanbags into hoops, discriminating low-contrast targets in peripheral vision, and memorising sequences of digits. In each version of the dilemma, participants first performed a single task multiple times over a range of difficulty levels to establish their skill level and limitations. For example, participants were asked to throw a beanbag into hoops over a wide range of distances. When the hoop is close, this task is trivially easy, and when it is far away, the task becomes challenging. In the second *decision* phase of these experiments, participants were given two *potential* goals, and they had to decide whether to try and accomplish both or focus on one or the other. In the throwing version of the decision phase, the participants’ goal is to throw a beanbag into one of two hoops, but they are not told which of the two is the target until after they choose where to stand. They could choose a spot equidistant between the two hoops, to maximise their chances of getting the beanbag in the hoops irrespective of which one is ultimately the target (divide). Alternatively, they could choose to stand next to one hoop or the other and be guaranteed success if it turns out to be the target (focus). Which of these strategies will lead to a higher probability of success depends on the distance between the hoops. If they are close together, the best choice is to stand at the midpoint between them. However, if the hoops are far enough apart that the expected odds of success in hitting either one from a central position is less than 50%, standing near one of the two hoop leads to better success rates: 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%.

Even though the logical solution to the dilemma is simple to understand and implement, the results from Clarke and Hunt (2016) demonstrate that in this and two other tasks (visual discrimination and memorising strings of digits), naive participants failed to maximise their potential accuracy on the task. Moreover, most of them did not modify their focus-divide choices with manipulations of task difficulty at all. A few participants consistently stuck to a ‘central’ strategy in which they always split their resources between both goals, leading them to perform poorly as difficulty increased. Most participants, however, exhibit a large amount of trial-to-trial variability, focusing on some trials and dividing on others, but not in a way that tracked with the demands of achieving the goals. This is important because it shows the limitation on performance in this task

cannot be easily described or explained as one particular bias or error (e.g. to always divide). This pattern of variable, sub-optimal choices in the focus-divide task is now well-established (James et al., 2017; Hesse et al., 2020; Morvan and Maloney, 2012) and several of the more obvious possible explanations have been ruled out (James et al., 2019; James et al., 2023). The tentative explanation advanced by Clarke and Hunt (2016) is that in daily life, prioritisation problems are usually difficult to solve because they involve many options and high uncertainty. Participants fail to recognise that the problem in this specific instance has a simple and easily implemented solution, so instead of trying to find a solution, participants default to heuristics, guesses, and trial-and-error strategies.

Recognising that a problem has a tractable solution is a necessary first step for solving it, and this recognition could be considered an ‘insight’. Insight is associated with the ‘aha’ moment that comes with solving some specific types of problems, classically known as *restructuring*, reflecting the idea that the underlying mechanism driving insight is a re-organisation that allows the same information to be represented in a different way. Weisberg (2006) argues that for a problem to be defined as an insight problem, there needs to be a discontinuity in the solution-finding process that is caused by restructuring, and this restructuring has to be necessary for a solution to be achieved. The focus-divide dilemma does not quite meet this definition of an insight problem, because participants can complete each decision with or without insight about the optimal strategy. That is, a failure to recognise the optimal solution does not cause an impasse in the focus-divide dilemma. Instead, participants can proceed through the experimental trials, repeatedly making choices and observing the consequences of those choices, but with a lower rate of success than they would have if they restructured the problem and implemented the optimal solution.

The well-known ‘dual-systems’ models (e.g. Kahneman, 2011; Sloman, 1996; Stanovich & West, 2000) propose that choices can be a consequence of either a rapid, automatic route or a more deliberative, logic-based route. These systems represent a very broad categorical distinction in cognitive reasoning and are normally applied to judgement and problem-solving contexts, which tend to pose complex, abstract scenarios. Classic animal learning studies draw analogous distinctions between choices guided by reflex versus those guided by reasoning, and may provide a more germane theoretical framework for the simple, repeated choices participants make in focus-divide tasks. Maier (1940), for example, argued for three different categories of solutions to choice problems: Variability, equivalence reactions, and reasoning. *Variability* is characterised as a functional tendency that is assumed to support trial-and-error learning, and stands in contrast with the repetitive, inflexible behaviour that is a

hallmark of cortical damage. Krechevsky (1937a, 1937b) argues that variability in choices allows the organism to explore the set of available choice options and their consequences within a particular problem space. 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. This process permits acceleration of the learning process when circumstances are familiar. *Reasoning* refers to the ability to spontaneously integrate information from outside the experimental context and use it to guide choices. A similar distinction was later drawn by Dickinson (1985), who referred to responses elicited by immediate stimulus conditions as *habits*, and contrasted these with responses that are guided by an understanding of anticipated consequences, which he referred to as *actions*. This dichotomy has held up well in the literature, with accruing evidence that habits and actions are subserved by distinct neural systems (e.g. O'Doherty et al., 2017).

Viewed from an animal learning perspective, the seemingly idiosyncratic series of choices observed in a typical focus-divide experiment is consistent with *Variability*, that is, an exploratory, trial-and-error strategy. Our question in this study is about whether and how we can transition people from *Variability* to *Reasoning* in the focus-divide task (to use Maier's (1940) taxonomy). Applying the correct solution in one circumstance and experiencing success could lead to the application of the same solution in circumstances with similar properties. This generalisation may be sufficient for the kind of 'aha' moment that is associated with discovering an abstract decision rule that can then be exploited in a range of similar circumstances. On the other hand, applying a decision rule carries risks. If applied to the wrong circumstances, a rule can lead to consistently suboptimal outcomes. Moreover, rule-following restricts the variability of responses that is the basis for learning, as seen in trade-offs of the classic 'explore-exploit' dilemma (e.g. Wilson et al., 2014). Given these risks, taking a conservative approach to avoid over-generalisation 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. Consistent with this view, experience in solving complex problems is usually context-specific and resistant to transfer (e.g. Gick & Holyoak, 1983; Markovits & Savary, 1992).

In the following experiments, we explore the threshold for transfer of decision rules across contexts by creating circumstances in which participants carry out the optimal strategy, and then provide a similar (Experiments 1A, B, and C) and the same (Experiment 2) context to observe whether the strategy persists. The experiments shed light on how readily participants abandon their trial-and-error strategy and embrace a consistent, rational rule. These experiments also help us establish the malleability of

people's failure to solve the focus-divide dilemma, that is, how easily it could be overcome with some guidance or reflection, and thereby how likely or unlikely it is to permeate decisions in everyday life. A common response when we explain the correct solution to participants during debriefing is that it seems obvious in retrospect, and many participants express chagrin at not having realised it themselves. This reinforces that the correct decision rule is only just out of reach for participants, and that a small nudge in the right direction would be sufficient for them to be able to arrive at this decision on their own. Our experiments provide conditions that we expected to lead them to this insight, but without explicitly stating the rule.

To foreshadow, participants can carry out the optimal solution when clearly guided to do so by the context, but show no benefits from this experience relative to control groups when the guidance is no longer available. To be confident that the source of this limitation is a failure to recognise the solution rather than a failure to execute it effectively, Experiment 3 presents a control experiment in which we test non-naïve participants under the same conditions used in Experiments 1 and 2. These participants explicitly know the rule for solving the focus-divide dilemma and can execute the rule easily. Taken together, the series of experiments suggests that even when rational decision rules are readily available, easily implemented and beneficial for performance, the trial-and-error strategy seems to be a surprisingly stable behavioural tendency.

Experiment 1A: Reaching, Then Throwing

Clarke and Hunt (2016) demonstrated that participants fail to modify their focus-divide choices to take task difficulty into account. The same pattern of variable, sub-optimal behaviour was exhibited across three diverse contexts: visual detection, throwing, and memorising strings of digits. However, in a fourth *reaching* experiment, Clarke and Hunt (2016) simplified the context to present a comparable, but trivial, decision that produced uniformly optimal behaviour. In this simplified version of the task, participants were presented with a long table upon which three pairs of coloured beanbags had been placed. The red beanbags were placed close to one another near the midpoint of the table, while the blue beanbags were placed at either end of the table. The two green beanbags were placed at intermediate position between the blue and red bags. Three chairs were placed at the table: one on the left end, one central and one on the right end (see Figure 1 for a photograph of this setup). On each trial, participants were told they would need to pick up a beanbag. After being told which colour beanbag they would need to pick up (but not which one), participants were asked to select and sit down in one of the three chairs. After they chose a chair, they were told which of the two beanbags (left or right) to pick




Figure 1. The three images down the left side show the task setups for all versions of the focus-divide dilemma presented in this series of experiments. In the reaching task, the participant selects a chair from which to pick up one of two beanbags of a specified colour on a long table. In the throwing task, two each of red, yellow, and blue hoops were taped to the ground, and participants had to choose a place to stand to throw a beanbag into one of the two hoops of a specified colour. In the detection task, a small white dot would appear in one of the two side boxes, and participants had to indicate whether it appeared in the upper or lower half of the box. In all three tasks, participants do not know which side (left or right) will be specified as the target until after they have sat in a chair, chosen a standing position, or fixated in one of the boxes, respectively. The intervention sequence for all the experiments is shown to the right. The effect of the intervention (red boxes) on position choices in the throwing (E1) and detection (E2) tasks is compared to the control group (blue boxes).

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 centre, the participant should select a chair close to one of the beanbags, ensuring they can at least reach that one. Clarke and Hunt (2016) found that all participants consistently used the optimal strategy in this more obvious version of the dilemma. Our approach here is to run the table/reaching version of the dilemma, immediately followed by the throwing version, to see if applying the correct solution to the table task would facilitate better decisions in the throwing version.

We expected that solving the table task would cause at least some participants to keep solving the problem the same way when they carried the beanbags outside to do the throwing task. This was based in part on participants consistently reporting that the solution seemed obvious in retrospect when we explained it in debriefing, leading us to expect the solution to be readily recognised and carried forward from the experience of performing it in the table task. Experimental evidence also indirectly supported this expectation. Claydon et al. (2024) found that uncertainty around the outcome is a critical aspect of poor performance in focus-divide dilemmas. This suggests the lack

of uncertainty in the length of one's reach (relative to throwing, detection and memory capacity) is the likely explanation for optimal choices in the table context. Claydon et al. (2024) used a computer-based version of the dilemma where it was possible to manipulate variability in the outcome parametrically. Participants who completed the decision task under less variable conditions first, and more variable conditions second, were closer to optimal than participants who completed the two conditions in the opposite order. The implication is that strategy was maintained as participants shifted from the more stable into a more variable condition. If the decisions can be transferred from a simple task context to a slightly more complex one, throwing task choices should be close to optimal for participants who just carried out the reaching version of the task.

Methods

Participants. Thirty-two students from the University of Aberdeen were recruited either via Online Recruitment System for Economic Experiments (ORSEE) or Sona Systems Research Management System (SONA). 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. A £5 remuneration was given to the ORSEE participants, and appropriate academic credits were awarded to the SONA participants. In this and all subsequent experiments, the experimental protocol was reviewed and approved by the Aberdeen Psychology Ethics Committee (except 1B below, which was approved by the Essex ethics board), and participants all gave signed, informed consent to participate.

Justification of the Sample. We use a Bayesian approach to data analysis. We apply the same justification for the sample size as used by James et al. (2023): they simulated experiments with a sample size from 3 to 24, resampling from distributions based on the standing positions of participants in the throwing experiment (E2) of Clarke and Hunt (2016). Simulated estimates of an underlying true shift of 5% (i.e. a mean shift of 0.05 in the normalised range of standing position) varied widely with a small N but stabilised after about 12 participants. This analysis suggests that we can expect to produce reliable estimates of the effect of an intervention with a sample size of 12 in each group, and that testing additional participants would have little impact on the results.

Materials and Procedure. Two experimenters conducted this study. The first experimenter carried out the reaching task (detailed below) with half of the participants. The other half of participants remained with the first experimenter for 5 to 10 min 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 (described below) with both groups of participants, and was blind to which condition they were in. This experimenter carried a clipboard with a pre-generated random order of trials (distances and directions) that was different for each participant, and used the list to instruct the participant and record their choices and throwing accuracy.

Reaching Task. Six waterproof beanbags of three colours (two red, two green, and two blue) were placed on a long table equidistant from each other (Figure 1). At the centre of the table were two red beanbags. Halfway across each side of the table were the green beanbags (one on each side). At each far end were the blue beanbags (one on each end). Three chairs were positioned alongside the table: one on each end (next to each blue beanbag) and one in the centre (next to the two red beanbags). To ensure that participants were aware of their own reaching span, they were first asked to sit on the chair positioned at the centre of the table. Sitting with their back resting on the chair they were

asked whether they could reach and touch the red, green and blue beanbags individually (the answer was yes, yes and no, respectively). 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 given the beanbags and instructed to go outside to the other experimenter for the second part of the experiment.

Throwing Task. This was an abbreviated version of the throwing task used in Clarke and Hunt (2016). An outdoor, sheltered area of concrete slabs just outside the psychology building (Figure 1) was used because the slabs (measuring 0.46×0.61 m) were useful for marking and recording hoop placement and standing positions, respectively. Six hoops with a diameter of 0.40 m were taped down in a row with three slabs' distance between them. The red hoops were 4.60 m apart (10 slabs), the yellow hoops were 8.28 m apart (18 slabs), and the blue hoops were 11.96 m apart (26 slabs). These hoop separations were selected on the basis of throwing performance measured in several previous experiments; participants should stand in the centre for the closest distance, and next to one of the hoops for the farthest distance, to achieve optimal performance. The colour of the beanbag the participant was handed on each trial informed them which pair of hoops were potential targets on that trial (i.e. when they were handed a red beanbag, they knew that one of the two red hoops was the target, and so on for the other colours). Participants took the beanbag and chose a place to stand. After choosing their standing position, they were told which hoop was the target (and they were informed on the first trial that the colour and hoop on each trial were determined based on the pre-generated list the experimenter had on their clipboard), 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 the order of colours specified by the trial list). Standing position and throwing success or failure were recorded. There were 15 trials for each distance (45 total).

Analysis. All data processing, analysis, and visualisation was carried out with R (v4.4.0) with the *tidyverse* packages (2.0.0. Wickham et al., 2019). We normalised standing position by dividing it by the distance between the hoop and the centre, so values reflect a proportion of this maximum distance. Statistical models were fit with *brms* and *cmdstanr*. As our normalised standing position measure is skewed, non-negative, and zero-inflated we will use




Figure 2. Results for each individual in Experiment 1a. Each dot is a trial, and each facet is a single participant. The distance from centre to each of the hoops is shown on the x axis. The participants shown in blue are the control group, and those in red were primed by making optimal decisions in the reaching task before completing the throwing task. The crosses show an ‘absolute’ optimal standard, which is standing at 0 at the smallest hoop separation and standing at 1 at the largest.

a multi-level hurdle-lognormal model to analyse the differences in decision between the near and far hoop conditions, and for the primed and control group (more details on this model are given in the ‘Results’ section). We used weakly informative $N(0,1)$ priors and maximal random effect structures for all models.

Results

Reaching Task Choice. All participants who carried out the reaching task did so optimally. That is, when they were told to reach for one of the two red or green beanbags, they sat in the central chair. When they were told they would need to reach for one of the two blue beanbags (too far away to reach from the central chair), they chose either the left or right side chair, resulting in a 50% success rate. No variance, either within, or between participants, was observed. This replicated the behaviour observed by Clarke and Hunt (2016).

Throwing Task Choice. Raw standing position choices in the throwing task are shown in Figure 2. The optimal solution to this problem would give standing position=0 for the closest hoop distance and standing position=1 for the furthest (with 1 being the distance from the centre to the hoop). This benchmark is shown as the crosses in Figure 2. 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 (shown in blue) that participants fail to do so. We can also see that no participant managed to execute the optimal strategy in this version of the focus-divide dilemma. From this we can conclude that experience of solving one form

of dilemma (the reaching task) does not lead to optimal behaviour in another (the throwing task).

We have shown every standing position choice for every participant as a separate point in Figure 2 to depict the large amount of variation in these choices, both within and between participants. Around a third of participants exhibit a strong central tendency, sticking with the mid-point between the hoops irrespective of the distance between them. The rest of the participants vary their behaviour from trial to trial. On a handful of trials, a few participants even chose to stand outside of the range of the hoops, leading to a normalised standing position greater than 1. This behaviour is rare but noteworthy both because it reinforces the wide variance in human choices in this task, and also because it affects the vertical axis of the plots (which would otherwise be 0–1). This pattern of sub-optimal and highly variable choices replicates the striking findings of Clarke and Hunt (2016).

The crosses on each facet of Figure 2 show where participants should have chosen to stand under a strictly optimal model. By eye, there are no obvious differences between the groups, but to drill into this more deeply we modelled the data using a multi-level hurdle-lognormal model. This model breaks choices down into (a) an initial decision about whether to stand at the midpoint, and (b) if not at the midpoint, how far away from the midpoint to stand. The effect of hoop distance should be reflected in both components, with the farther separation making participants both less likely to stand at centre, and if not standing at centre, standing near the hoop (i.e. close to 1). The results of this analysis are shown in Figure 3. We can see that the control group is equally likely to stand in the central position irrespective of hoop separation (95% Highest Posterior Density Interval (HPDI) of [0.4, 0.79] for the




Figure 3. Summary of the fixed effects from the Bayesian hurdle-lognormal model. The saturation of the coloured vertical bands indicates 50%, 80%, and 95% highest posterior continuous intervals for (left) the probability of selecting a central standing position; and for (right) the average distance of the standing position from the centre, specifically on the trials where participants chose not to stand at the centre.

near hoops and [0.25, 0.72] for the far hoops). The primed group, while not following the optimal strategy, do appear to be more likely to stand in the centre when the hoops are near to one another compared to when they are far (95% HPDI of [0.64, 0.95] and [0.07, 0.47], respectively). For the second part of the model (how far from the centre do they stand), neither the primed nor the control group shows evidence that their choices about where to stand are based on the hoop separation (for the control group [0.35, 0.63], for near hoops [0.16, 0.44], for the far hoops, compared with [0.30, 0.73] and [0.35, 0.55] for the primed group). Taking both components of the model into account, we could potentially infer that our primed group was moderately closer to optimal, but the size and inconsistency of this difference suggest it could equally be a consequence of random sampling (i.e. some people just stick to the centre more than others, and we happened to have more of that kind of person in our control group). These results suggest that, counter to our expectations of a large and unambiguous improvement, the reaching task did not make participants substantially better at making throwing task decisions.

Throwing Performance. Throwing accuracy for the three hoop separations in Experiment 1A is shown in Figure 4. The two groups have similar accuracy, aligning with the lack of an effect of the reaching task intervention in the choice data above. In previous versions of the focus-divide dilemma (and in Experiments 2 and 3 below), we included a session to measure participant accuracy with a single task over a range of difficulty so that we could compare their actual throwing performance to a simulated optimal version of their own performance based on single-task data. In this experiment, we were primarily interested in comparing the two groups to each other, so we did not include a single-task session. But we can assume that participants will be accurate on average 50% of the time if

they stand at one hoop in the far condition. Both the primed and control groups fall short of this standard. We cannot estimate what their optimal accuracy would have been for the other two distances, but the far distance is the condition in which there are the largest potential gains for executing an optimal strategy, and both primed and control groups are similarly far from optimal in this condition.

Discussion

We had expected that solving the focus-divide dilemma correctly in the reaching task would produce a large and unambiguous shift towards solving the dilemma optimally in the subsequent throwing version of the task. The results of Experiment 1A show that priming participants with the reaching task did not produce the large shift towards optimality in the throwing task that we were expecting. However, some participants in the primed group did not stand in the central position when the hoops were far apart, which is consistent with being closer to optimal. As this effect is inconsistent and small, but in the expected direction, and there are large individual differences in both groups, we ran two conceptual replications of this experiment in Experiments 1B and 1C to rule out the possibility that there is a robust effect of the intervention that we missed in Experiment 1A due simply to chance. These two experiments reinforced the conclusion that there is no substantial or consistent difference in standing position choices in the throwing task between participants who are primed with reaching relative to control groups.

Experiments 1B and 1C: Replications

Experiment 1B includes both the reaching task intervention used in Experiment 1A and a second intervention in which participants completed a logic problem based on the focus-divide dilemma. The rationale is to confirm the




Figure 4. Throwing accuracy in Experiment 1A. The primed group (red) performed the reaching task first, before completing the throwing task. Boxplots show the median (black line), and their vertical extent is the interquartile range. Optimal accuracy at 6 m would have been 0.5.

results of Experiment 1A and also test whether framing the problem verbally might encourage the participants to use a strategy of *reasoning* (Maier, 1940) that would more readily transfer to the throwing task. A third group completed arithmetic problems. Experiment 1C is similar to 1A, except we asked all participants to complete the throwing task twice, once before and once after the reaching-task intervention. We compared the before/after change to a control group, who completed the throwing task before and after completing a Sudoku puzzle. This mitigated some of the problems with comparing groups of participants on a task that elicits highly variable choice behaviour that could mask or produce spurious group differences. Performing throwing first also might encourage participants to connect the solution they execute in the reaching task to the throwing context they just experienced. The aim of these further experiments was to reinforce (or challenge) the conclusion we have drawn from Experiment 1A, that the optimal solution to the focus-divide dilemma does not readily generalise from one situation to another.

Methods

Participants

Experiment 1B. Twenty-four participants (17 females, mean age 21) from the University of Essex were randomly assigned to 1 of 3 groups: a reaching task group (9 participants), an arithmetic group (8 participants), and a logic puzzle group (7 participants).

Experiment 1C. Twenty participants (11 females, mean age, 22.8) from the University of Aberdeen were randomly assigned to either complete the reaching task or the control condition (10 in each).

Justification of the Sample. Having already established in Experiment 1A that the large effect we predicted is unlikely, the primary goal of these experiments is to corroborate or challenge that conclusion. A small convenience sample is sufficient for this purpose.

Materials and Procedure. The throwing and reaching task components of Experiment 1B were similar to Experiment 1A, with some minor differences in the setting due to being carried out at different sites. For Experiment 1C, they were the same as Experiment 1A unless otherwise stated.

Experiment 1B. Rather than engage in small talk as in Experiment 1A, here participants completed 10 simple arith-




Figure 5. Summary of the fixed effects from the Bayesian hurdle-lognormal model of the data from Experiments 1B and 1C. The saturation of the vertical coloured bands indicates 50%, 80%, and 95% highest posterior density intervals for (left) the probability of selecting a central standing position; (right) average distance from the centre, on trials in which participants chose not to stand at the centre.

metic questions (e.g. “ $27-12=$ ___”). This was compared to the experimental group, who completed the reaching task as described in Experiment 1A. We also introduced a third group, in which participants completed five logic problems. The first two of these presented a hypothetical version of the focus-divide dilemma. The rationale was to frame the focus-divide dilemma like a standard insight problem, which might lead at least some participants to a generalisable solution to the problem. Given the participants are recruited from a U.K. student population, the scenario dealt with revising (studying) for exams:

You find out that you have an exam tomorrow. You have [10HOURS] to revise and there will be two questions where you must only answer one. You know what the two topics are that the questions will be on but you are not confident about either. Note down how you will divide your time in revision.

A second version of the same question substituted ‘10 min’ for ‘10hrs’. The students should revise both topics when they have plenty of time, and be selective when they only have 10min. These were supplemented with three additional, unrelated logic puzzles. Participants in both the logic and arithmetic condition had 10 min to complete the questions. They then completed the throwing task as described in Experiment 1A. Participants consistently solved the reaching task correctly, but not the logic task, so this group can be considered an additional control group in the analysis.

Experiment 1C. Participants carried out this experiment over two sessions, approximately 1 week apart. The first session had two parts. In Part 1, participants stood in the middle of the area and threw 12 beanbags to a single hoop target placed at each of four different hoop distances {1.38, 3.22, 5.06, 6.90 m} in increasing order. The beanbags were cleared out of the way after each trial. They then threw to four new distances {2.30, 4.12, 5.98, 8.74 m} in the oppo-

site direction, for a total of 96 trials. In Part 2, participants then carried out a block of the throwing task, as detailed in Experiment 1, but this time with four, rather than three separations {4.6, 8.28, 11.96, 15.64 m}. Participants carried out six trials for each distance, in a random order. In the second session, participants were randomly assigned to two groups. Half the participants carried out the reaching task, which was the same as that used in Experiment 1A. The other (control) half were given a Sudoku to complete, which participants worked on for 5 min. Both groups then completed the decision trials again, and then a second session of single-hoop throwing measurement, so that improvements in throwing ability over the course of the experiment could be accounted for in calculating optimal standing position in the second session. Finally, participants in the Sudoku group completed the reaching task at the end of the experiment, to confirm that they were all indeed able to successfully execute the optimal strategy in this context (which they were).

Analysis. Standing position choices in the throwing sessions of the two experiments were separately analysed in the same manner as Experiment 1A. We present the results of the two models in the same figure to ease comparisons.

Results

These experiments were analysed using the same Bayesian hurdle-lognormal model as Experiment 1A. The model fits are summarised in Figure 5. We can see that there is a more consistent tendency to stand at the central position when the hoops are close to one another across all conditions in these samples. Because of this, the estimates for the non-central standing positions have high uncertainty, as there are very few examples of this behaviour. When the hoops are far apart, people still stand in the central position

Table 1. Differences in parameter estimates between groups in Experiments 1B and 1C.

Contrast	95% HPDI for difference in probability of standing at centre in far condition	95% HPDI for difference in distance from centre in far condition
Experiment 1B		
Arithmetic minus reaching	[−0.330, 0.178]	[−0.071, 0.475]
Arithmetic minus logic	[−0.172, 0.360]	[−0.227, 0.361]
Reaching minus logic	[−0.030, 0.365]	[−0.400, 0.151]
Experiment 1C		
Reaching minus Sudoku	[−0.245, 0.299]	[−0.042, 0.199]

Note. There are two parameters reported for each contrast, representing the two components of the hurdle-lognormal model (tendency to stand at centre, and distance from centre). We have reported this interval for the far condition only because this is the condition in which the differences can be expected to be largest between an optimally-performing group and the standard/control groups.




Figure 6. Throwing accuracy over distance for the three groups in Experiment 1B and two groups in Experiment 1A, with each group as a facet. Note that there were two medium distances in Experiment 1C which have been merged in this plot. Each point represents a participant.

around a third of the time in Experiment 1B (95% HPDIs of [0.13, 0.48], [0.23, 0.65] and [0.29, 0.71] for the three groups, respectively), and a 10th of the time in Experiment 1C (95% HPDIs of [0.08, 0.28] for the control (Sudoku) group and [0.05, 0.16] for the reaching group). In trials in which participants chose a non-central position, they move around 0.25 to 0.75 units away from the central midpoint towards one of the hoops in the far condition.

In Experiment 1C, the preference to stand at the midpoint between the two hoops in the near condition is stronger than the preference observed in both Experiments 1A and 1B, and when the hoops are far from one another, participants are less likely to stand in the centre. This is closer to an optimal strategy for the throwing experiment than we have seen in all other versions of this decision dilemma, a point we return to in the discussion. Nonetheless, standing position choices are still far from optimal, leaving room for the reaching intervention to improve strategies relative to the control (Sudoku) group. To establish whether there are improvements in any of the groups, the distribution of the differences in the parameter estimates between groups is reported in Table 1 as 95% HPDIs. When these intervals include 0 (as all of them do),

this indicates there is no strong difference between groups on that parameter.

In summary, there is little evidence that our intervention causes a robust change in behaviour in the focus-divide dilemma. For completeness, throwing accuracy of all the participants in both experiments is reported in Figure 6 (the high degree of variability across participants and small sample makes points a better representation than aggregate data). From the plots, the accuracy reinforces the interpretation based on modelling the standing position data: the groups appear broadly similar, and no single group stands out as having achieved higher accuracy than the others.

Discussion

The results of these three similar experiments indicate that performing the reaching task optimally is not an effective intervention for shifting participants towards more optimal throwing task choices. The conclusion we drew from Experiment 1A continued to hold true for Experiments 1B and 1C: Making optimal decisions in a more obvious version of the task does not lead to systematic improvements in the subsequent throwing task. It is an important caveat

that the participants sampled in Experiment 1C (both in control and intervention groups) were closer to optimal than all the other participant groups reported in this paper, or in previous versions of the throwing experiment (Clarke & Hunt, 2016; James et al., 2017; James et al., 2023). These participants were part of the same population of Aberdeen University psychology students from whom we sampled for Experiment 1A. Although we excluded participants who completed one from participating in the other, in the context of the typical psychology experiment (computer-based tasks and/or questionnaires), throwing beanbags into hoops outdoors is an unusual experimental procedure which may have generated discussion within the cohort. It is our suspicion that participants who were debriefed may have disclosed their understanding of the optimal strategy to other students in the cohort, raising the general level of awareness for the population as a whole. Experiment 3 below demonstrates that having explicit knowledge of the optimal strategy does enable participants to implement this strategy effectively, and we think that may have been why we saw our population's performance improve overall. Nonetheless, participants in Experiment 1C were still far from the ceiling in executing an optimal strategy, leaving room for the reaching task intervention to have an effect. The fact that the reaching task intervention did not have an effect in this group is consistent with the results from Experiments 1A and 1B and suggests that standing position choices in the throwing task are highly variable across participants and generally suboptimal, even in a group who have recently solved the problem optimally in a simpler context.

Experiment 2: Guiding Optimal Eye Movements

In Experiment 1, we did not find evidence that optimal choices during the reaching task readily influenced standing position choices in the throwing task. Participants solved the reaching choice problem but did not generalise this behaviour to new circumstances. Our expectation that it would generalise came from two sources: (a) the solution seems widely agreed to be 'obvious in retrospect', based on conversations with past participants during debriefing, and so we expected participants would be able to recognise and implement the solution when they had just carried it out, and (b) in the experiments reported by Claydon et al. (2024), performing focus-divide choices under conditions of low variance in outcome improved performance under subsequent conditions of high variance, which we considered to be analogous to reaching (where the effect of choosing a chair on reaching the beanbags is certain) versus throwing (where the effect of choosing a standing position on hitting the target is less certain). One plausible explanation for the unexpected lack of an effect of reaching on throwing in Experiment 1 is the

change in context. When participants finished the reaching task and went outdoors for the next part of the experiment, they were given the beanbags to take with them and use in the throwing task, and had a similar layout of colours, with red as the 'close' condition and blue as the 'far'. The stable object and colours across both tasks were expected to serve as a cue for the spontaneous integration we described in the introduction, which we thought would facilitate better choices in standing positions. But in Claydon et al. (2024), they were presented with the same choice (where to park a firetruck) across low and high variance conditions. 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. More generally, it has long been known that learning and memory are highly context-specific (e.g. Carr, 1925; Tulving & Thomson, 1973), a reasonable tendency considering the risk of overgeneralising a learned environmental contingency beyond its limits. Thus, although the group who were exposed to the reaching task understood and applied the correct logic, the change in circumstances (inside to outside, reaching to throwing) may have been too different for them to apply the same logic in the new situation, so they returned to solving the problem using variability.

In the next experiment, we therefore eliminated the change in the task. To implement this, we used the detection task version of the focus-divide dilemma. This is the original version of the paradigm as presented by Clarke and Hunt (2016), adapting a task devised by Morvan and Maloney (2012) to understand fixation choices. In this variant, participants are shown three boxes in a row across the screen (Figure 1). The participant first chooses a box to fixate, and once they are looking in their chosen location, the target (a small grey circle) appears in one of the two outer boxes. When the boxes are close together, the participant can choose the centre box, and use their peripheral vision to detect the target regardless of whether it is in the left or right box. As the distance between the boxes increases, the target is no longer visible from the centre box. For large distances, therefore, choosing the left or right box yields a higher chance of success. As with the other versions of the focus-divide dilemma, participants largely fail to implement this strategy. Most participants fixate the centre box on some trials and one of the side boxes on other trials, but their choice of boxes does not vary with the distance between the boxes (see also W. R. G. James et al., 2019).

We guided half the participants to execute optimal fixation decisions. Specifically, these participants completed a training session where they were cued to fixate a particular box on each trial. The target would only appear after they had fixated the cued box. The cue directed the participants' fixations to the location that would optimise their detection accuracy, given their own visual acuity as measured in a preceding sensitivity mapping phase. After completing

this session, they completed a session with no cues, where they freely chose which box to fixate. 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. We changed to visual detection from throwing because this variant of the dilemma is a computer-based task, so each participant can easily complete hundreds of trials. This variant also allows us to instruct the optimal behaviour on each trial in a consistent manner and minimises the possibility of experimenter-driven effects. With the larger number of trials and more limited set of three options to choose between, we thought participants might have a stronger chance of learning the optimal choices. Finally, there is also evidence that solutions to insight problems can be provoked by guiding the eyes to relevant locations (Grant & Spivey, 2003) or the body in a pattern that is consistent with the solution (Thomas & Lleras, 2009). This suggests that a pattern of behaviour can lead participants to restructure information in a way that is consistent with that behaviour. This experiment's methods and analysis plan were pre-registered on the Open Science Framework: <https://osf.io/deh4j/overview>.

Methods

Participants. Twenty-four participants (mean age 21, range 18–26, 19 females, 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 course 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.

Justification of the Sample. As for Experiment 1, we will apply a Bayesian analysis to model uncertainty around the effect of the intervention. The estimate of $N=12$ or more as providing a reliable estimate of uncertainty provided by James et al (2023) therefore also applies here.

Materials and Stimuli. The stimuli and equipment were similar to those used by Clarke and Hunt (2016; Experiment 1). All stimuli were presented against a uniform grey background (50% white). A row of boxes was presented along the horizontal meridian, each with an edge length of 1.0° visual angle and lighter than the background (75% white). The target was a small, light grey dot (80% white) that appeared for 500 ms either in the top (dot-up configuration) or the bottom (dot-down configuration) of one of the boxes. The 'page-up' and 'page-down' keys of a standard keyboard were used as response keys, 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 R2009b on a PowerMac running Apple Macintosh Pro OS X (version 10.6.4) software using Psychtoolbox (Brainard, 1997; Pelli, 1997) and EyeLinkToolbox functions (Cornelissen et al., 2002). Stimuli were presented on a 24-inch Sony Trimaster EL OLED monitor with a resolution of 1,080p and a refresh rate of 60 Hz. A chin rest with forehead bar assured a viewing distance of 54 cm. The right eye was tracked using a desktop-mounted EyeLink 1000 (version 4.594; SR Research Ltd, Mississauga, ON, Canada) to record eye position at 1,000 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 5 times in a row, or 10 times cumulatively for each block of trials. Following standard procedures for eye-tracking experiments, the experiment took place in a dimly-lit, quiet room.

Procedure and Design. The experiment took place over two sessions. Session 1 lasted approximately 1 hr 30 min, and Session 2 lasted approximately 1 hr. There were three phases (see Figure 1): 'An Acuity Mapping Phase', a 'Directed Eye Movement Phase', and a 'Decision Phase'. Participants were randomly assigned to one of two groups: the control group completed the 'Acuity Mapping Phase' and the 'Decision Phase' in Session 1, and the primed group, called the *instructed* group in this experiment, completed the 'Acuity Mapping Phase' and the 'Directed Eye Movement Phase' in Session 1. In Session 2, both groups completed the same 'Decision Phase'. In all phases, the task was to indicate the target's location in the box (upper or lower) 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 acuity 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). The three phases are described below.

Acuity Mapping Phase. During 4 blocks of 96 trials each (384 trials total), 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. The order of distances was randomised. 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

tant, 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 500ms 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). The purpose of this phase was to fit a logistic regression line to each individual's acuity over distance, and use that to estimate the distance from fixation at which each person's accuracy falls below 75%. This distance is the switch point (E_0). We use this in the decision phase, as described below, to set the locations of boxes for each participant relative to their own E_0 .

Decision Phase. The Decision phase consisted of four blocks. During each block, boxes were presented in nine 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 seven distances ($E_0 [-3^\circ, -2^\circ, -1^\circ, \pm 0^\circ, +1^\circ, +2^\circ, +3^\circ]$). Each distance was presented 10 times per block (90 trials per block, total 360 trials). Trial order was randomised. 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 700ms of stable fixation, three boxes were presented along the horizontal meridian (see Figure 1, bottom row, for an illustration). 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 circle was presented for 500ms in either the left or the right box (never the centre).

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 400ms and then reappeared) and participants were instructed to execute a saccade to this box and remain fixated there until the target circle 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.

Analysis. The choice of which box to fixate was analysed using a similar approach to the one used in Experiment 1. The main difference is that whereas in Experiment 1 participants were free to stand at any location of their choosing, their behaviour in Experiment 2 was constrained so that they could only fixate one of the three boxes. This allows us to simplify our analysis and use a binomial general linear model to describe participants' choices, rather than a hurdle-lognormal model. Another difference between experiments is in the number of separations tested: in the previous experiments there were only three hoop distances in the decision phase, which were treated as categorical. In this experiment, the boxes are presented with nine different separations, which were adjusted to each participant's visual acuity, so we treat this independent variable as continuous. As before, we use a multi-level model. Note that we use a less conservative prior for the slopes, as we expect a near-vertical slope when participants are correctly following the ideal decision strategy.

The analysis we report deviates from the pre-registered analysis, which specified using one-sided *t*-tests to compare the deviation of accuracy from optimal accuracy between the two groups. Instead of this frequentist approach, we adopted a Bayesian analysis to bring it in line with the analysis in Experiment 1. There are no contradictions between the conclusions we draw from these results relative to a frequentist approach, and we have reported the pre-registered tests and plots in the Supplemental Appendix for full transparency.

Results

The Bayesian model's posterior is summarised and compared against the empirical data in Figure 7 (*left two plots*). In Block 1, half the participants were cued to the correct location (the instruction condition), and all the participants in that group successfully followed the instructions (as can be seen in the red lines). The estimate for the *no instruction* condition is extremely wide, reflecting the high degree of variance between participants when they are left to make their own decisions. Consequently, there is a very large difference in accuracy between the two groups in Block 1, as can be seen in the plot on the right of Figure 7. This demonstrates that executing the optimal strategy has large benefits for performance.

In Block 2, in which neither group was provided with guidance telling them which box to choose to fixate, there is no evidence that the participants who were presented with instructions in Block 1 are continuing to implement the optimal strategy. The model fits for Block 2 show that




Figure 7. (Left two plots): Each line represents a different participant's proportion of fixations on the central box as a function of target separation. The coloured bands indicate the 90% HPDI for the fixed effect component of our generalised linear model. Note that in Block 1, there is a red (instruction condition) band in this plot, but the model estimate is very certain (narrow), making it hard to see under the individual lines. (Middle): The model's posterior distribution of the slopes (estimating the effect of target separation on fixation choices) in Block 2 of the experiment. There is no evidence they differ from 0 except when we instructed the optimal strategy in Block 1. (Right): Boxplots illustrating the difference in overall accuracy between conditions and blocks.

both groups of participants exhibit a high degree of variance. We can examine the size of the difference between the two groups in Block 2 by examining the model's posterior density function for the effect of target separation as shown in Figure 7 (third plot). The 95% HPDIs for the instructed group is $[-0.04, 0.10]$, while the no-instruction group gives $[-0.12, 0.03]$, and the posterior difference is $[-0.03, 0.18]$. As all these intervals cross 0, we can conclude there is no overall effect of box separation on fixation choices in Block 2, and no increase in the size of the effect of box separation in Block 2 among the group who were guided to adjust their choices with box separation in Block 1. This failure to benefit from instruction is also reflected in discrimination accuracies (fourth plot): the accuracy gains that come from instruction in Block 1 are eliminated in Block 2.

Discussion

As has been observed in the detection version of the focus-divide dilemma previously (Clarke & Hunt, 2016; James et al., 2019; Morvan & Maloney, 2012), participants show a large degree of variation in their choices of whether to fixate the centre or side boxes in this task, and none appeared to be using the distance between the boxes to decide where to fixate, even after extensive practice. When we instructed participants where to look (cueing them to look at the central box at close distances, and at one of the side boxes when the boxes were too far into their peripheral vision to discriminate the target easily from centre), these instructions conferred considerable benefits to discrimination accuracy. Importantly, however, the participants who had been guided to make optimal fixation choices in the first block returned to highly variable and sub-optimal fixation choices in the second block and looked indistinguishable from the participants who did an

equivalent number of trials but with no instructions. The results are consistent with Experiment 1 in suggesting we have not been successful in nudging participants away from their default of highly variable and sub-optimal choice behaviour, despite giving them direct experience with a far more successful alternative.

Experiment 3: Optimal Choices in Non-Naive Participants

So far in this series of experiments, we have given participants direct experience with solving the focus-divide dilemma, with the expectation that this would facilitate at least some of these participants to continue making optimal choices in subsequent encounters with the same problem. Across the preceding four experiments, we found no evidence of such facilitation. We had assumed that it would be simple for participants to implement the optimal solution once they had acquired some experience with it, but these results led us to question that assumption; perhaps there is a barrier or complexity to implementing the optimal solution that we had not considered. We therefore ran a replication of the throwing and detection tasks on a small group of non-naive participants. This included three of the authors, so that we could experience and report on making these choices from a first-hand perspective.

In the throwing task (3A), participants made a choice about where to stand to throw a beanbag into one of two possible hoops, similar to Experiments 1A, B, and C above. They are only told which hoop is their target after they have made their choice. The detection task (3B) is similar in structure to Experiment 2: participants make a choice about where to look to detect a probe that can appear in one of two boxes. In both tasks, the distance between the possible target locations (hoops/boxes) is varied, and the logic is the same: When they

are close together, stand/look in a central location, equidistant between the two potential target locations; when they are too far apart to throw to/see reliably from a central position, stand/look close to one potential target location. The results demonstrate that the optimal strategy can be easily and effectively implemented when it is explicitly known.

Methods

Participants

Experiment 3A: Throwing. Four participants (three females, all right-handed) completed the throwing task. Three are authors on this paper (AH, WJ, JR), and the final participant was a member of the lab, familiar with the paradigm but not involved in any project making use of it.

Experiment 3B: Detection. Five participants (four females, all right-handed) completed the detection task. Three were authors on this paper (AH, WJ, JR), and the other two participants were lab members, familiar with the paradigm. The authors had run themselves through at least some trials of the detection task before in the process of developing the experiment. The two other participants knew of the strategy but had not done the experiment before.

All participants were aware of the optimal strategy, had normal or corrected-to-normal vision, and were unaware of their personal switch-point. All participants provided informed consent.

Justification of the Sample. As the goal was to verify for ourselves that the solution was easy to implement, this was an appropriate sample. We included non-authors to be reassured that our perception that it was simple was not due to our extensive experience with it.

Materials and Procedure

Experiment 3A: Throwing. The procedure was similar to that used in Clarke and Hunt (2016, Experiment 2), and the throwing task in Experiment 1, with the following modifications. A first phase was conducted, in which flat hoops with a diameter of 0.40 m were placed at a range of distances away {1.38, 3.22, 4.14, 5.06, 6.90, and 9.20 m}. Participants threw 12 bean bags into a hoop at each distance in each of 2 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%.

In the second phase, participants were again asked to throw beanbags 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. Six distances were randomly selected from a range of distances based on the performance of participants in previous experiments (0.46–11.5 m). The random selection of distances within this range ensured that

participants would be unable to rely on their prior 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 colour 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 3 new separations for a further 45 trials. The colour of the beanbag 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 colour 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 3B: Detection. The participants completed a procedure that matched Session 1 of the control group as described in Experiment 2 above. In other words, they completed the Acuity Mapping Phase (384 trials) followed by the Decision Phase (360 trials). The distances between boxes, like in the throwing experiment, were randomly selected from the range of distances used in previous experiments to ensure the non-naive participants would have to rely on their own visual acuity to decide where to fixate, rather than their knowledge of how the experiment had been set up for naive participants.

Results

Choices. We first visually compared each participant's choice behaviour to an individualised estimate of their optimal strategy (i.e. 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: we fit a logistic regression line to performance over distance to estimate where each person's accuracy falls below 50% (throwing) and 75% (detection). This point is E_0 ; the location of the step in the function shown in Figure 8. For the throwing experiment, the participant should choose to stand in the centre for distances where accuracy from a central position is expected to be greater than 50%, and should stand near one hoop for distances where accuracy from centre is less than




Figure 8. Results for each expert participant in the throwing experiment (3A) on the top row. Each dot is light grey and represents the 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 normalised to the distance from the centre to the hoop. The second row shows the results for each expert participant in the detection experiment (3b). Each dot represents the proportion of fixations made towards one of the side boxes at each separation between fixation targets (delta). The dark blue line represents the optimal fixation strategy for each participant.

50%. The blue line shows this optimal performance, and black circles show actual standing position. It is clear from this figure that these expert participants made standing position choices that were close to optimal.

Similarly, for the detection experiment, the participant should choose to look at the centre box when the separation between boxes is small enough that expected accuracy from the centre is $>75\%$. For expected accuracy from the centre 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 8. 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 perfectly followed the blue lines in Figure 8). The size of the difference between actual proportion correct and each participant's optimal performance ranges from underperforming by 0.089 to outperforming by 0.081 . Overall, the mean difference from optimal was -0.017 for the throwing task, and 0.036 for the detection

task, suggesting participants overall were close to optimal in their performance (Figure 9).

Discussion

The successful implementation of the optimal strategy by non-naïve participants in both tasks allows us to more confidently explain the results of Experiments 1 and 2 as a failure to pick up the decision rule, rather than the rule being known but difficult to implement. The results also provide a behaviour-based upper bound on optimal choices in these two tasks. Although participants' choices were broadly consistent with a near-optimal strategy, 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 equalled the accuracy expected under an optimal strategy. The results demonstrate that implementing the optimal strategy is straightforward when it is known.

General Discussion

Experiments 1 and 2 clearly demonstrate that guiding naive participants to perform optimally did not help them




Figure 9. 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.

to make better decisions when the guidance was removed. Experiment 3 demonstrates that optimal performance is straightforward to attain when the participant has an explicit understanding of the domain-general principle that should guide the decision. Even though the decision rule is easy to understand and follow and would boost performance considerably, people do not appear to spontaneously construct or exploit this rule, nor do they apply it more broadly even when they have just implemented it in another context (in Experiment 1) or have experienced the benefits of using it (in Experiment 2). Instead, participants appear to default to relatively variable, sub-optimal patterns of responding.

The experiments were designed to reveal how readily participants recognise and adopt rational decision rules in a series of related choices. Based on the results, we conclude that a variable trial-and-error approach to making choices is a common and persistent tendency. While most participants had some degree of bias towards one choice option or another, there was also a large amount of within-participant, within-condition variability. A small minority of participants responded the same way on nearly every trial; for example, 4/32 participants in Experiment 1A always stood in the centre to throw, and 2 out of 24 participants in Experiment 2 always chose to fixate the centre box. This could reflect a fixed decision rule, but not one which maximises success in the task. Effort minimisation (Irons & Leber, 2016), stress (Shors & Dryver, 1992), and distraction

(e.g. Wolford et al., 2004) could contribute to rigid or stereotyped behaviour. The rigid sub-optimal 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 study by de Wit et al. (2018), however, failed to induce habits in humans under similar protocols to those used previously in rats. Humans 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 far more common 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.

The results of Experiment 2 are at odds with previous research showing that when participants make motor movements that are aligned with the solution to a problem, this can guide them to solve the problem correctly (Alibali et al., 2011; Grant & Spivey, 2003; Thomas & Lleras, 2009). These previous studies suggest motor behaviour can translate to abstract thought to promote solution-finding in difficult problems. On the other hand, our results are consistent with much of the literature on the difficulties of achieving *analogical transfer*, which is the term Gick and Holyoak (1983) coined to describe the idea that participants

might be able to extract a general schema for solving problems from a specific concrete instance of that problem, and then use that schema to solve future instances of that problem. Analogical transfer was not spontaneously achieved by their participants under a range of different conditions with increasingly explicit directions about how to use the solution to one problem to solve a second (although being prompted to describe the similarity between two problems facilitated transfer to a third). This resistance to transfer is echoed in perceptual learning (e.g. Crist et al., 2001). As with analogical transfer, conditions can be carefully engineered to produce transfer (e.g. Ahissar & Hochstein, 1997) but spontaneous transfer of learning across even small changes in context is not the default. The results from Experiment 2 suggesting that participants fail to transfer a successful solution to the focus-divide dilemma across contexts align with this general theme. We can add to this from our experiments the observation that participants default instead to a highly variable and exploratory pattern of decision making.

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 recognised 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 action 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. For example, we move our eyes around three 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 that precise moment. Solving this problem optimally, although possible, requires complex computations even in a highly simplified and predictable environment (Najemnik & Geisler, 2008). In a complex and unpredictable environment, a more efficient approach to these kinds of 'small' problems may be to solve them with *variability* (Krechevsky, 1937a, 1937b), 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 et al., 2016,

2017). A similar process of random selecting from a population of possible responses may guide other forms of decision, preventing stereotyped behaviour while avoiding over-thinking of minor choices.

Applying a consistent decision rule based on knowledge will restrict choice variability. Restricting this variability may stunt 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. 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, participants make variable but sub-optimal choices over the series of trials, consistent with a trial-and-error approach. The failure to spontaneously adopt or maintain an optimal strategy suggests this trial-and-error approach is dominant and persistent. Although this explanation for our results is speculative, a similar tradeoff between learning and logic has been previously proposed in the context of *probability matching* (Gaissmaier & Schooler, 2008). In probability matching, participants are asked to repeatedly guess which of two events will occur on each trial. Participants tend to match the ratio of their guesses to the probability of the events, so for example if the probability of event A is .8 and the probability of event B is .2, they will guess event A on about 80% of trials. Of course, this is a sub-optimal strategy for maximising guessing accuracy, resulting in an overall success rate (in this example) of 0.68, compared to the 0.8 participants could attain if they guessed the event with the higher probability on every trial. Probability matching has been interpreted as a cognitive limitation (e.g. West & Stanovich, 2003). However, Gaissmaier and Schooler (2008) noted that maximising 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-matched (in a random sequence block) were more likely to detect and exploit a pattern in the sequence when one was introduced in a separate block. This is consistent with the notion that applying a constant rule (such as always guessing the more likely event) can optimise 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, truly 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.

Some choice variability is also likely to arise 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 emphasise 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 maximising potential gains and minimising risk and energy expenditure in an environment that, over the long term, is stable. Many so-called ‘sub-optimal’ choices, like probability matching (Edwards, 1956), have been suggested to result from trying to predict short-term, idiosyncratic variations in local context (e.g. Gao & Corter, 2015). Similarly, opting for larger reward later over a smaller one now is only optimal if the environment is stable. However, environments are rarely stable. Incorrectly assuming they are stable could lead to negative outcomes, such as giving up a short-term sure reward in favour of a longer-term uncertain reward that never comes to fruition. Kolling et al. (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 introduction presented the focus-divide dilemma as a routine problem we resolve in daily life, but our experiments are a simplified version of the real-world situations in which people make decisions about how to allocate their resources between competing goals. The competing tasks and goals of daily life are rarely matched in their difficulty or other features, and the resource requirements and probability of successful outcomes are often difficult to estimate accurately. In the experiments presented here, we have simplified the problem to make it as predictable and achievable as possible, with the assumption that if participants can’t solve the problem in its simplest possible form, it is unlikely they are solving it in more complex forms. But simpler versions of problems do not always provoke better decisions than complex ones. One example of this comes from Nowakowska et al. (2024), who showed that search strategies for simple visual features were sub-optimal and heterogeneous, but the same search problem presented with complex objects instead of simple features provoked uniformly optimal strategies. It is possible that simplifying the focus-divide dilemma removes features that provide cues to better choices, leading to an overly pessimistic prediction about how well participants would resolve these dilemmas when they encounter them in vivo. It would be interesting to

find or engineer more familiar or high-stakes problems (e.g. in financial investment or resource triage) to understand the potential impacts of failures to resolve the focus-divide dilemma in practice. The current results suggest that limitations of this kind will be resistant to short-term learning or interventions, and that overcoming them in one context is unlikely to benefit other circumstances.

Conclusion

Choices are governed by a complex set of individual and contextual factors and their interaction, even for seemingly simple decisions. In a focus-divide dilemma, we confirm that persistently variable and suboptimal choices are typical, despite the existence of a simple and intuitive decision rule with clear performance benefits. In these experiments, we have created conditions in which participants directly engaged in using this optimal decision rule, and were then returned to the standard focus-divide task context and allowed to freely choose. Participants returned to making variable and idiosyncratic decisions, and compared to control groups with equivalent practice, there was no discernible change in choice strategy. Given how universal and persistent this behavioural tendency is, the benefits of variability in choice may, in many circumstances, outweigh the costs of finding and applying consistent decision rules.

Acknowledgements

The authors would like to thank Jessie Grimsey and Michelle Tran for collecting data for Experiment 1b; Anca Popescu for collecting data for Experiment 1c; and Eleanor Mackay for collecting the data for Experiment 2. We would also like to thank Justin Claydon, Alex Anwyl-Irvine, Ellen Angus, Manjiri Bhat, Olga Amato, and Anna Nowakowska for their contributions. This research was made possible through support from the James S. McDonnell Foundation (Scholar Award to ARH).

ORCID iDs

Amelia R. Hunt  <https://orcid.org/0000-0002-4122-9499>
 Warren R. G. James  <https://orcid.org/0000-0002-4163-462X>
 Alasdair D. F. Clarke  <https://orcid.org/0000-0002-7368-2351>

Funding

The authors received no financial support for the research, authorship, and/or publication of this article.

Declaration of Conflicting Interests

The authors declared no potential conflicts of interest with respect to the research, authorship, and/or publication of this article.

Supplemental Material

Supplemental material for this article is available online.

References

- Ahissar, M., & Hochstein, S. (1997). Task difficulty and the specificity of perceptual learning. *Nature*, 387(6631), 401–406. <https://doi.org/10.1038/387401a0>
- Alibali, M. W., Spencer, R. C., Knox, L., & Kita, S. (2011). Spontaneous gestures influence strategy choices in problem solving. *Psychological Science*, 22, 1138–1144. <https://doi.org/10.1177/0956797611417722>
- Brainard, D. H. (1997). The psychophysics toolbox. *Spatial Vision*, 10(4), 433–436. <https://doi.org/10.1163/156856897x00357>
- Carr, H. (1925). *Psychology: A study of mental activity*. Longmans, Green.
- 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. <https://doi.org/10.1167/16.7.4>
- 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. <https://doi.org/10.1177/0956797615611933>
- 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. <https://doi.org/10.1167/17.11.12>
- Claydon, J., James, W. R., Clarke, A. D., & Hunt, A. R. (2024). The role of framing, agency and uncertainty in a focus-divide dilemma. *Memory & Cognition*, 52(3), 574–594. <https://doi.org/10.3758/s13421-023-01484-6>
- Cornelissen, F. W., Peters, E. M., & Palmer, J. (2002). The Eyelink Toolbox: Eye tracking with MATLAB and the Psychophysics Toolbox. *Behavior Research Methods, Instruments, & Computers*, 34(4), 613–617. <https://doi.org/10.3758/bf03195489>
- Crist, R. E., Li, W., & Gilbert, C. D. (2001). Learning to see: Experience and attention in primary visual cortex. *Nature Neuroscience*, 4(5), 519–525. <https://doi.org/10.1038/87470>
- de Wit, S., Kindt, M., Knot, S. L., Verhoeven, A. A., Robbins, T. W., Gasull-Camos, J., Evans, M., Mirza, H., & Gillan, C. M. (2018). Shifting the balance between goals and habits: Five failures in experimental habit induction. *Journal of Experimental Psychology: General*, 147(7), 1043. <https://doi.org/10.1037/xge0000402>
- Dickinson, A. (1985). Actions and habits: The development of behavioural autonomy. *Philosophical Transactions of the Royal Society B*, 308, 67–78. <https://doi.org/10.1098/rstb.1985.0010>
- Edwards, W. (1956). Reward probability, amount, and information as determiners of sequential two-alternative decision. *Journal of Experimental Psychology*, 52, 177–188. <https://doi.org/10.1037/h0047727>
- Gaissmaier, W., & Schooler, L. J. (2008). The smart potential behind probability matching. *Cognition*, 109, 416–422. <https://doi.org/10.1016/j.cognition.2008.09.007>
- Gao, J., & Corter, J. E. (2015). Striving for perfection and falling short: The influence of goals on probability matching. *Memory & Cognition*, 43, 748–759. <https://doi.org/10.3758/s13421-014-0500-4>
- Gick, M. L., & Holyoak, K. J. (1983). Schema induction and analogical transfer. *Cognitive Psychology*, 15(1), 1–38. [https://doi.org/10.1016/0010-0285\(83\)90002-6](https://doi.org/10.1016/0010-0285(83)90002-6)
- Grant, E. R., & Spivey, M. J. (2003). Eye movements and problem solving: Guiding attention guides thought. *Psychological Science*, 14, 462–466. <https://doi.org/10.1111/1467-9280.02454>
- Hesse, C., Kangur, K., & Hunt, A. R. (2020). Decision-making in slow and rapid reaching: Sacrificing success to minimize effort. *Cognition*, 205, Article 104226. <https://doi.org/10.1016/j.cognition.2020.104426>
- Irons, J. L., & Leber, A. B. (2016). Choosing attentional control settings in a dynamically changing environment. *Attention, Perception, & Psychophysics*, 78, 2031–2048. <https://doi.org/10.3758/s13414-016-1125-4>
- James, W. R. G., Clarke, A. D. F., & Hunt, A. R. (2017). *Failure to use probability of success when deciding whether to pursue one goal or two* [Conference session]. Proceedings of the 39th annual meeting of the Cognitive Science Society. Cogsci, pp. 2279–2283.
- James, W. R. G., Hunt, A. R., & Clarke, A. D. F. (2023). Six of one, half dozen of the other: Suboptimal prioritizing for equal and unequal alternatives. *Memory and Cognition*, 51, 486–503. <https://doi.org/10.3758/s13421-022-01356-5>
- James, W. R. G., Reuther, J., Angus, E., Clarke, A. D. F., & Hunt, A. R. (2019). Inefficient eye movements: Gamification improves task execution, but not fixation strategy. *Vision*, 3(3), 48. <https://doi.org/10.3390/vision3030048>
- Kahneman, D. (2011). *Thinking, fast and slow*. Macmillan.
- Kolling, N., Wittman, M., & Rushworth, M. F. S. (2014). Multiple neural mechanisms of decision making and their competition under changing risk pressure. *Neuron*, 81, 1190–1202. <https://doi.org/10.1016/j.neuron.2014.01.033>
- Krechevsky, I. (1937a). Brain mechanisms and variability I: Variability within a means-end-readiness. *Journal of Comparative Psychology*, 23, 121–138. <https://doi.org/10.1037/h0061929>
- Krechevsky, I. (1937b). Brain mechanisms and variability II: Variability where no learning is involved. *Journal of Comparative Psychology*, 23, 139–163. <https://doi.org/10.1037/h0056131>
- Maier, N. R. F. (1940). The behavior mechanisms concerned with problem solving. *Psychological Review*, 47, 43–58. <https://doi.org/10.1037/h0058466>
- Markovits, H., & Savary, F. (1992). Pragmatic schemas and the selection task: To reason or not to reason. *The Quarterly Journal of Experimental Psychology*, 45(1), 133–148. <https://doi.org/10.1080/14640749208401319>
- Morvan, C., & Maloney, L. T. (2012). Human visual search does not maximize the post-saccadic probability of identifying targets. *PLOS Computational Biology*, 8(2), Article e1002342. <https://doi.org/10.1371/journal.pcbi.1002342>
- Najemnik, J., & Geisler, W. S. (2008). Eye movement statistics in humans are consistent with an optimal search strategy. *Journal of Vision*, 8(3), 4. <https://doi.org/10.1167/8.3.4>
- Nowakowska, A., Clarke, A. D., Reuther, J., & Hunt, A. R. (2024). Variable search for orientation, uniformly optimal search for identity. *Journal of Experimental Psychology: General*, 153(2), 495. <https://doi.org/10.1037/xge0001509>
- O'Doherty, J. P., Cockburn, J., & Pauli, W. M. (2017). Learning, reward, and decision making. *Annual Review of Psychology*, 68, 73–100. <https://doi.org/10.1146/annurev-psych-010416-044216>

- Pelli, D. G. (1997). The VideoToolbox software for visual psychophysics: Transforming numbers into movies. *Spatial Vision*, 10, 437–442. <https://doi.org/10.1163/156856897x00366>
- 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. <https://doi.org/10.3758/bf03332056>
- Sloman, S. A. (1996). The empirical case for two systems of reasoning. *Psychological Bulletin*, 119(1), 3.
- Stanovich, K. E., & West, R. F. (2000). Advancing the rationality debate. *Behavioral and Brain Sciences*, 23(5), 701–717. <https://doi.org/10.1017/s0140525x00623439>
- Summerfield, C., & Tsetsos, K. (2015). Do humans make good decisions? *Trends in Cognitive Sciences*, 19, 27–34. <https://doi.org/10.1016/j.tics.2014.11.005>
- Thomas, L., & Lleras, A. (2009). Swinging into thought: Directed movement guides insight problem solving. *Psychonomic Bulletin and Review*, 16, 719–723. <https://doi.org/10.3758/pbr.16.4.719>
- Tulving, E., & Thomson, D. M. (1973). Encoding specificity and retrieval processes in episodic memory. *Psychological Review*, 80(5), 352. <https://doi.org/10.1037/h0020071>
- 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. <https://doi.org/10.3758/bf03194383>
- Wickham, H., Averick, M., Bryan, J., Chang, W., McGowan, L. D. A., François, R., Grolemund, G., Hayes, A., Henry, L., Hester, J., Kuhn, M., Pedersen, T. L., Miller, E., Bache, S. M., Müller, K., Ooms, J., Robinson, D., Seidel, D. P., Spinu, V., . . . Yutani, H. (2019). Welcome to the Tidyverse. *Journal of Open Source Software*, 4(43), 1686. <https://doi.org/10.21105/joss.01686>
- Wilson, R. C., Geana, A., White, J. M., Ludvig, E. A., & Cohen, J. D. (2014). Humans use directed and random exploration to solve the explore–exploit dilemma. *Journal of Experimental Psychology: General*, 143(6), 2074–2081. <https://doi.org/10.1037/a0038199>
- Wolford, G., Newman, S., Miller, M. B., & Wig, G. (2004). Searching for patterns in random sequences. *Canadian Journal of Experimental Psychology*, 58, 221–228. <https://doi.org/10.1037/h0087446>