Decision making in slow and rapid reaching:
Sacrificing success to minimize effort

Constanze Hesse[1*], Karina Kangur[1], & Amelia R. Hunt[1]

06/08/2020

RUNNING TITLE: "Decision making in slow and rapid reaching"

*Correspondence should be addressed to:

Dr. Constanze Hesse

School of Psychology

University of Aberdeen

King's Campus

AB24 3FX

United Kingdom

e-mail: c.hesse@abdn.ac.uk

phone: +44 (0)1224 273215

---

[1] School of Psychology, University of Aberdeen, Aberdeen, UK

Abstract

Current studies on visuomotor decision making come to inconsistent conclusions regarding the optimality with which these decisions are made. When executing rapid reaching movements under uncertainty, humans tend to automatically select optimal movement paths that take into account the position of all potential targets (spatial averaging). In contrast, humans rarely employ optimal strategies when making decisions on whether to pursue two action goals simultaneously or prioritise one goal over another. Here, we manipulated whether spatial averaging or pre-selection of a single target would provide the optimal strategy by varying the spatial separation between two potential movement targets as well as the time available for movement execution. In Experiment 1, we aimed to determine the time needed to reach for targets with small and large separation between them and to measure baseline strategies under low time pressure. Given generous time limits, participants did not employ a pure averaging approach but instead tended to pre-select the target that was easiest to reach and corrected their movement path in-flight if required. In Experiment 2, a strict time limit was set such that the optimal strategy to reach the correct target successfully depended on the separation between the potential targets: for small separations, there was enough time to employ averaging strategies, but higher success for larger separations required pre-selecting the final target instead. While participants varied in the strategies they preferred, none of them flexibly adjusted their movement strategies depending on the spatial separation of the targets. In Experiment 3, we confirm the bias towards targets that are easiest to reach and show that this comes at the expense of overall task success. The results suggest a strong tendency for humans to minimize immediate movement effort and a general failure to adapt movement strategies flexibly with changes in the task parameters.

1. **<u>Introduction</u>**

Imagine you are a goal keeper during a penalty kick. The decisions you have to make are complex and numerous. Before the kick, you need to choose a place to stand. Should you stand in the centre of the net or off-centre? Should you stand closer to the kicker, or inside the net? Should you start your dive before the kick and risk guessing the wrong direction, or wait to react to the kick and risk responding too slowly? When diving, which foot do you lead with? Will you turn your body towards the ball to get more distance, or face forward to create a larger blocking surface? Do you try to grab the ball, smack it away, or bounce it off another body part? The goal keeper's many dilemmas illustrate the cascade of motor decisions taking place at many different levels of deliberation and control. Some decisions can be epistemic, such as where to stand before the kick: these decisions can take into account explicit, conscious knowledge about how far you can dive in each direction and what you know about the opposing player's skills and history to increase your odds of success. Other decisions, like which foot to lead with when you start your movement, may automatically follow from an unconscious estimate of how to get your body to intersect with the ball, based on long hours of training.

Researchers have long been interested in both aspects: the efficiency of motor control and the optimality of human decision making. However, these domains have largely been investigated separately from each other, and are traditionally viewed as belonging to different fields (i.e. motor control vs. cognitive psychology). To understand how the goalkeeper's dive is executed, for example, researchers could measure the speed and trajectory of blocking movements in response to different kicks to understand visuomotor constraints on performance. In contrast, research on strategic decisions, such as where to stand before the kick, measure deviations from optimal choices to understand the cognitive biases and heuristics that limit our success. Yet, it is clear from our goal keeper example that decisions exist on a continuum, and strong interconnections must exist between the cognitive systems that drive strategic decisions and sensorimotor control. During the last decade there has been a considerable increase in research attempting to bridge the gap between cognition and sensorimotor systems (for review see Gallivan, Chapman, Wolpert, & Flanagan, 2018; Wispinski, Gallivan, & Chapman, 2018).

One strand of studies that has provided initial evidence for the view that our action execution is strongly influenced by cognitive processes is based on the investigation of eye and hand movements in the presence of visual distractors (e.g., Findlay & Walker, 1999; Nakayama, Song, Finkbeiner, & Caramazza, 2007; Song & Nakayama, 2008, 2009; Van der Stigchel, Meeter, & Theeuwes, 2006; Walker, McSorley, & Haggard, 2006). These studies demonstrated that the shape of the selected movement path toward a target is affected by the presence of distractor

objects and their specific properties, suggesting that the visuomotor system initially represents both targets and distractors as potential movement locations. This demonstrates that the spatial properties of movement trajectories can provide insights into the underlying target selection process and can therefore, more generally speaking, reveal the internal cognitive processes (such as attention and decision making) that underlie our actions.

Based on these findings, numerous studies have investigated eye and hand movements in the presence of multiple potential targets and have confirmed that the motor system seems to prepare for different possible actions simultaneously, resulting in spatial changes of the observed trajectory to a final movement target. These spatial changes in the selected movement path have been considered to reflect the unsuccessful suppression of actions that were planned prior to execution but that were not called for (Cisek, 2012; Cisek & Kalaska, 2010; Kable & Glimcher, 2009). The notion that multiple potential movements are prepared by the CNS in parallel and compete for execution has been coined the "affordance competition hypothesis" by Cisek and Kalaska (2010). The behavioural effects of the competition between multiple action possibilities on subsequent movements to a target object are probably most effectively illustrated in so-called "go before you know" tasks (e.g., Chapman et al., 2010a; Gallivan & Chapman, 2014; Gallivan et al., 2011; Stewart, Gallivan, Baugh, & Flanagan, 2014). In these tasks, participants are presented with two (or more) potential action targets. They are then asked to reach to a final target as quickly as possible, but the target is only revealed once participants have initiated their movements (e.g., after the release of a start button during reaching). What is usually observed in these tasks is a movement path that is initially directed between the potential movement locations and is corrected in-flight toward the final target as soon as it has been revealed (for review see, Gallivan & Chapman, 2014). In other words, the initial movement is directed toward the average direction of all potential movement targets. This is known as *spatial averaging*.

Currently, there are two different explanations for the observation that the visuomotor system seems to take into account the location of all possible movement targets when planning the initial direction of the movement. The first is based on evidence from neurophysiological studies indicating that neurons in sensorimotor regions represent multiple potential targets and actions before a final decision is made (e.g., Cisek & Kalaska, 2005; Cisek & Kalaska, 2010). In line with the affordance competition hypothesis specified above, it has been suggested that the averaging behaviour is a consequence of the simultaneously specified competing motor plans and can therefore be considered unintentional behaviour (e.g., Gallivan, Logan, Wolpert, & Flanagan, 2016; Stewart, Baugh, Gallivan, & Flanagan, 2013). The second, more recent, explanation is that these averaged movements are actually deliberate and may represent one single optimal movement

plan (e.g., Haith, Huberdeau, & Krakauer, 2015; Hudson, Maloney, & Landy, 2007; Nashed, Diamond, Gallivan, Wolpert, & Flanagan, 2017; Wong & Haith, 2017). In other words, the observed averaging behaviour in "go-before-you-know" tasks may represent a deliberate and optimal strategy that participants implement to deal with the uncertainty of the task (Wong & Haith, 2017). The strategy can be considered optimal as the movement-related costs are minimised by initialising a movement toward the average location of all potential movement targets, which ensures that the required in-flight corrections are minimal once the final target is revealed.

The issue with both explanations is that they are difficult to disentangle as the predicted movement trajectories are identical. A few recent studies have tried to address this issue by experimentally creating situations in which averaging behaviour is no longer beneficial (e.g., Haith et al., 2015; Wong & Haith, 2017). For example, Wong & Haith (2017) showed that participants abandoned averaging behaviour if they had to execute extremely fast movements. That is, averaged movements were found to be produced only at slower speeds, when there was sufficient time to correct them in flight. Similarly, Haith et al. (2015) found no averaging behaviour if the separation between potential targets was very large or a barrier was put between them. Most importantly, however, the observation that a change in high-level task requirements can eliminate averaging behaviour is clearly at odds with the idea that these movements are a result of competing motor plans and are thus automatic and unintentional. Based on these findings, it has been argued that averaged movements in decision tasks can be best understood within an optimal control theory, according to which a single motor plan is specified with the aim of minimising movement costs while at the same time increasing the chances of task success (Haith et al., 2015).

The suggestion that humans show nearly optimal behaviour when presented with a visuomotor decision task presents an interesting contrast with findings demonstrating clearly suboptimal performance in similar, but more deliberative, decision making tasks that require people to specify and select a movement strategy prior to movement onset. This seems to be true for a range of different tasks (Clarke & Hunt, 2016; Morvan & Maloney, 2012; Nowakowska, Clarke, & Hunt, 2017). For example, Clark and Hunt (2016) manipulated the distance between two potential targets in a beanbag throwing task based on participants' baseline throwing performance. Participants' task was to choose a standing position prior to the experimenter revealing which of the two targets participants had to hit with their beanbag. The rationale of this task is simple: If the two potential targets are separated by a short distance, participants should choose to stand midway between them (similar to spatial averaging in reaching) as this would ensure that they could easily hit either target after it is revealed by the experimenter. If the two targets are separated by a large distance, such that standing in the middle would make it difficult

to successfully hit either target (as established through participants' baseline performance), participants should choose a standing position close to one of the potential targets, thus ensuring at least a task success of 50%. However, it was repeatedly found that very few participants ever select this (optimal) strategy with changes in distance between the targets, thereby falling far short of the maximum throwing success they could have achieved. Instead, participants seem to show highly variable standing position choices, resulting in inferior task performance (see Clarke & Hunt, 2016). The same failure to adjust choices with changes in task difficulty was observed in two other task contexts (memorizing digits and detecting targets), suggesting a general tendency that cuts across many contexts, from eye and hand movements to more complex decisions. Thus, in tasks that require participants to select movements in the presence of multiple potential goals there is a clear mismatch between more abstract and deliberative decisions on the one hand, and more implicit sensorimotor decision behaviour on the other, with the former appearing to be far from optimal, and the latter close to optimal.

A similar mismatch between visuomotor and deliberative/cognitive decisions has been revealed and discussed in the visuomotor literature using a paradigm investigating decision making under risk using eye and hand movements (e.g., Jarvstad, Hahn, Warren, & Rushton, 2014; Maloney, Trommershäuser, & Landy, 2007). In this paradigm, participants are asked to reach within a tight time constraint to a target region that incurs a small (usually monetary) reward. What makes this task a motor decision task is that the target can be overlapped with a penalty region that, when touched, incurs a small (monetary) loss. While it would be optimal to hit the centre of the target in the absence of a penalty region, in its presence participants should shift their movement end-point away from the target centre and the penalty region. The beauty of this task lies in the fact that the optimal hitting point (that is, the location that yields the maximum expected gain in points and money) can be precisely determined for each participant using their own motor precision and cost functions. Consequently, participants' performance can be easily compared to that of an optimal decision maker (Trommershäuser, Maloney, & Landy, 2003a, 2003b, 2008). The initial studies consistently found that participants select strategies that are close to optimal when making these speeded and risky visuomotor decisions, in contrast with the sub-optimal behaviour usually observed in traditional economic decision making tasks (Trommershäuser et al., 2008). More recent studies have, however, identified a number of boundary conditions for optimal behaviour to occur, specifically suggesting that consistent feedback and experience are essential (Neyedli & LeBlanc, 2017; Neyedli & Welsh, 2013, 2014, 2015) as well as relatively constant and simple gain landscapes (Jarvstad et al., 2014; Wu, Trommershäuser, Maloney, & Landy, 2006). Based on these findings, it has been argued that the perceptuomotor systems may not be as

superior to higher-level cognitive systems as initially thought when it comes to optimal decision making (Jarvstad et al., 2014; Neyedli & LeBlanc, 2017).

What distinguishes this task from the go-before-you-know tasks outlined previously is the presence of externally imposed payoffs. Neyedli and Welsh (2014) varied both the distance between the target and the penalty region, as well as the associated pay-off values, on a trial-by-trial basis. Interestingly, they found that while participants seem to optimally shift their movement endpoints as the distance between the two circles varied, this was not the case for trial-by-trial variations in the associated penalty values. Based on this finding, they argued that spatial parameters, as an intrinsic property of the visual stimuli, are more tightly linked to the motor system than pay-off parameters, which need to be interpreted by the cognitive system. If spatial parameters are indeed easier for the motor system to optimize, then in go-before-you-know tasks, changes in appropriate task strategies depending on target distance might be more readily applied to optimally meet the demands of the task at hand.

In our current study, we combined the logic of Clarke and Hunts' choice task of pursuing one goal versus two, with a "go-before-you-know" reaching paradigm as employed in previous studies on sensorimotor decision making. Specifically, we varied the distance between the potential movement targets whilst limiting the amount of time available to participants to reach the final pointing target specified after movement onset. We thereby selected the available time to reach the final movement target such that averaging behaviour was feasible for targets separated by a short distance but not for targets separated by a large distance. The primary aim of Experiment 1 was to determine the approximate movement time limit that would allow participants to apply averaging behaviour for near distance targets but not for far distance targets. In addition, this experiment also allowed us to examine reaching behaviour in sensorimotor decision making under quite liberal timing conditions. We then applied a tight response deadline (as determined by Experiment 1) in Experiment 2 to assess whether participants would be able to flexibly adapt their movement strategies on a trial-by-trial basis based on the expected likelihood of reaching the target before the deadline. Participants were told that trials where the wrong target was hit or the correct target was not reached in time would have to be repeated and were regularly reminded how many trials they had successfully completed and how many they had left to do. We defined optimal performance as the strategy that would allow participants to maximize their probability of success

on any single trial, and thereby also to finish the experiment in the smallest possible number of performed trials[1].

We were particularly interested in the types of errors individual participants would make in Experiment 2: specifically, we determined the proportion of trials ending at the wrong target, and the proportion of total trials that failed to reach the target in time. Previous studies usually excluded movement errors from further analysis, which amount to the exclusion of a very large number of trials, and even entire participants who were unable to complete enough trials correctly within the given time-constraints (see Gallivan & Chapman, 2014 for detailed discussion). These exclusions may potentially lead to an incomplete picture of the (unsuccessful) strategies that participants employ when doing the task. For example, one could assume that for speeded movements as required in Wong & Haith's (2017) experiment, only straight movement paths would actually allow participants to reach the targets in time. Therefore, it is unsurprising that the trajectories of the correct movements are straight. To understand how participants solve the problem, it would be necessary to look at the frequency and types of errors they make. Optimal performance in our Experiment 2 requires participants to employ an averaging strategy for targets presented close to each other, but a pre-selection strategy (i.e. straight movement path) for targets separated by a large distance. If participants fail to adapt their strategies to the distance between the targets, we should see consistent use of a single strategy across all distances; this would mean more time-out errors at the far distances (if they over-apply an averaging strategy), and more movements to the wrong target at the close distances (if they over-apply a pre-selection strategy).

By examining error trials, we revealed, similar to Clark & Hunt (2016), that none of our participants (N=20) flexibly adapted their behavioural strategy to the different task demands. However, they did show a strong tendency to minimize effort on a trial-by-trial basis. To confirm these observations, in Experiment 3 we tested participants' decision-making behaviour in a visuomotor choice task by encouraging them, through the wording of our instructions, to select and move to one of the two potential targets, while still reinforcing speed and accuracy by repeating movement errors and time-outs. We found that participants readily adopted a pre-selection strategy and consistently chose targets that were associated with reduced biomechanical effort.

---

[1] This differs from studies investigating decision making under risk where optimal choices are monetarily rewarded and optimality is defined as the maximum gain based on the specific costs associated with the outcome of an action (see also Wolpert & Landy, 2012).

2. Experiment 1 – Decision task without time-pressure

2.1 Methods

2.1.1 Participants

Twelve undergraduate and postgraduate students of the University of Aberdeen (aged between 18 to 28 years, mean age: 20.3 years, 6 males) who were naïve to the purpose of the experiment participated in the study. Undergraduate students received course credits for their participation. They all had normal or corrected to normal visual acuity and were right-handed by self-report. The experiment was approved by the School of Psychology Ethics Committee at the University of Aberdeen and all volunteers provided written informed consent prior to participating.

2.1.2 Stimuli and Apparatus

Participants were seated centrally at a wooden table in a dimly lit lab. Reaching movements were measured using a TrakStar™ electromagnetic motion tracking system (Ascension Technology Corporation, NDI) at a sampling rate of 240 Hz. One marker was attached to the nail of participants' right index finger and secured in place using medical tape. Prior to the experiment the set-up was calibrated such that the start position of the hand, which was placed centrally at the lower edge of the table (see Figure 1), corresponded to the origin of the Cartesian coordinate system.

The stimulus presentation display was a wooden box (L: 122 x W: 9 x H: 6 cm) with a frosted light-diffusing Perspex cover sheet (3 mm thick). The cover could be lifted, and six green LED lights with a diameter of 5 mm could be fixed in place at varying distances along a rail inside the box. The frosted Perspex cover ensured that the LEDs and mountings were not visible to the participant when switched off but that the LED light would be visible through the cover when switched on. The stimulus display was fixed centrally on the table at a distance of 30 cm from the start position of the hand (see Figure 1).

The six LED lights were placed at different locations: 6 cm, 12 cm, and 48 cm to the left (L) and right (R) of the centre of the stimuli display resulting in three different reaching distances of about 30.6 cm, 32.3 cm and 56.6 cm, respectively. The experiment was programmed in Matlab.

2.1.3 Procedure

Participants sat at the table with their right index finger placed at the start position (5x5 mm plastic knob). They were instructed to reach swiftly and as accurately as possible to the target light following an auditory go-signal (100 ms duration at 1000 Hz). In all trials, the go-signal and the

visual target(s) were presented simultaneously. All participants had to perform two different tasks: A baseline task followed by a decision task.

In the baseline task (single target condition), one of the six LEDs lit up at the presentation of the go-signal and participants were instructed to reach to this target. They had 500 ms to initiate their movement after the go-signal and were encouraged to reach the target within 600 ms once the movement had started (see data analysis for more detail). Movements that exceeded the reaction time limit were classified as error and repeated at a later random position within the experiment (10.7% of all trials). Each of the six target locations (06L, 06R, 12L, 12R, 48L, 48R) was presented 12 times in randomised order, resulting in a total of 72 reaching trials. The task was preceded by six practice trials (one to each location in random order) to familiarise participants with the task. At the end of the trial, participants received an auditory warning signal if their movements were slower than 600 ms (100 ms at 400 Hz).
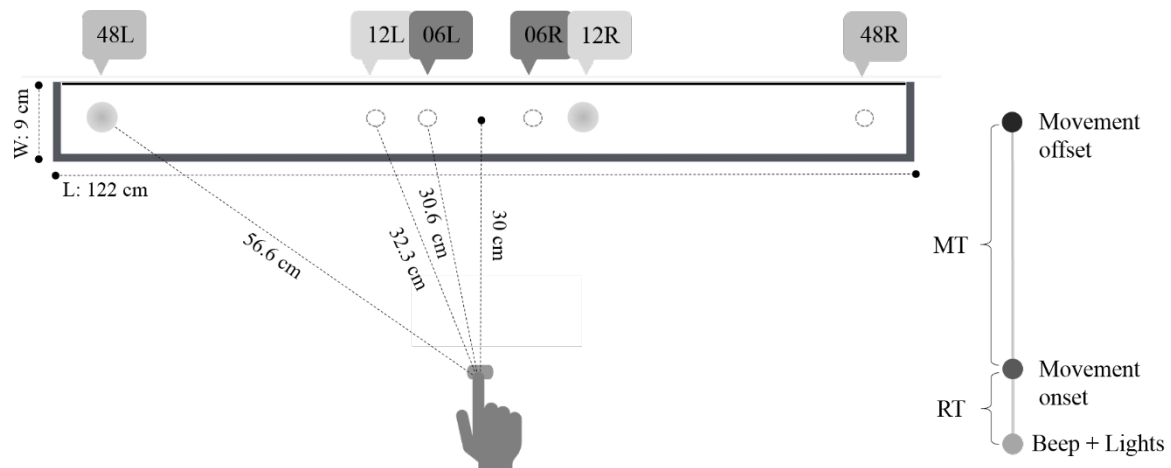
Figure 1: Experimental setup (not drawn to scale) demonstrating the stimulus display and target positions. This example demonstrates a decision trial. In the baseline task, only one LED was turned on at a time.

Following the baseline task, all participants performed the decision task (double-target condition). Here, two potential target positions (always one left and one right of the midline) lit up together with the presentation of the go signal. Once a movement was initiated (i.e. the finger had moved 2.5 cm away from the starting position in 3D space), one of the lights switched off and participants were instructed to point to the light that remained switched on. Light-pair combinations are labelled such that the first digit and letter combination indicated the final target that would stay lit-up (i.e. with the number indicating its distance from midline in cm and letter the side it was presented at) and the second digit-letter combination indicates the light that would be switched off. For example, the 06L-12R and the 12R-6L condition were initially visually identical up to movement initiation, after which the light located 6 cm to the left of the midline

10

would become the target location in the first case, but the light located 12 cm to the right of the midline would be the target location in the second case. There were 18 possible trial combinations in total (06L-06R, 06L-12R, 06L-48R, 12L-06R, 12L-12R, 12L-48R, 48L-06R, 48L-12R, 48L-48R, 06R-06L, 06R-12L, 06R-48L, 12R-06L, 12R-12L, 12R-48L, 48R-06L, 48R-12L, 48R-48L). Each of the 18 light-pair combinations was presented 12 times in random order throughout the experiment, resulting in a total of 216 trials (usually performed in about 40 min). Before the start of the decision task, participants completed 10 random practice trials.

2.1.4 Data analysis

All movement data was derived from the electromagnetic marker placed on participants' index finger. The positional data was filtered offline using a second order dual-pass Butterworth filter with a cut-off frequency of 15 Hz. The resultant velocity was calculated from the filtered position data. Movement onset was defined as the moment the marker exceeded a resultant velocity of 0.05 m/s. Movement offset was defined using a combination of a spatial criterion and a velocity threshold. Specifically, we looked for the first frame at which the resultant velocity dropped below a threshold of 0.05 m/s after the finger had travelled at least 25 cm in the y-direction. Movement time (MT) was defined as the time between movement onset and movement offset.

Trajectories were normalised and processed using the methods suggested by Chapman et al. (2010a) (for a detailed discussion see also Gallivan & Chapman, 2014). The trajectory data was normalised between movement onset and offset relative to reach distance in the y-direction (the distance from the start position to the pointing device in y-coordinates) as this dimension was identical in all movement conditions. Space normalisation was done using functional data analysis (FDA) tools (Ramsay & Silverman, 2005). Adopting the procedures of Chapman et al. (2010), we fitted 6[th] order b-splines to each of the three dimensions (x, y, z) of the extracted reach trajectory with a knot at every data point. From the trajectories, now mathematically defined, we extracted the horizontal location (x-dimension) and times at 200 points now equally spaced along the reaching distance (y-dimension). Trials of the same condition were averaged for each participant and mean trajectories were calculated by subsequently averaging data across all participants. In order to determine when trajectories for initially visually identical conditions differed significantly in the lateral dimension, we ran functional comparisons using two-level repeated-measures functional ANOVAs (fANOVA) and adjusted significance levels using a Bonferoni correction for the number of initially visually identical displays tested (i.e. $p < 0.05/9$). The fANOVA is an adaptation of the traditional ANOVA designed to be specifically applied to continuous data to test if and where functionally defined measures differ. It provides a functional F-statistic indicating

which points differ significantly between two continuous measures across conditions. Note that the functional F-statistic and p-values can only be meaningfully interpreted in their entirety (i.e., trajectory differences in the lateral dimension). Hence, in line with previous studies applying this procedure, we report significant differences by placing significance bars to the side of our trajectories in the figures (e.g., Chapman et al., 2010a; Gallivan & Chapman, 2014).

Finally, to determine and compare to which light participants directed their movement first, we also calculated the hand's initial starting direction in baseline and decision trials. The initial start direction was calculated as the angle in degrees between the direction to the pointing device (y), the hand's starting position, and the position of the hand (measured in x and y) when it had moved 50 mm in the y-direction toward the pointing device (for similar procedure see Brouwer, Middelburg, Smeets, & Brenner, 2003). Previous work has suggested that in-flight trajectory corrections usually only begin 100-150 ms after movement initiation (Brenner & Smeets, 1997; Day & Lyon, 2000; Soechting & Lacquaniti, 1983) and that hence, the first 100 ms of the movement (usually corresponding to about 70 mm of reaching distance) reliably reflect the effects of cue-processing before any corrections are implemented (see Gallivan & Chapman, 2014 for an in-depth discussion).

2.2 Results and Discussion

2.2.1 Baseline Task

This task was employed to determine the average time participants needed to reach the targets presented at the three different distances at the left and right side of body midline as well as to determine the shape of the movement paths when movements are planned and executed to a single target. Figure 2 shows the median MTs in each condition averaged across participants as well as the median MT for each participant separately. Note that 15 out 864 baseline trials (1.7%) had to be excluded from this analysis as the end of movement could not be determined reliably using our criteria. Descriptively, MTs were, as expected, longer for targets presented on the left-hand side of the display as compared to the right, and longer for targets presented at the largest eccentricity. The effect of eccentricity was particularly pronounced when movements had to be executed toward the left, thus requiring our right-handed participants to move their arm and hand across the body. These observations were confirmed statistically using a 3 (distance of LED) x 2 (side of LED) repeated measures ANOVA. This revealed significant main effects of both LED side, $F(1,11)=39.06$, $p<.001$, $\eta_p^2=.78$, and LED distance, $F(2,22)=104.34$, $p<.001$, $\eta_p^2=.91$, as well as a significant interaction effect, $F(2,22)=19.15$, $p<.001$, $\eta_p^2=.64$. The figure also demonstrates that participants were able to reach all target lights within the time limit of 600 ms.
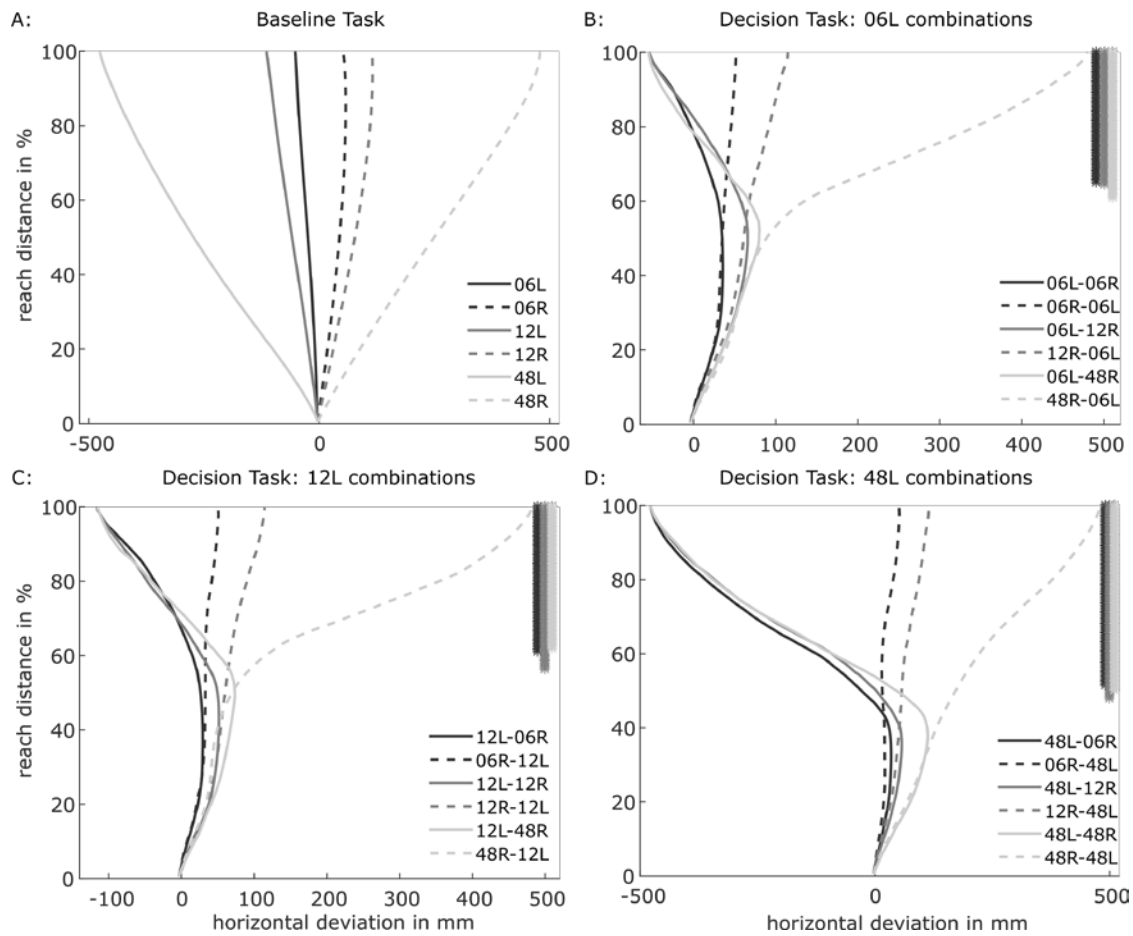
Figure 3A shows the mean normalised trajectories averaged across participants. As expected, participants performed relatively straight movements from the start position to the target position when they moved to any of the six different LED lights in the single target condition.
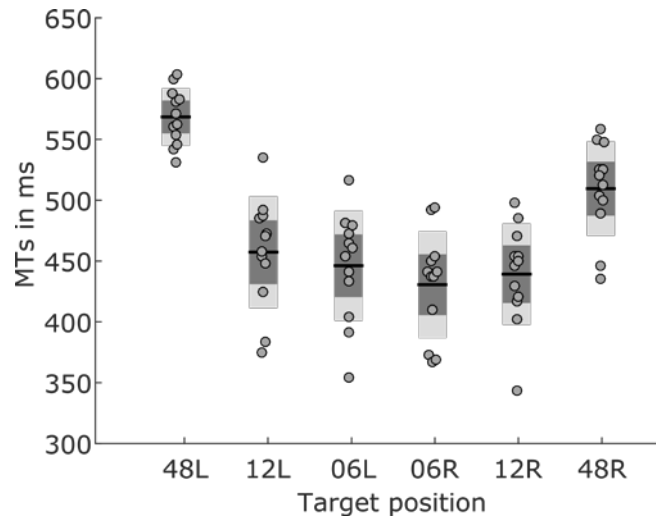
Figure 2: MTs in the single target condition for all six LED light positions. The black solid line indicates the mean value across all participants and the dark grey shaded area 95% confidence intervals. The light grey shaded areas indicate ± 1 SD and each of the dots depicts the median MT of one participant (N=12).

2.2.2 Decision Task

The averaged normalised trajectories for all dual-target conditions are displayed in Figure 3 B-D. Similarly to Chapman et al. (2010a), we found that for all initially visually identical conditions, reaches only started to differ after about 50-65% of the reaching distance was covered. For example, Figure 3B shows all light combinations where one of the initial LEDs was always 6 cm to the left of the display (Figure 3C for 12 cm to the left and Figure 3D for 48 cm to the left respectively). Solid lines indicate trials in which the left LED became the final pointing target after movement initiation, while dashed lines refer to trials where the right LED became the final pointing target. In initially visually identical conditions (indicated in the Figure by lines of the same shade of grey), participants initiated their movements similarly, as would be expected, and then adjusted the initial trajectory in flight in response to the final target location. However, the trajectories also suggest that the employed averaging strategies were not necessarily optimal with respect to keeping a minimal distance to the two possible targets on either side of the display. Firstly, all trajectories show a general bias toward the right-hand side. Secondly, the starting direction varied primarily with the distance of the LED light presented at the right side of the display, and not with distances on the left-hand side (in particular for central light locations). To confirm these observations statistically, we calculated the initial starting direction of the trajectory (see Methods section for more information).

Figure 3: Averaged normalised trajectories showing the lateral deviation of the hand over the reaching distance. A: Single-Target condition (baseline trials). B-D: Dual-Target conditions (decision trials): Lines of the same shade (solid and dashed) indicate conditions that are initially visually identical at the start of the trial; solid lines indicate the final target was on the left, and dashed lines indicate the right. The different shades of grey indicate conditions where the position of the right light varied (left light position is fixed in each panel). B: One of the lights was always 6 cm to the left, paired with any other right position. C: One of the lights was always 12 cm to the left, paired with any other right position. D: One of the lights was always 48 cm to the left, paired with any other right position. Please note that the x-axes vary between graphs. Significance bars at the side indicate when the two trajectories of the initially visually identical conditions started to differ from each other.

We first ran paired-samples t-tests comparing the start angles for each of the nine pairs of initially visually identical conditions. This analysis showed that, as expected and in line with our trajectory data (Figure 3), there were no differences for any of the tested pairs (all $t(11) \leq 2.39$, $p \geq .035$, Bonferroni-corrected significance threshold of $p = .006$) suggesting again that participants initiated their trajectories in the same direction for initially visually identical trials. Hence, we averaged the start directions for all initially visually identical conditions for further analysis (see Figure 4). Positive angles indicate deviations toward the right and negative angles would indicate a deviation toward the left of the midline (but never occurred). Bonferroni-corrected one-sample t-tests against zero confirmed that all start directions, except for the combination of the right 6 cm light

and the left 48 cm light (i.e., average start angle for R06-L48 and L48-R06), were different from zero and directed toward the right-hand side of the display (all t(11)=3.26-6.12, p≤.008). Most importantly, however, paired-sample t-tests comparing the initial start direction in the decision trials with those measured in the baseline condition revealed that for all combinations containing a central right light (at 6 or 12 cm laterality) starting angles did not significantly differ from the baseline trials in which movements were directed to the 6 or 12 cm target on the right respectively (all t(11)≤0.83, p>.43, indicated by the diamond symbols in Figure 4). As expected, all starting angles in the decision trials significantly differed from the starting angles observed for left-hand side targets in the baseline trials (L06= -5.5°, L12=-18.6°, L48=-56.0°, all t(11)=6.15-19.36, p<.001). If, however, a far distance right-hand side light (48R) was presented, starting angles were significantly smaller than in the baseline condition (all t(11)=3.22-7.71, p≤.008). In summary, these findings suggest that if central right-side lights were present as potential targets, participants tended to initiate a movement toward those right-side lights, which were corrected in flight if the left light was revealed as the final movement target. In other words, for central targets, our participants employ a pre-selection and correction strategy rather than an averaging strategy.

This interpretation was further confirmed by a 3 (left light position: 6, 12, 48 cm) x 3 (right light position: 6, 12, 48 cm) repeated measures ANOVA on the starting directions. The distance of the left light did not affect the starting angle of the movement, $F(2,22)=0.50$, $p=.62$ $\eta_p^2=.04$, , while the distance of the right light did have a small effect, $F(2,22)=3.62$, $p=.044$, $\eta_p^2=.25$. As shown in Figure 4, starting angles were consistently larger (i.e. further to the right) if the right light was located at 12 cm (25.2° ± 5.6°) as compared to 6 cm (18.2° ± 5.7°), p<.001. However, when the right light was located at 48 cm, the starting angle only further increased when the left-hand side target was also presented at the largest eccentricity of 48 cm. This was reflected in a significant interaction between the effect of distance for right and left light locations, $F(4,44)=7.27$, $p<.001$, $\eta_p^2=.40$. Note that the observation that the start direction is furthest to the right in the R48-L48 and L48-R48 conditions is again inconsistent with the hypothesised averaging approach. If participants had perfectly spatially averaged, their movements should be roughly aimed toward the middle of the display (at 0°) in those conditions. Thus, it seems that the general bias towards the right light positions becomes more pronounced as the left-hand side target gets more difficult to reach.
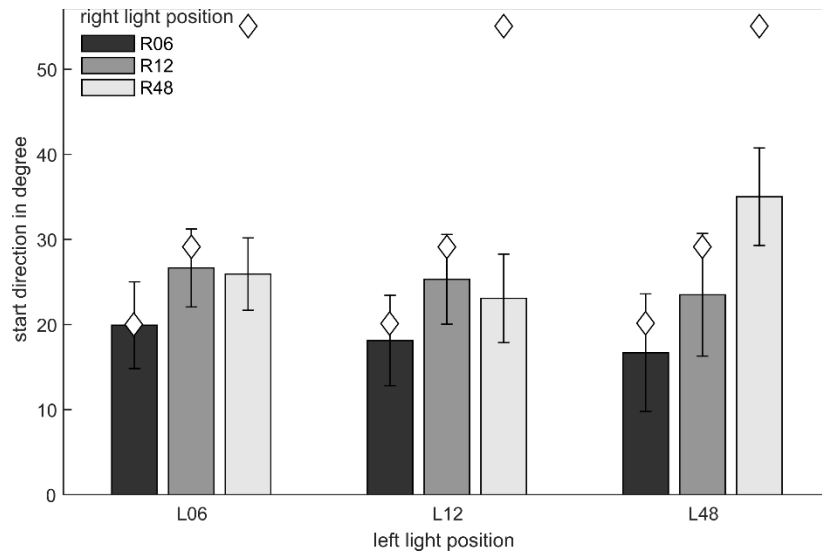
Figure 4: Start direction in degrees averaged across the initially visually identical conditions as a function of the LED light locations on the left- and right-hand side in the decision trials. Diamond symbols indicate the start angle in degree as measured in the baseline condition for the targets presented on the right-hand side at the corresponding distance. Error bars depict ±1 SEM between subjects.

The notion that movements tend to be initiated toward right lights and were corrected in flight when the left light became the final target is further supported by the fact that there was a considerable increase in movement times when the left lights became the target in the decision trials (see Figure 5). While movements were overall much slower in the decision task than in the baseline task, participants reached the right target lights considerably faster than the left target lights at all eccentricities.
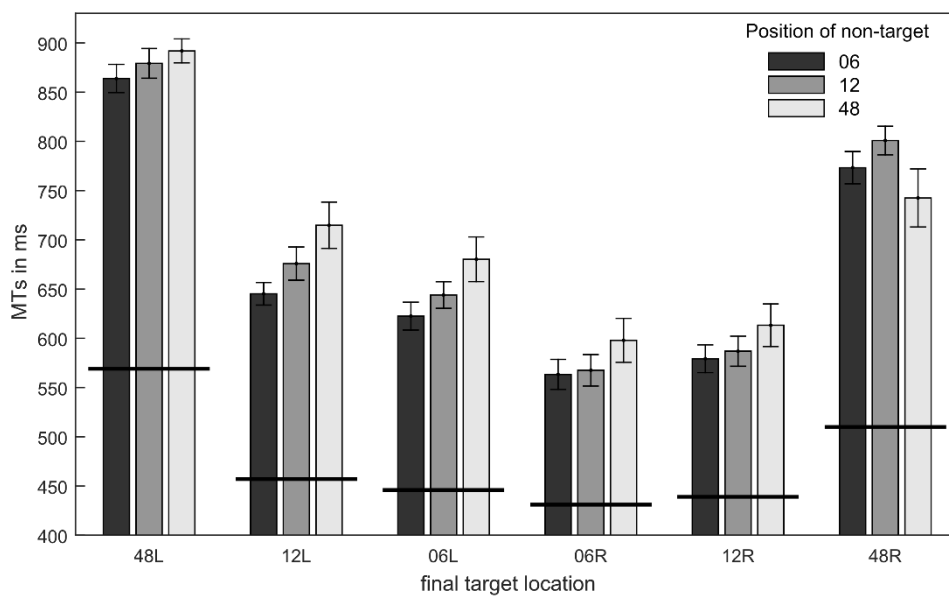
Figure 5: MTs for all different trial types of the decision task. Error bars depict ±1 SEM between subjects. Black horizontal lines indicate MTs in the baseline condition for the respective target distances.

In summary, our participants prefer to pre-select and initiate movements to positions that can be reached more easily (i.e., which have a biomechanical advantage) and more quickly. When the unselected light becomes the final target, participants adjust their movements in flight to reach the correct final location. We provided our participants with a generous amount of time (1000 ms) so this pre-select-and-correct strategy still allowed them to reach the final target within the time limit.

We next investigated how participants select their movement targets in a decision task under strict time pressure. In previous research, participants under time pressure tend to adopt an averaging strategy (e.g., Chapman et al., 2010a), meaning that the movement path is selected such that the distance of the hand to both potential targets is kept minimal until the final target is revealed. Here, we were particularly interested in whether and how participants' strategies would change if the time limit is selected such that averaging is a feasible strategy for lights presented at the central locations but cannot be successfully employed in conditions involving far-eccentricity lights. Participants were informed that trials with errors (i.e. pointing to the wrong LED or movements exceeding the time limit) would need to be repeated. Throughout the experiment, they were encouraged to complete as many trials as possible (i.e. by reaching to the correct final target within the available time limit) and were regularly told how many successful trials they had completed so far and how many remained. Thus, the optimal strategy (averaging for near distance and pre-selection for far distance light combinations) both maximizes the chance of success on any given trial, and minimizes the number of overall trials performed. Based on the findings of Experiment 1 showing that participants are able to reach far eccentricity lights within 600 ms in the baseline condition (see Figure 2) but not in the decision trials (see Figure 5), we decided to select this as the maximum MT-limit in our subsequent experiments. Note that this time limit should also encourage participants to adopt an averaging approach for near distance targets (rather than a pre-select and correct strategy) as this limit was exceeded in most conditions (see Figure 5).

3. Experiment 2 – Decision task under time-pressure

3.1 Methods

3.1.1 Participants

Twenty-three undergraduate and postgraduate students of the University of Aberdeen took part in this experiment. Three participants had to be excluded from data analysis – one who was unable to execute movements fast enough and two who turned out to not be naïve to the purpose of the experiment. Thus, the final sample consisted of 20 participants (aged between 19 to 29 years, mean

age: 22 years, 8 males). All participants had normal or corrected-to-normal visual acuity and were right-handed by self-report. They all provided informed consent prior to participating and were debriefed and reimbursed with £5 after participation. The study was approved by the School of Psychology Ethics Committee at the University of Aberdeen.

3.1.2   Setup, Stimuli and Procedure:

The setup was identical to that of Experiment 1 and the participants performed the exact same two tasks (i.e. baseline task and decision task). However, in contrast to Experiment 1, where tasks were performed under no real time pressure (i.e. maximum movement time of 1000 ms), participants were now allowed only 600 ms to reach the correct target light after movement initiation in the decision task. That is, trials were classified as errors and recycled if: 1) MTs exceeded 600 ms, 2) RTs exceeded 500 ms, or 3) they landed on the wrong location (i.e. the extinguished light, see section 3.2.2 for further information). Participants were told that these error trials had to be repeated (at a random position later in the experiment) and were given verbal feedback about their performance by the experimenter after each trial. The experimenter also regularly reported the total number of correct trials achieved, as additional encouragement for participants to obtain as many correct trials as possible. The experiment ended after 216 correct trials or after 60 min of testing time, whichever happened first. This means the number of trials varied by participant (see results section).

The initial set of trials, before recycling, contained 24 repetitions of each of the nine initially visually identical conditions, in a random order, like in Experiment 1. Unlike Experiment 1, however, we now determined randomly at the beginning of each trial which of the two LEDs would remain lit up and which would switch off (rather than balancing the number of left and right targets). This change was made to prevent participants from making predictions about the position of the upcoming target based on the number of trials they already had performed in one condition. For example, in Experiment 1, participants had a preference to move to the 6R light if it was presented together with the 48L light. Therefore, participants are likely to make more mistakes (i.e. too slow movements) when the 48L light becomes the target. If we recycled all these trials, eventually there would be a point near the end of the experiment at which the 48L light would always become the final target for this combination. To avoid this, we randomly selected which light would be the target on each trial. If a trial of a visually identical condition had to be recycled, whether the left or right light was the target would be determined for this trial anew each time it was recycled, until 24 correct trials were completed for this specific combination of lights. The total number of trials was again 216, though due to recycling of error trials, participants

completed far more than this number in the allotted hour of testing (range was 293 to 434 trials, see Figure 7), and most participants did not complete all 216 trials correctly. Before the start of the decision task, participants performed 10 practice trials.

3.2 Results and Discussion

3.2.1 Baseline Task

Figure 6A shows the MTs in each condition averaged across participants as well as the median MT for each participant separately. Across all 1440 trials, 132 were recycled during the experiment in this task as participants exceeded the RT limit (9.1%) and 10 out 1440 trials (0.7%) had to be excluded from this analysis as the end of movement could not be determined reliably using our criteria. Similarly to Experiment 1, the 3 (distance of LED) x 2 (side of LED) repeated measures ANOVA confirmed a significant main effect of LED side, $F_{(1,19)}=76.47$, $p<.001$, $\eta_p^2=.80$, reflecting that participants were quicker to reach lights on the right side. There was also a main effect of LED distance, $F_{(2,38)}=355.87$, $p<.001$, $\eta_p^2=.95$, with MTs increasing for larger target eccentricities. Again, the interaction effect was also significant, $F_{(2,38)}=30.06$, $p<.001$, $\eta_p^2=.61$, indicating a larger increase in MTs with distance on the left hand side relative to the right. Importantly, all participants could reach all targets within the allocated time limit of 600 ms (see Figure 6A).
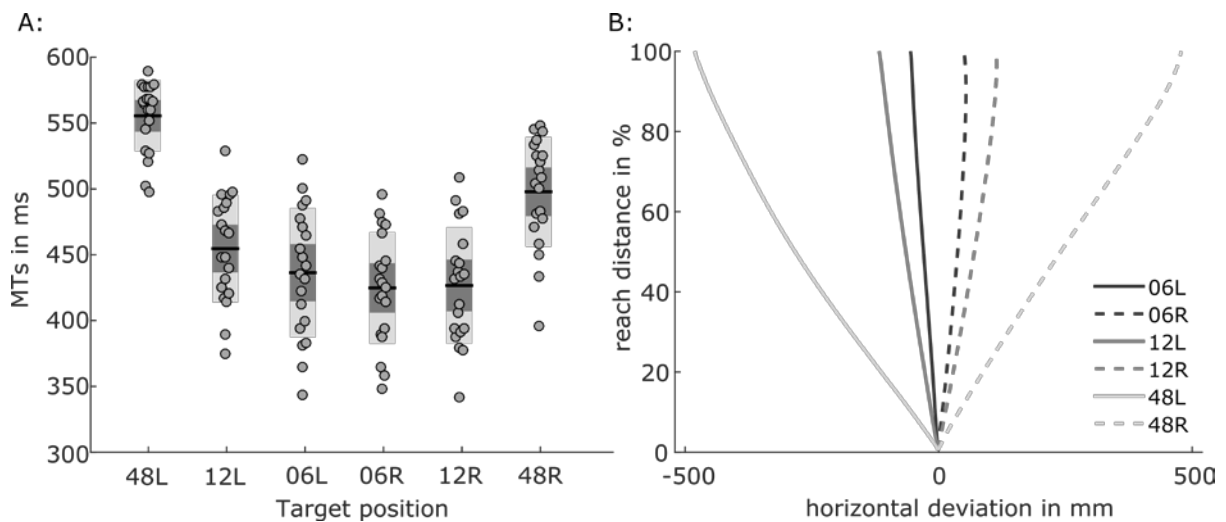
Figure 6: A: MTs in the single target condition for all six target positions. The black line indicates the mean value across all participants and the dark grey shaded area 95% confidence intervals. The light grey shaded areas indicate ± 1 SD and each of the dots depicts the median MT of one participant (N=20). B: Average normalised trajectories in the baseline task showing the lateral deviation of the hand over the reaching distance.

Figure 6B shows the mean normalised trajectories averaged across participants. As expected and similar to Experiment 1, participants performed relatively straight movements from the start

position to the target position when they moved to any of the six different targets in the single target condition.

3.2.2 Decision task

During the decision task, all "error trials" were saved but went back into the trial pool and participants were made aware that error trials would need to be repeated (with the aim of achieving 216 correct trials). Error trials were later sorted offline into time-out trials (MTs longer than 600 ms) and position error trials (i.e. participants pointed to the light that turned off). In order to be classified as a position error the lateral distance of the movement endpoint had to be within ±30 mm of the non-target position (i.e. the light that turned off). Trials that could not be classified using these criteria (e.g., error-trials ending the movement at a different position than either of the two potential targets), or for which the MTs could not be determined reliably, were excluded from further analysis (about 4.5% of all trials).

Figure 7 (upper panel) shows the percentage of correct trials, time-out errors and position errors for each participant averaged across all movement conditions as well as the total number of trials they performed. The figure illustrates that the majority (i.e. 80%) of participants made a large number of time-out errors but comparatively few position errors (except for Participants 4, 6, 12 and 17).

Note that we purposefully selected our MT-limit such that participants should be unsuccessful in using an averaging strategy when one of the potential targets was presented at the furthest eccentricity (i.e. 48 cm to the left or right of the midline). Therefore, for trials involving a far-eccentricity light, a pre-selection strategy should ensure a success rate of around 50%, while (perfect) averaging (i.e. going initially exactly in the middle between the two targets) will result in a success rate close to 0% (well below 50% chance-level). Hypothetically, there are two ways of accomplishing a maximum (50%) success rate for trials involving far-eccentricity lights through pre-selection. Firstly, if participants take biomechanical constraints into account and consistently select the easier/quicker to reach target (e.g., the right light in the 06R-48L/48L-06R conditions), they should always move to the light that is presented closer to the midline. Consequently, they should make position errors in all situations in which the far left light becomes the final target location (e.g., 0% correct in the 48L-06R condition). In contrast, they should make virtually no position errors in situations in which the near or mid-eccentricity right light becomes the final reaching target (e.g., 100% correct in the 06R-48L condition). Thus, the average success rate across the two visually identical conditions should result in about 50% correct.
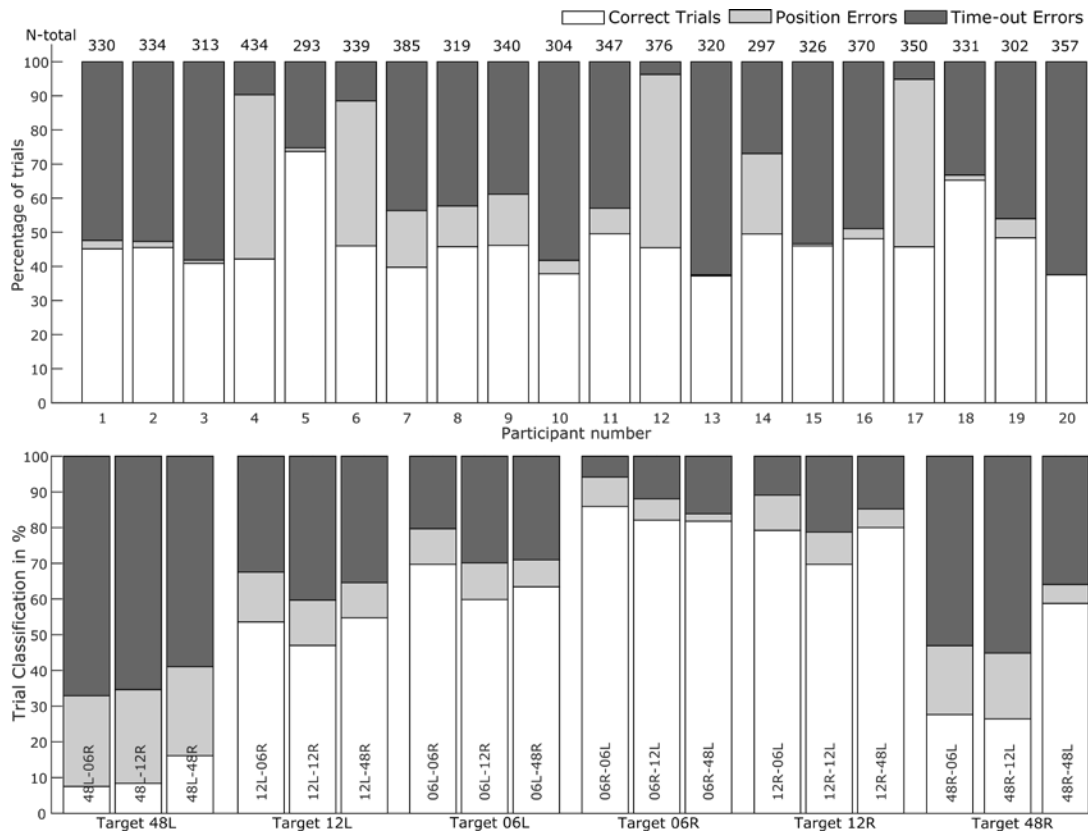
Figure 7: Upper panel: Percentage of correct trials, position errors and time out errors for each of the 20 participants averaged across all movement conditions. The numbers on the top of each bar show how the total number of completed trials per participant. Lower panel: Percentage of correct trials, position errors and time-out errors in each of the 18 movement conditions, averaged across all participants.

Alternatively, participants could randomly pre-select one of the two lights (as they were aware that final targets were determined randomly) without taking potential movement constraints into consideration. This strategy should also result in about 50% of position error and 50% of correct trials across the two initially visually identical conditions involving far eccentricity lights independent of the final location of the target (e.g., 50% correct in 48L-06R condition and 50% correct in 06R-48L condition).

The lower panel of Figure 7 shows the percentage of correct trials, position errors and time out errors for all 18 movement conditions. It becomes clear that participants do not opt for a random pre-selection strategy as they achieve many more correct trials when the near-or mid eccentricity light becomes the final target in the initially visually identical conditions involving a far eccentricity target (e.g., 81.8% correct for 06R-48L condition vs. 7.5% correct for the 48L-06R condition). Moreover, in all far eccentricity trials, participants also commit proportionally more time-out errors than position errors (more than 50% of time-out errors in all trials where the far eccentricity light becomes the final target). This might, at first glance, suggest that participants fail to adopt either of the two optimal strategies for trials involving far-eccentricity targets. However,

21

time-out errors may have two different causes: First, participants may employ a strategy of, at least partially, spatially averaging their movements to account for both possible target positions even for trials featuring a far eccentricity light (sub-optimal). Alternatively, participants may have pre-selected the wrong light but attempted to correct their movement path in flight as soon as the final target was revealed. The latter strategy would still be a very reasonable (and optimal) approach as it provides participants with a chance of getting the trial correct; they would commit a position error with certainty if they were to continue their pre-planned movement without correction.

To further test which of the two explanations (partial averaging vs. correction strategy) is more likely, we calculated the average movement trajectories separately for both types of error trials (i.e. time-out and position errors)[2].
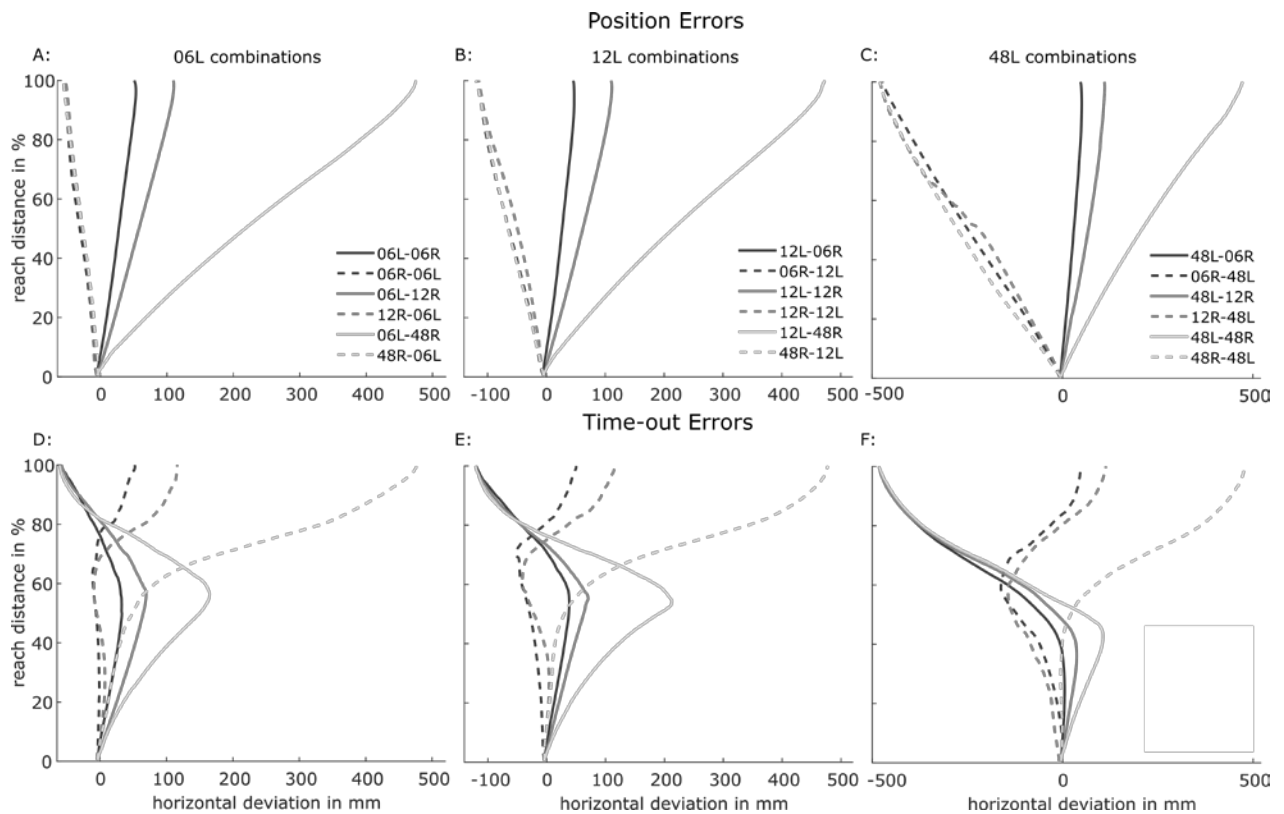

Figure 8: Average normalised trajectories in the different error-trials of the decision task showing the lateral deviation of the hand over the reaching distance. A-C: Trajectories for position errors in which participants reached to the wrong final target location (i.e. light that was switched off). D-F: Trajectories in trials that were classified as time-out errors (i.e. exceeded MTs of 600 ms). The left light was at 6 cm (A, D), 12 cm (B, E) or 48 cm (C, F) paired with any

---

[2] Note, that for our current research question we were primarily interested in the type and frequency of errors participants made. The correct trial trajectories are difficult to interpret because the strategies that can lead to a correct response can be constrained or flexible, depending on the individual participants' reaching speed and the spatial configuration of targets. For more information on correct trials (depending on participants' preferred strategy), refer to the supplementary material of this article.

other light on the right-hand side. Note that the different conditions contain a varying amount of trials and that not all participants made all types of errors in the different conditions.

Figure 8 shows that trajectories that resulted in the commitment of a position error (upper row) look descriptively similar to those recorded in the baseline task (i.e. straight movement path from start position to target location). For time-out errors (lower row), trajectories for initially visually identical trials (depicted as the same grey shade) tended to differ right from the beginning, suggesting that movements were, at least partly, pre-planned and corrected their trajectories later.

To find out whether movements in error trials were indeed pre-planned to just one location (similar to baseline trials), we calculated the start directions of the movements separately for position and time-out trials. If time-out errors would be the mere consequence of a preselection and correction strategy, start angles should be identical to those observed for position errors. Figure 9 confirms that the starting directions for position error trials are comparable to those measured for the respective locations in the baseline trials, suggesting that participants pre-selected one light location and planned their movement accordingly. Note that we refrain from any statistical analysis, as there were overall very few position error trials (most participants only committed position errors in some conditions but not others, and some participants did not commit any position errors at all, cf. Figure 7). For time-out errors, the average start directions tended to vary for initially visually identical trials, but also seemed to differ from the start direction as determined in the baseline trials (for right-sided targets). In particular, when far eccentricity lights were presented, the average start direction was considerably attenuated, suggesting that participants did not consistently pre-select one of the target lights.
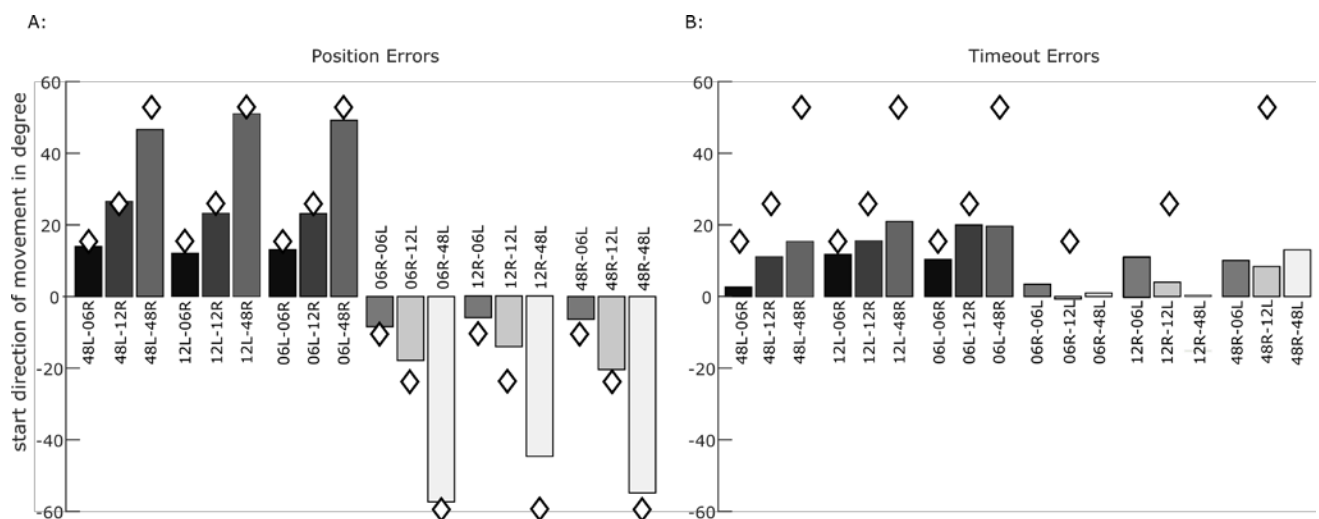
Figure 9: A: Average start direction of the movement for position error trials. Note that only four participants (out of twenty) made a considerable number of position errors. B: Average start direction for time-out error trials. Diamonds indicate the average start direction in the corresponding baseline condition. Note that for time-out errors we only indicated the average start direction toward right-sided targets in the baseline condition as participants showed

a right-bias in all conditions. Positive values indicate a start direction to the right and negative values to the left-hand side.

To recapitulate, spatial averaging would have been the optimal strategy for conditions where both lights were presented close to the midline, but this strategy leads to consistent failure for widely-spaced lights. Pre-selecting targets is the better strategy when the lights are widely spaced, because it ensures a success rate of about 50%. The fact that starting directions in time-out trials were quite different from those observed in the baseline trials, in particular in conditions involving far eccentricity lights, indicates that participants did not consistently apply such a pre-selection strategy in those trials, but may have at least partially averaged[3]. Most importantly, even though there were differences in the types of errors different participants made (see Figure 7, and discussed in more detail below), there was no indication that any of the participants consistently switched between averaging and pre-selection strategies to maximize success across different separations.

While none of the participants could be considered optimal, individual differences in strategy choices complicate a straightforward conclusion about what strategy people, as a group, adopt. Figure 10 shows examples of these differences: Participant 4, who committed primarily position errors and Participant 20, who only committed time-out errors. This figure nicely illustrates that participants seem to stick to their preferred strategy and do not adjust depending on overall task success and trial outcome. Participant 4 committed a large number of position errors even in the near distance target combinations, which could have been easily avoided by applying an averaging strategy instead. This participant also showed a strong preference for right sided targets and almost never (pre)selected the 48L target if present (confirming the preference for easy-to-reach targets observed in Experiment 1). In contrast, Participant 20 committed no position errors at all, but consistently failed to reach the target on trials where either light at 48 cm became the final reaching target. However, for the condition with two nearest eccentricity targets they rarely made any mistakes. This indicates an over-application of spatial averaging strategy (with a bias towards the easier to reach target). If the time out errors of this participant were a consequence of a consistent pre-selection and correction strategy, they should have performed about 50% correct across visually identical trials, which is clearly not the case (see Figure 11). Instead, they consistently timed out in these trials, even after repeating far-eccentricity trials up to 30 times without success (Figure 10).

---

[3] Refer to supplementary material for more information and histogram plots of start directions for position errors and time-out errors.
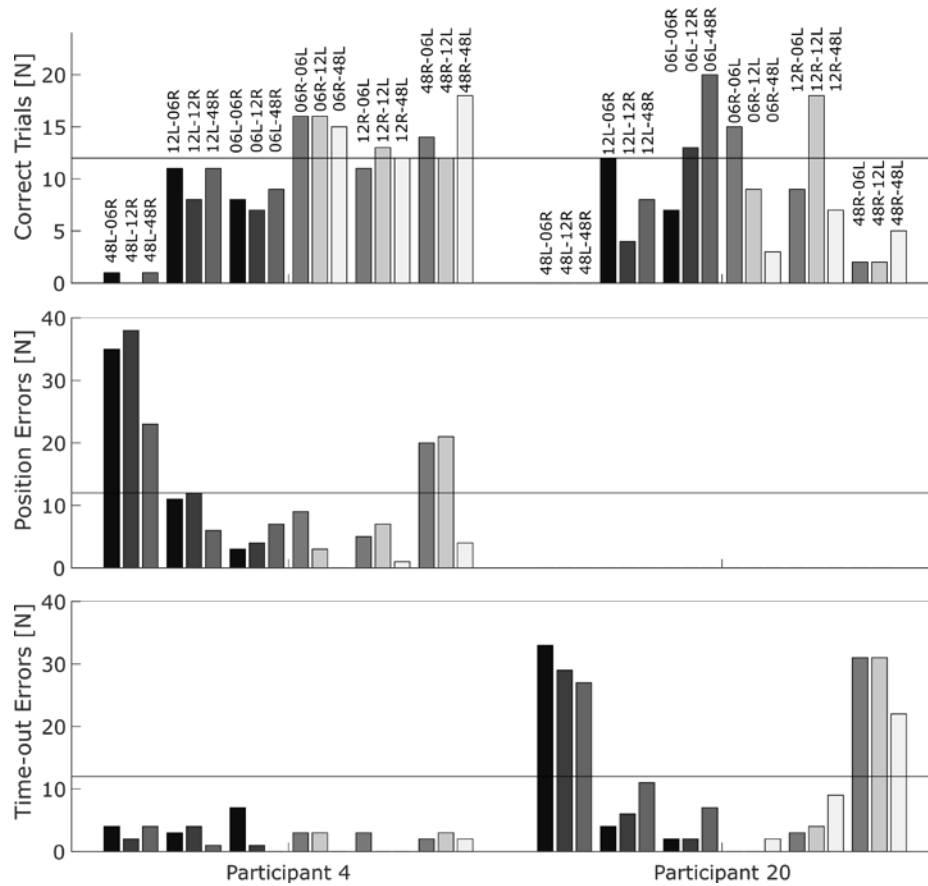
Figure 10: The number of correct trials, position errors and time out errors for two example participants who applied different strategies when performing the task. Participant 4 tends to pre-select and therefore makes a large number of position errors. Participant 20 tends to spatially average, leading to a large number of time-out errors. Neither participant adapted their strategy depending on the spatial separation of the targets, resulting in a relative low number of correctly completed trials (see also Figure 7 and 11).

The average success rates, across all participants, for all initially visually identical trials are shown in Figure 11, which also illustrates how performance rates varied for participants that used different strategies. As expected, Participant 4, who primarily pre-selected targets in all trial combinations, performed close to 50% correct across all visually-identical conditions, independent of the target separation (white diamonds in Figure 11). In contrast, as one would expect from an over-application of an averaging approach, Participant 20 performed well above chance in trials featuring near distance targets (80-90% correct), but well below the possible 50% chance level (achievable with a pre-selection approach) in trials featuring a 48 cm target (black triangles in Figure 11). Note that due to motor noise, we cannot expect 100% accurate performance in near eccentricity trials even when averaging is used consistently.

Statistically comparing participants' average performance level in all five initially visually identical trials containing far eccentricity targets against the maximum possible success rate of 50% further confirms that participants, on average, perform significantly worse than the theoretically

possible 50% (all t(19) >2.88 p<.01), except for the combination where the 06 left light and the 48 right light are presented together (45.5% correct trials, t(19)=1.24, p=.23).
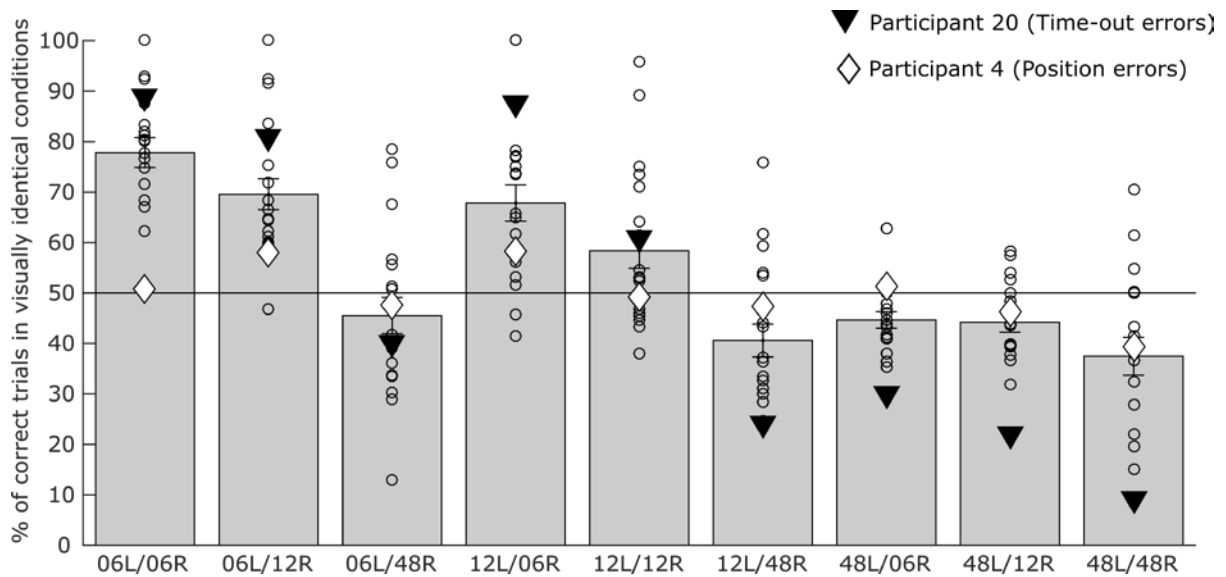
Figure 11: Average success rates for all initially visually identical conditions. The largest proportion of correct trials should be achieved by consistently applying an averaging approach in trials where targets are separated by a small distance (close to 100%) while a pre-selection approach should allow a success rate of about 50% in trials featuring a far eccentricity light. Open circles show data of individual participants. White diamonds show the percentage of correct trials for Participant 4 who committed primarily position errors and black triangles the data of Participant 20 who only committed time out errors (see text for further information).

Taken together, these findings suggest that participants are not consistently and flexibly selecting movement strategies that would guarantee them the highest likelihood of success (i.e. reaching the final target in time). However, the fact that participants showed a large bias for lights presented closer to the midline and on the right-hand side of the display indicates that they do take biomechanical constraints into account when selecting their movements (i.e. preparing movements to positions that are easier and/or quicker to reach). Most participants tend to partially pre-select and partially spatially average their movements, rapidly and efficiently choosing to execute movements towards the more expedient of the two possible targets, but with trajectories that are modulated by the position of the non-selected light to make corrections more likely to be successful, if they are required (see also supplementary material). For close targets, this is a successful strategy, but this modulation by the non-selected light is often observed even for trials involving far-eccentricity targets, where these deviations from a straight trajectory have the effect of lowering performance below the maximum accuracy of 50% that could be attained through pre-selection alone. A minority of participants adopted a pre-selection strategy at far eccentricities, but they over-apply this strategy by also pre-selecting at close eccentricities, thereby also making errors for targets that should have been easy to reach within the time limit.

This failure to modify strategy with the increase in eccentricity is reminiscent of the results observed in Clarke and Hunt (2016), in which participants did not modify their decisions with changes in task difficulty to maximize their response accuracy. In that study, as in this one, this failure to modify strategy across difficulty has an especially detrimental effect in difficult conditions, in which participants could ensure a performance level of 50% if they consistently focus on executing one of two equally likely responses. Instead, many participants attempt to prepare for both responses simultaneously and fail frequently. Those few who do successfully apply this strategy tend to over-apply it, and their performance for the easier trials suffers as a result.

However, as outlined in the introduction, controversy exists over whether averaging behaviour represents a deliberate optimal strategy or is a result of unintentional averaging of competing motor plans. Therefore, an alternative explanation for the results of Experiment 2 is that the persistent spatial averaging for light combinations involving far eccentricity targets is not a deliberate (but suboptimal) strategy, but a reflexive consequence of the competition between two potential targets. To address this question more directly we conducted a final experiment, in which we encouraged participants to pre-select one target position in our task instructions. If participants are able to over-ride the tendency to (partially) average, we should see straight trajectories to the selected position and an accuracy around 50% across all initially visually identical light pairs. If this pattern is observed, we can use these data to further explore the extent to which participants take biomechanical constraints into account when they are encouraged to pre-select a target location.

4. Experiment 3: Encouraging participants to pre-select a target

4.1 Methods

4.1.1 Participants

Twenty undergraduate and postgraduate students of the University of Aberdeen took part in this experiment (age-range 18 to 37 years, mean age: 21 years, 8 males). All participants were right-handed by self-report and had normal or corrected-to-normal visual acuity. They all provided informed consent prior to participating. Participants either received course credits or were reimbursed with £5 after participation. The study was approved by the School of Psychology Ethics Committee at the University of Aberdeen.

4.1.2   Setup, Stimuli and Procedure:

The setup and task were the same as those described for Experiment 2: Participants performed both a baseline task (single target condition) followed by a decision task (dual-target condition). In both tasks, they were allowed a maximum MT of 600 ms to reach the final target after they had started their movements. The only differences to Experiment 2 were that we now instructed participants to reach to one of the lights as quickly and accurately as possible (while still trying to get as many trials as possible correct) and that we now presented exactly 12 trials for all of the 18 movement conditions in random order throughout the experiment (resulting again in a total of 216 trials). This latter adjustment was introduced to encourage participants to also select far eccentricity lights as targets in the decision trials, because participants were made aware that each position would become the final target during the experiment equally often. Again, the experiment was aborted after 60 minutes, resulting in a varying number of trials per participant.

4.2   Results and Discussion

4.2.1   Baseline Task

As in Experiment 1 and 2, participants first completed a single-target condition in which they were required to point to one of the target LEDs as soon as it switched on. Across all 1440 trials, 133 were recycled during the experiment in this task as participants exceeded the RT-limit (9.2%) and 6 out 1440 trials (0.4%) had to be excluded from this analysis as the end of movement could not be determined reliably using our criteria. Figure 12 shows MTs and Figure 14A the normalised and averaged trajectories across all participants. We confirmed again that MTs increased with pointing distance, $F(2,38) = 409.16$, $p<.001$, $\eta_p^2 =.96$, and were longer for lights presented on the left-hand side of the display, $F(1,19) = 98.69$, $p<.001$, $\eta_p^2 =.84$. The increase in MTs with reaching distance was again longer for left-sided targets as indicated by an interaction effect, $F(2,38)=65.62$, $p<.001$, $\eta_p^2 =.78$. As in Experiment 1 and 2, participants were well able to reach the targets within 600 ms, and showed straight reaching paths from the start position to the target.
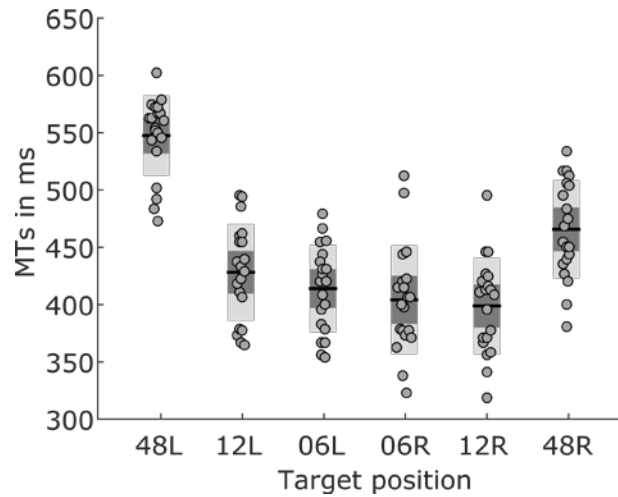
Figure 12: MTs in the single target condition for all six target positions. The black line indicates the mean value across all participants and the dark grey shaded area 95% confidence intervals. The light grey shaded areas indicate ± 1 SD and each of the dots depicts the median MT of one participant (N=20).
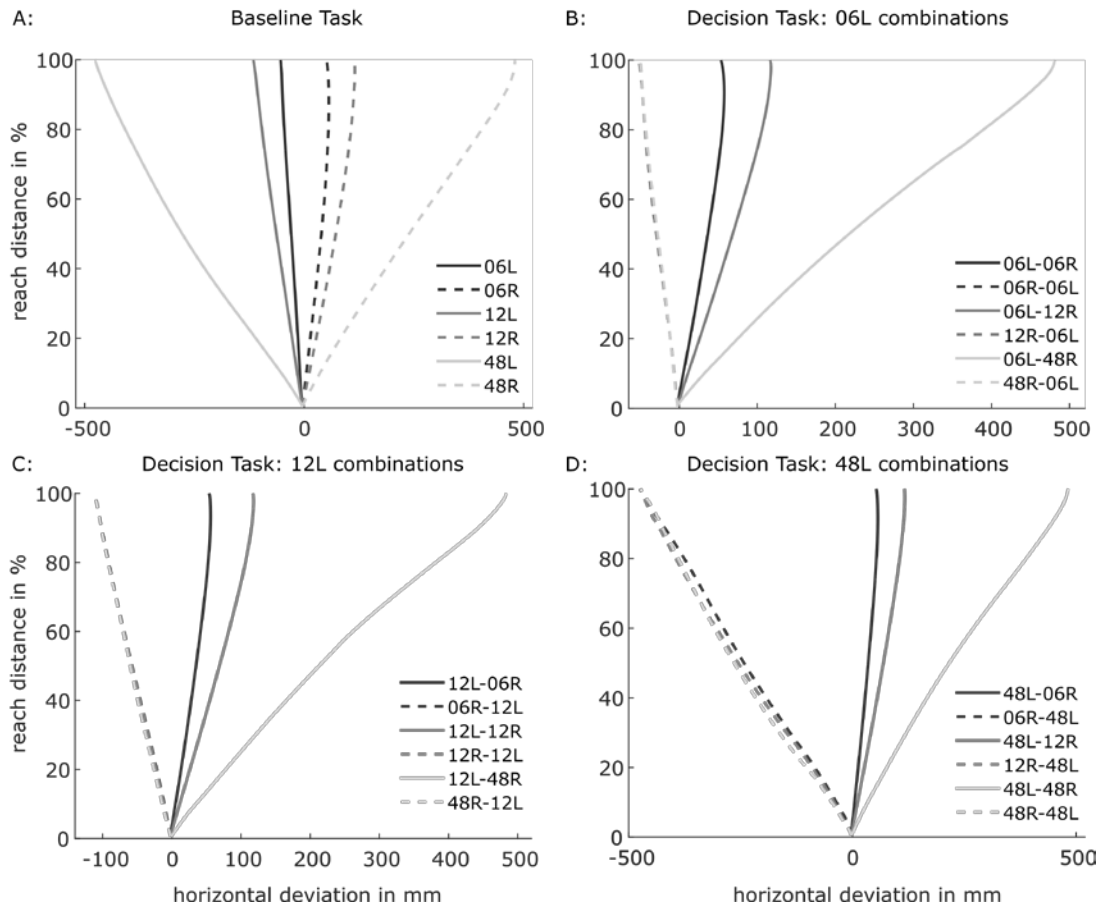
4.2.2 Decision task

As in Experiment 2, all trials that either exceeded the movement time limit or were executed to the wrong location were classified as errors and participants were aware that any error trials would need to be repeated at a later random position within the experiment (with the aim of completing all 216 trials). We slightly changed the instruction compared to Experiment 2 by encouraging participants to go for one of the lights whilst still trying to get as many correct trials as possible throughout the experiment. Figure 13 shows that this instruction had quite a drastic effect on the types of errors participants made, relative to Experiment 2 (cf. Figure 7).
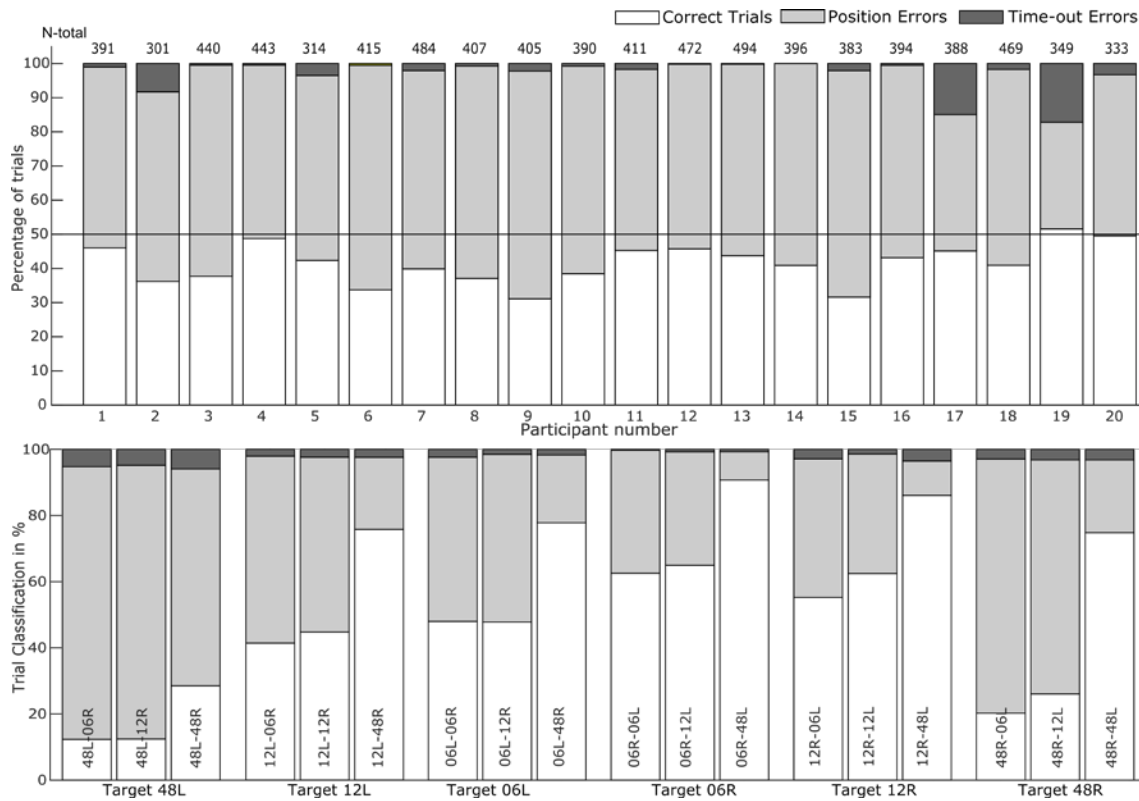
Figure 13: Upper panel: Percentage of correct trials, position errors and time out errors for each of the 20 participants averaged across all movement conditions. The numbers on top of each bar show how many trials each participant performed during the experiment. Lower panel: Percentage of correct trials, position errors and time-out errors in each of the 18 movement conditions averaged across all participants.

In contrast to Experiment 2, the number of position errors increased considerably, while very few time-out errors were committed overall. As the number of time-out errors across the different participants and trial types was negligible, we focus our analysis on the position errors.

The normalised trajectories for trials in which participants pointed to the wrong location are depicted in Figure 14B-D, illustrating that the movement paths in these conditions were virtually identical to those observed in the corresponding baseline conditions. Most importantly, movement paths were unaffected by the position of the second potential target light. This was further confirmed by a comparison of the start directions between position error trials and baseline trials as depicted in Figure 15. Generally, movements were clearly initiated in the direction of the non-target light. Overall, as a consequence of implementing this "selection" strategy, success rates for reaches in all nine initially visually identical conditions were between 49% to 55% and did not differ significantly from 50% (all t(19)=-0.5-2.20, p≥.04).

Figure 14: Averaged normalised trajectories showing the lateral deviation of the hand over the reaching distance. A: Single-target condition (baseline task). B-D: Trajectories of position errors in the decision task: Lines of the same shade (solid and dashed) indicate conditions that are initially visually identical at the start of the trial. The different shades of grey indicate conditions where the position of the right light varied (left light position is fixed in each panel). B: One of the lights was always 6 cm to the left, paired with any other right position. C: One of the lights was always 12 cm to the left, paired with any other right position. D: One of the lights was always 48 cm to the left, paired with any other right position. Please note that the x-axes vary between graphs.

What is most noticeable when looking at the distribution of position errors across the different trial types (Figure 13, lower panel) is that they are not equally likely for all conditions, suggesting that participants are strategically (not randomly) selecting their movement targets. That is, if both lights were located centrally and were quite easy to reach (e.g., L06-R06 and R06-L06 conditions), participants showed only a small preference for the right-sided light, and were likely to choose either position as indicated by a similar number of position errors in both trial types (L06-R06: 49.6 ± 5.2 % vs. R06-L06: 37.2 ± 3.7 , t(19) =2.09, p=.05). In contrast, if one of the lights was more difficult to reach, such as in the L48-R06 condition, participants committed 82.4 ± 3.7% position errors if the target was located at far left-hand side but only 8.5 ± 2.7% position errors if the final target was located at the near right-hand side (06R-48L), t(19)=12.76, p<.001. In fact, out of our twenty participants, three never even attempted to reach for a light on the far left in this condition.
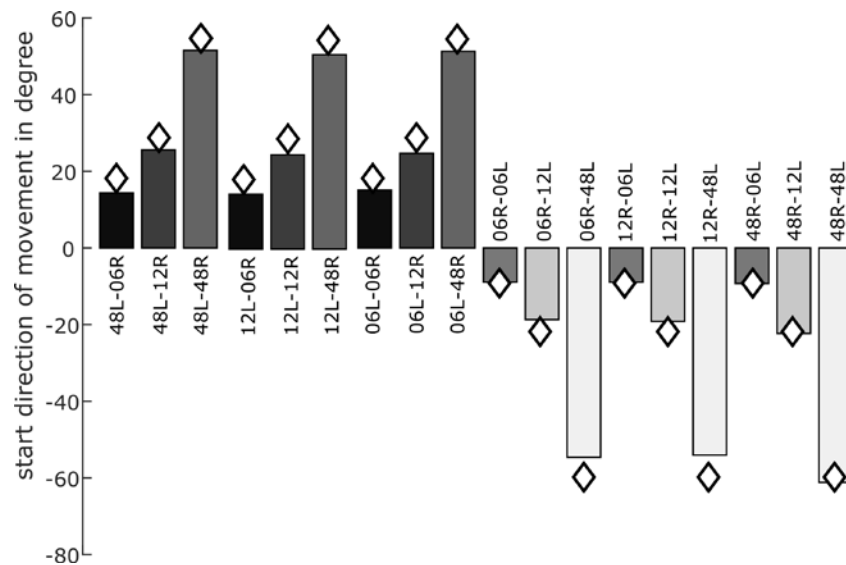
Figure 15: Average start direction of movement in position error trials. Diamonds indicate the average start direction in the corresponding baseline condition, averaged across all participants.

These findings suggest that the ease with which a target can be reached is reflected in how often participants select it over the alternative target in the decision task. To explore the relationship between the time it takes participants to reach each target and the probability of selecting each of these targets when doing the task, we calculated two distinct measures: First, we determined, for each of the visually-identical conditions, the difference in MT when reaching for each target as measured in the baseline condition (for example, if it took a participant on average 420 ms to reach to the 06L target but 450 ms to reach to the 12R in the baseline condition, this indicates a 30 ms advantage for this participant of reaching for the left light in the L06-R12 and R12-L06 decision trials). In a second step, we calculated how much more likely each participant was to select one light over the other in the initially visually identical decision trials. For this, we calculated the difference in the percentage of correct trials achieved in each of the nine initially visually identical conditions (e.g., if a participant performed 65% of correct trials in the L06-R12 condition (with 35% being position errors) and 45% of correct trials in the R12-L06 conditions (with 55% being position errors), we took the difference to indicate that this participant was 20% more likely to select the L06 light over the R12 light as their target). Figure 16 shows the relationship between the preference for a light in the decision trials and the MT advantage in the baseline condition for all participants and conditions. In fact, there was a very strong correlation, $r(180)=.78$, $p<.001$, with participants showing a clear preference for selecting lights as targets that were quicker to reach.
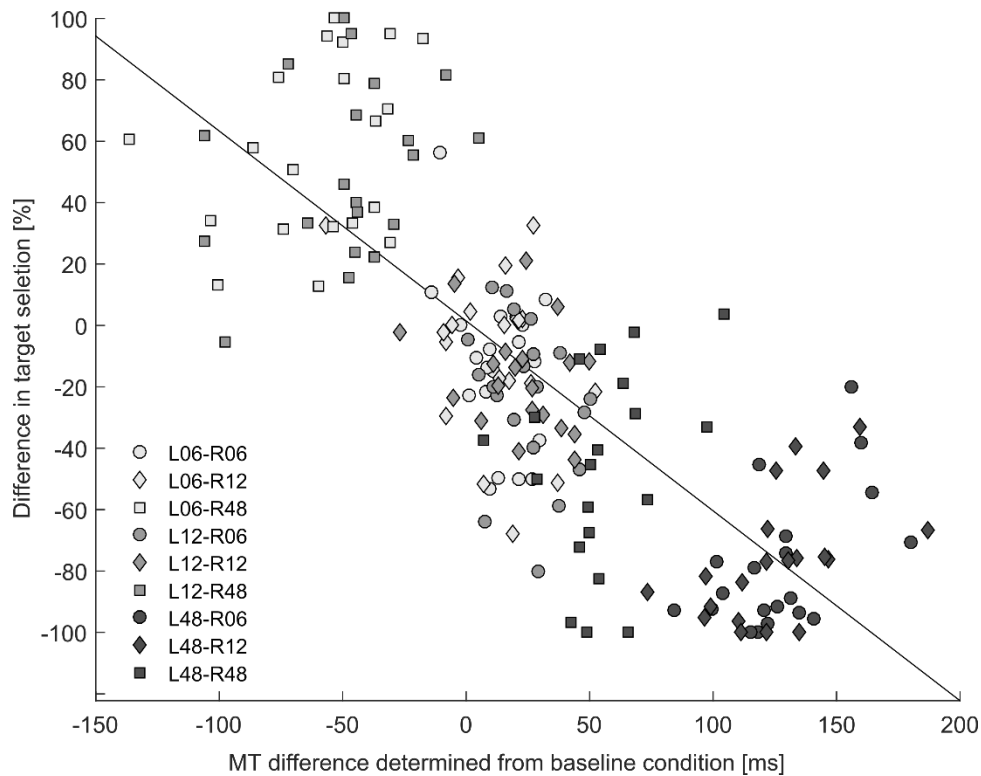
Figure 16: Correlation between MT difference in the baseline condition (x-axis) and preference to select targets in the decision trials (y-axis). Note that negative numbers on the y-axis indicate a preference for the right hand side light and positive number a preference for lights on the left hand side. Each symbol reflects one participant in one of the decision conditions.

In other words, the findings from this experiment show that participants have a very strong preference to select lights as targets that they can reach more quickly (and easily), suggesting that humans are extremely efficient in making economical decisions in a speeded visuo-motor choice task.

5. <u>General Discussion</u>

The aim of this study was to investigate which strategies humans would employ when the optimal movement strategy depended on the exact stimulus arrangement and varied from trial to trial. We demonstrate that participants can quickly and consistently select the more expedient of two targets, but on a broader level, their behaviour does not flexibly adjust to changes in the spatial configuration by shifting between averaging and pre-selection strategies on a trial-by-trial basis to enhance success.

Previous studies report seemingly contradictory findings regarding the optimality with which humans select and perform their movements. On the one hand, studies that investigated decision making in tasks that required some insight into the task demands have come to the conclusion that humans do not tend to select optimal strategies at all with respect to increasing their task probability of success (Clarke & Hunt, 2016; Morvan & Maloney, 2009; Nowakowska et

al., 2017). On the other hand, experiments that studied more implicit visuomotor decision making in fast reaching tasks seem to suggest that humans usually select movements that optimize their performance (Chapman et al., 2010a; Haith et al., 2015; Trommershäuser et al., 2008; Wong & Haith, 2017; but see also Jarvstad et al., 2014, Neyedli & LeBlanc, 2017). Our results use a fast reaching task and the results are consistent with the former, rather than the latter, conclusion. Even though our participants could be encouraged, through instruction, to adopt one strategy over the other, there was no evidence that participants could flexibly adapt their fast reaching strategies on a trial-by-trial basis to maximize task success. Instead, estimates of the relative economy of the immediate movement options seem to be the primary driver of target selection and motor execution.

To investigate whether participants modify their strategies in accordance with the specific demands of the task, we adapted the "go-before-you-know" task by including a pair of far eccentricity targets (i.e. 48 cm from midline) at either side of the display, in addition to the closer targets previously tested (at 6 and 12 cm from midline, similar to Chapman et al., 2010a). Experiment 1 served two main purposes: First, to find out what a suitable time limit would be to ensure that most participants would be able to employ an averaging strategy for close eccentricity targets but not for trials featuring far eccentricity targets; Second, to test the occurrence of averaging behaviour in a decision task with a more generous time limit. We made a number of noteworthy observations in this experiment. First, we replicated that movement paths for initially visually identical light combinations started to diverge only after about 50% of the reach distance was covered (see also Gallivan & Chapman, 2014). However, most of the observed trajectories did not seem to reflect an averaging strategy. Instead, for all conditions featuring near-eccentricity lights on the right hand-side of the display (R06 and R12), the initial start directions of the hand were very similar to those measured in the baseline condition for these targets. In other words, rather than spatially averaging, our participants seemed to employ a "pre-select and correct-in-flight" strategy, with movements being directed toward the right near and mid eccentricity targets. Only in trials in which one of the lights was presented on the far right side (R48) was the initial start direction attenuated in line with an averaging approach (but still biased toward the right hand-side of the display).

Even though these findings might, at first glance, appear incompatible with the previous literature, we would argue that they do not challenge the notion that averaging does occur under specific circumstances, but instead show the limits under which such behaviour can be observed. In fact, Gallivan and Chapman (2014) already pointed out that rapid movement time requirements have to be in place in order to observe spatial averaging reliably. They also noted that in their

experiments about one in four participants were unable to meet the strict temporal and spatial demands of the task and hence had to be excluded from analysis. By employing a much more generous time limit in our study, we were able to use a less restricted sample of participants and show that they seem to prefer a different strategy, i.e. pre-selection and correction, if time allows.

For light pairs including far-eccentricity targets, movement paths in Experiment 1 more closely followed the predictions of the averaging account, deviating significantly from the trajectories in the single-target baseline condition. This is likely because the time limit was more constraining for these longer movements. Similar to most previous studies employing "go-before-you-know" tasks, we also observed a strong right-hand-side bias in those conditions. For example, in the condition where both far eccentricity targets were presented (R48-L48 and L48-R48) and the average start direction would be straight ahead (start angle of 0º), movements were strongly biased, with a start angle of about 35º to the right. One fairly obvious explanation for this right hand-side bias is that, similar to previous studies, we only tested right-handed participants. For them, movements to the right are more comfortable as the arm does not need to move across the body (for discussion see also Gallivan & Chapman, 2014). Furthermore, ipsilateral movements are biomechanically easier, as moving the arm across the body midline requires a larger displacement of the centre of mass of the limb (Carey, Hargreaves, & Goodale, 1996). In addition, to reach the far eccentricity targets on the left hand-side, a flexion of the elbow together with an additional movement of the shoulder (and potentially even the trunk) was necessary as compared to a mere elbow extension for right-sided reaches. Hence, more muscle groups have to be recruited for the execution of contralateral movements, which is in turn associated with a higher energy expenditure (Happee & Van der Helm, 1995). Finally, movements to the right side are not only less biomechanically costly but are also executed (and usually initiated) much faster and more accurately than movements directed to the left of the body midline (Carey et al., 1996).

One might argue that our failure to observe clear averaging behaviour in Experiment 1 may be due to employing not only a much more liberal movement time criterion, but also a more liberal reaction time criterion (i.e. time available for movement planning). Previous studies by Chapman and colleagues that consistently found evidence for averaged movement trajectories constrained participants' reaction times to 325 ms after target presentation (e.g., Chapman et al., 2010a, 2010b). As our pilot studies indicated that most of our participants were unable to meet this reaction time criterion in most of the trials, we opted for a more generous reaction time limit of 500 ms. We think the reason that our participants needed more time for movement initiation can be attributed to the use of a slightly different setup. Specifically, we used a relatively narrow target presentation box placed on the table in front of participants instead of presenting pointing

targets on a vertical touch screen display. Consequently, participants had to control their movements in all three spatial dimensions and plan both the direction and extent of their reach. In contrast, movements toward a vertical screen only need to be controlled in direction; their extent is constrained by the placement of the screen. During movement planning the specification of amplitude is known to add to overall reaction times (Rosenbaum, 1980). Note that Wong and Haith (2017) also used longer RT limits in one of their experiments and found that it did not affect the overall movement pattern.

Based on the findings of Experiment 1, and the previous observation that averaging seems to only occur consistently if it is beneficial and increases the chances of task success (Haith et al., 2015), in Experiment 2 we aimed to investigate whether participants would flexibly adjust their strategies (between averaging and pre-selection) to the requirements of each trial. To test this, we set a much tighter movement time constraint such that far eccentricity targets could only be reached in time if pre-selected but not by employing an averaging strategy. In other words, at the far eccentricities, a pre-selection strategy would ensure a minimum error rate of about 50%. Therefore, the optimal strategy would be spatial averaging for near-eccentricity targets and pre-selection for far-eccentricity targets.

In contrast to previous studies, in which usually only correct trials are analysed while error trials are discarded, we were particularly interested in the type of errors participants would make, as these are informative with regard to the strategy employed. An optimal strategy would be revealed by an increased proportion of position errors (relative to time-out errors) in trials involving a far-eccentricity target (i.e. complete pre-selection of the wrong target), reflecting a shift from averaging to pre-selection, and/or in the case of time-out errors, a similar start direction for targets in the baseline and decision trials (i.e. reflecting a pre-selection and correction strategy). However, our data supported neither of those two possibilities. Overall, the number of position errors remained relatively small and did not vary systematically across different spatial conditions, suggesting that (most) participants over-applied an averaging strategy, and did not change their movement strategy based on the spatial separation of the presented target combination. For the large number of time-out errors, in particular in trials where the far-eccentricity light became the final pointing target, the analysis of start directions showed that participants seemed to initiate their movements toward an intermediate direction rather than committing toward the wrong target location initially.

In summary, the findings from Experiment 2 show that participants do not employ, or switch flexibly between, different optimal strategies that would increase their chances of success for each target combination. The few participants who did apply a pre-selection strategy did so

indiscriminately, and as a result they underperformed in the close-target trials (performing at or only slightly better than chance level), compared to the participants who over-applied the averaging strategy and underperformed in the far-target trials (performing well below chance level). These findings seem contradictory to a recent study by Wong and Haith (2017) where it was found that participants were able to flexibly adjust their strategies and employed averaging in trials in which there was sufficient time to do so but tended to avoid that strategy if time limits were more restrictive. The most plausible explanation for this discrepancy, in our view, is that the strategy that would be considered optimal in Wong and Haith's experiment was the same for each block of trials, while ours required participants to flexibly shift between averaging and pre-selection based on the difference in spatial separation of targets on each trial. Our results suggest participants are not assessing distance between targets and time constraints and applying an appropriate strategy that will maximise accuracy given these constraints. Rather, they appear to select and apply a motor strategy very broadly, for a set of trials, and do not shift to a different strategy easily. Also consistent with this interpretation, one of several interesting observations made by Wong and Haith (2017) was that participants' reach trajectories for single-target baseline trials were also affected by the timing constraints of the task, with movements to single targets being more direct during blocks with a tighter time limit than those without. In their experiment, unlike in the current ones, single-target and dual-target trials were randomly intermixed within a block. The fact that the single-target trials showed a smaller, but similar, effect of time pressure as dual-target trials is consistent with our interpretation that participants flexibly alter their movement strategies to meet instructions or task demands, but only broadly, for blocks of trials, rather than on a trial-by-trial basis.

In addition, our findings also nicely align with, and extend, the literature on visuomotor decision making under risk. Visuomotor decisions in these recent studies have not always been found to be as optimal as initially claimed, with deviations from optimality occurring as the task becomes more complex (e.g., Jarvstad et al., 2014; Neyedli & LeBlanc, 2017; Neyedli & Welsh, 2013, 2014; Wu et al., 2006). Interestingly, Neyedli and Welsh (2014), similar to us, found that participants are unable to select the optimal pointing location when penalties vary on a trial-by-trial basis, and yet they still found near optimal adjustments to trial-by-trial variations in the distance between reward and penalty region. These findings led them to suggest that spatial information may be processed more efficiently by the motor system than reward information. Here, we show that if spatial variations necessitate a change in movement strategy, rather than just a small adjustment in movement endpoint, participants show similar problems with flexibly selecting optimal strategies as they do for variations in cognitive parameters (such as reward).

In contrast to the decision-making under risk studies, we relied on the intrinsic reward of succeeding on each trial and of completing the experiment quickly. To motivate our participants, we informed them after each error that they would need to repeat that trial, and regularly providing a running tally of how many trials they had completed so far. Anecdotally, this appeared to keep our participants engaged in the task, but it is important to consider whether participants under-performed because of limitations in their abilities, or if they could have exhibited near-optimal reaching behaviour with stronger motivation. Other studies presenting a similar decision problem in different contexts have provided extrinsic rewards, such as money for correct trials (e.g., Morvan & Maloney, 2012), and found similar (suboptimal) results (see also Nowakowska, Clarke, Sahraie, & Hunt, 2019). Recently, James, Reuther, Angus, Clarke, and Hunt (2019) created a video-game-like environment to engage and motivate participants to make better decisions in an eye movement task where participants had to decide where to look to detect a target that could appear in one of two possible locations. Interestingly, they found gamification improved accuracy in the detection task, demonstrating its effectiveness in encouraging participants to try harder. Importantly, however, fixation strategies remained unchanged, even though much better detection performance could have been achieved with changes in fixation strategies. Applied to the current results, additional incentives may lead to faster and more accurate actions, but we do not think the failure to optimally adjust strategy across different target separations is due to a failure in motivation, and is therefore unlikely to change with higher incentives. More generally, it is intriguing that some components of task performance respond well to incentives, while the strategies seem less responsive.

Finally, there may be an alternative explanation for why most of our participants failed to successfully implement a pre-selection strategy in Experiment 2, namely the suggestion that averaging is not a deliberate strategy but the result of an automatic process where the motor system prepares competing action plans to all possible targets simultaneously (Cisek, 2012; Cisek & Kalaska, 2005). A number of studies have already shown that straight movement paths can be observed in the presence of multiple targets, calling the "automaticity" of these competing action plans into question (Chapman et al., 2010a; Haith et al., 2015; Wong & Haith, 2017). In addition, in most studies (including ours) there seem to always be a few participants who do not show any averaging behaviour at all (e.g., Stewart et al., 2013; Stewart et al., 2014; Wong & Haith, 2017) further challenging the notion that this behaviour is completely automatic. Nevertheless, we aimed to further check this possibility in Experiment 3 by encouraging participants to move to one of the two target lights. Here, we no longer observed averaging behaviour: participants were able to pre-select one target over the other, suggesting it is not an automatic process. Importantly,

Experiment 3 not only allowed us to test for the automaticity of averaging behaviour, but also allowed us to measure participants' movement preferences for certain target locations. We observed a clear preference for targets that are faster (and potentially easier) to reach, with a very clear negative correlation between how long a participant would take to reach a target and how likely they were to start a reach towards that target when given the choice.

Based on the outcomes of Experiment 2, it seems natural to argue that participants are poor in selecting and switching between different optimal movement strategies. However, the issue is that it all depends on how we chose to define "optimality". Clearly, participants do not seem to be optimal with respect to prioritising task success on a trial-by-trial basis, thereby minimising the number of trials they have to perform to complete the experiment. However, participants may aim to minimise different costs instead that they value more highly. Namely, our results, together with others from the motor control literature, indicate that participants may primarily aim to minimise the biomechanical and energetic costs of their movements and thus the immediate effort involved (Rosenbaum & Gaydos, 2008; Rosenbaum & Gregory, 2002). Studies by Cos and colleagues (Cos, Bélanger, & Cisek, 2011; Cos, Medleg, & Cisek, 2012) suggest that the biomechanical properties of a movement strongly determine which actions we select. In particular, humans seem to prefer shorter movements (Guigon, Baraduc, & Desmurget, 2007), movements that require the least metabolic energy (Goble, Zhang, Shimansky, Sharma, & Dounskaia, 2007) and movements that follow a path of maximum mobility (Cos et al., 2012). These preferences seem to be nicely confirmed by the finding of our Experiment 3, where we observed a strong preference of right hand-side over left hand-side targets and near eccentricity over far eccentricity targets. In fact, we found that even when participants were aware that they had to reach to all targets equally often over the course of the experiment, they had a strong tendency to select the target that could be reached in the shortest amount of time. Thus, while participants may have been poor in deliberately chosen movement strategies that would minimise the amount of trials performed over the course of the experiment (Experiment 2), they seemed to be extremely efficient in selecting movements that would reduce movement time and required effort on each trial.

Within the set of results presented here, in addition to having a bias towards the more motorically-expedient of the two targets, participants also exhibited a range of different movement strategies: spatial averaging, pre-selection with correction-in-flight, pre-selection with no correction, and behaviour that appeared to blend different strategies, partially averaging but with a bias to one target. We can therefore be confident that all of these strategies are options that humans can select between on any given trial. What we did not observe, however, was an optimal

selection between these strategies to maximise the likelihood of successfully reaching the target in a given trial. In light of the efficiency with which target selection minimizes motor effort, the failure to use task difficulty and time constraints to select the best strategy is even more surprising. Yet, this finding aligns with Clarke and Hunt's (2016) similarly surprising observation that people do not make optimal choices between trying to accomplish one goal and trying to accomplish two (where an optimal choice would be to pursue both goals only when you have sufficient time or resources), extending this conclusion to include choices about how to reach towards multiple possible targets under time constraint. Their tentative explanation is that under most circumstances, estimating the likelihood of success for all options and choosing the best outcome is computationally taxing, so people instead rely on shortcuts, like random selection. It is only when there are two targets, and the likelihood of success with each one is simple to estimate, that random selection can very clearly lead to worse performance than what could be achieved if an optimal (or even an approximation of optimal) strategy was adopted. As more targets are added, and the ability to estimate success becomes more difficult, the difference in outcome between optimal and random choices quickly diminishes. People do not seem to recognize when the situation has become simple enough that optimal choices can be easily estimated, and instead persist in using shortcuts.

Movement trajectories in the present experiment were not selected at random by participants, but clearly biased to the target that was easier to reach. It is important to note that biases and heuristics can boost performance above a completely random baseline, but without the effort required for computing an optimal strategy. This idea is formalised in a model of eye movements during visual search (Clarke, Green, Chantler, & Hunt, 2016), in which the sequence of eye movements was selected at random, but from a population of eye movements participants made from that region of the search area (see also Clarke, Stainer, Tatler, & Hunt, 2017). This random walk incorporates natural tendencies of the saccade system to, for example, make eye movements of a particular size and angle, and to saccade toward the centre more than other regions. Even though the model does not use any information about which regions might contain more or less information about the target, this model matched human search behaviour just as well as a model which does (Ackermann & Landy, 2013; Najemnik & Geisler, 2005). Applying this model's general framework to reach trajectories, the tendency to select closer targets on the right-hand side will, under most circumstances, produce better results than selecting targets at random. When faced with uncertainty about which light will be the target, participants may select at random from a population of possible movement options, and the population is biased towards the easier choices, and can be influenced by top-down factors such as instructions.

Returning back to our goal-keeper analogy, our results suggest that the goal-keeper's cognitive decision-making during a penalty kick may be far from optimal. That is, the goalkeeper is unlikely to optimally combine relevant information (about, for example, the other player, their own skill, the weather conditions, fatigue and injury) to make the best possible decision about where to stand, which way to dive, and when to start their movement. On the bright side, however, the kicker will have similar limitations, so both players are likely to choose a strategy somewhat at random, even though better choices exist that could improve their odds of success. Our results also suggest that these players can count on their own motor biases to help them rapidly select movements that play to their own strengths.

<u>Acknowledgement</u>

<u>Supplementary Material</u>

Data available at: https://zenodo.org/record/3604284

References

Ackermann, J. F., & Landy, M. S. (2013). Choice of saccade endpoint under risk. *Journal of Vision, 13*(3), 27-27.

Brenner, E., & Smeets, J. B. J. (1997). Fast responses of the human hand to changes in target position. *Journal of Motor Behavior, 29*(4), 297-310.

Brouwer, A.-M., Middelburg, T., Smeets, J. B., & Brenner, E. (2003). Hitting moving targets. *Experimental Brain Research, 152*(3), 368-375.

Carey, D. P., Hargreaves, E. L., & Goodale, M. A. (1996). Reaching to ipsilateral or contralateral targets: within-hemisphere visuomotor processing cannot explain hemispatial differences in motor control. *Experimental Brain Research, 112*(3), 496-504.

Chapman, C. S., Gallivan, J. P., Wood, D. K., Milne, J. L., Culham, J. C., & Goodale, M. A. (2010a). Reaching for the unknown: multiple target encoding and real-time decision-making in a rapid reach task. *Cognition, 116*(2), 168-176.

Chapman, C. S., Gallivan, J. P., Wood, D. K., Milne, J. L., Culham, J. C., & Goodale, M. A. (2010b). Short-term motor plasticity revealed in a visuomotor decision-making task. *Behavioural Brain Research, 214*(1), 130-134.

Cisek, P. (2012). Making decisions through a distributed consensus. *Current Opinion in Neurobiology, 22*(6), 927-936.

Cisek, P., & Kalaska, J. F. (2005). Neural correlates of reaching decisions in dorsal premotor cortex: specification of multiple direction choices and final selection of action. *Neuron, 45*(5), 801-814.

Cisek, P., & Kalaska, J. F. (2010). Neural mechanisms for interacting with a world full of action choices. *Annual Review of Neuroscience, 33*, 269-298.

Clarke, A. D. F., Green, P., Chantler, M. J., & Hunt, A. R. (2016). Human search for a target on a textured background is consistent with a stochastic model. *Journal of Vision, 16*(7), 4.

Clarke, A. D. F., & 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*(1), 64-74.

Clarke, A. D. F., Stainer, M. J., Tatler, B. W., & Hunt, A. R. (2017). The saccadic flow baseline: Accounting for image-independent biases in fixation behavior. *Journal of Vision, 17*(11), 12.

Cos, I., Bélanger, N., & Cisek, P. (2011). The influence of predicted arm biomechanics on decision-making. *Journal of Neurophysiology, 105*, 3022-3033.

Cos, I., Medleg, F., & Cisek, P. (2012). The modulatory influence of endpoint controllability on decisions between actions. *Journal of Neurophysiology, 108*, 1764-1780.

Day, B., & Lyon, I. (2000). Voluntary modification of automatic arm movements evoked by motion of a visual target. *Experimental Brain Research, 130*(2), 159-168.

Findlay, J. M., & Walker, R. (1999). A model of saccade generation based on parallel processing and competitive inhibition. *Behavioral and Brain Sciences, 22*(4), 661-674.

Gallivan, J. P., & Chapman, C. S. (2014). Three-dimensional reach trajectories as a probe of real-time decision-making between multiple competing targets. *Frontiers in Neuroscience, 8*, 215.

Gallivan, J. P., Chapman, C. S., Wolpert, D. M., & Flanagan, J. R. (2018). Decision-making in sensorimotor control. *Nature Reviews Neuroscience, 19*, 519-534.

Gallivan, J. P., Chapman, C. S., Wood, D. K., Milne, J. L., Ansari, D., Culham, J. C., & Goodale, M. A. (2011). One to four, and nothing more: nonconscious parallel

individuation of objects during action planning. *Psychological Science, 22*(6), 803-811.

Gallivan, J. P., Logan, L., Wolpert, D. M., & Flanagan, J. R. (2016). Parallel specification of competing sensorimotor control policies for alternative action options. *Nature Neuroscience, 19*(2), 320-328.

Goble, J. A., Zhang, Y., Shimansky, Y., Sharma, S., & Dounskaia, N. V. (2007). Directional biases reveal utilization of arm's biomechanical properties for optimization of motor behavior. *Journal of Neurophysiology, 98*(3), 1240-1252.

Guigon, E., Baraduc, P., & Desmurget, M. (2007). Computational motor control: redundancy and invariance. *Journal of Neurophysiology, 97*, 331-347.

Haith, A. M., Huberdeau, D. M., & Krakauer, J. W. (2015). Hedging your bets: intermediate movements as optimal behavior in the context of an incomplete decision. *PLoS Computational Biology, 11*(3), e1004171.

Happee, R., & Van der Helm, F. C. T. (1995). The control of shoulder muscles during goal directed movements, an inverse dynamic analysis. *Journal of Biomechanics, 28*(10), 1179-1191.

Hudson, T. E., Maloney, L. T., & Landy, M. S. (2007). Movement planning with probabilistic target information. *Journal of Neurophysiology, 98*, 3034-3046.

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.

Jarvstad, A., Hahn, U., Warren, P. A., & Rushton, S. K. (2014). Are perceptuo-motor decisions really more optimal than cognitive decisions? *Cognition, 130*(3), 397-416.

Kable, J. W., & Glimcher, P. W. (2009). The neurobiology of decision: consensus and controversy. *Neuron, 63*(6), 733-745.

Maloney, L. T., Trommershäuser, J., & Landy, M. S. (2007). Questions without words: A comparison between decision making under risk and movement planning under risk. In W. Gray (Ed.), *Integrated Models of Cognitive Systems* (Vol. 29, pp. 297-313). New York, NY: Oxford University Press.

Morvan, C., & Maloney, L. (2009). Suboptimal selection of initial saccade in a visual search task. *Journal of Vision, 9*(8), 444-444.

Morvan, C., & Maloney, L. T. (2012). Human visual search does not maximize the post-saccadic probability of identifying targets. *PLoS Computational Biology, 8*(2), e1002342.

Najemnik, J., & Geisler, W. S. (2005). Optimal eye movement strategies in visual search. *Nature, 434*, 387-391.

Nakayama, K., Song, J. H., Finkbeiner, M., & Caramazza, A. (2007). Hand trajectories reveal cognitive states. *Journal of Vision, 7*(9), 568-568.

Nashed, J. Y., Diamond, J. S., Gallivan, J. P., Wolpert, D. M., & Flanagan, J. R. (2017). Grip force when reaching with target uncertainty provides evidence for motor optimization over averaging. *Scientific reports, 7*(1), 11703.

Neyedli, H. F., & LeBlanc, K. A. (2017). The Role of Consistent Context in Rapid Movement Planning: Suboptimal Endpoint Adjustment to Changing Rewards. *Journal of Motor Behavior, 49*(6), 697-707.

Neyedli, H. F., & Welsh, T. N. (2013). Optimal weighting of costs and probabilities in a risky motor decision-making task requires experience. *Journal of Experimental Psychology: Human Perception and Performance, 39*(3), 638-645.

Neyedli, H. F., & Welsh, T. N. (2014). People are better at maximizing expected gain in a manual aiming task with rapidly changing probabilities than with rapidly changing payoffs. *Journal of Neurophysiology, 111*(5), 1016-1026.

Neyedli, H. F., & Welsh, T. N. (2015). Experience and net worth affects optimality in a motor decision task. *Motor Control, 19*(1), 75-89.

Nowakowska, A., Clarke, A. D., & Hunt, A. R. (2017). Human visual search behaviour is far from ideal. *Proceedings of the Royal Society B: Biological Sciences, 284*(1849), 20162767.

Nowakowska, A., Clarke, A. D. F., Sahraie, A., & Hunt, A. R. (2019). Practice-related changes in eye movement strategy in healthy adults with simulated hemianopia. *Neuropsychologia, 128*, 232-240.

Ramsay, J., & Silverman, B. (2005). *Functional Data Analysis*. New York: Springer.

Rosenbaum, D. A. (1980). Human movement initiation: specification of arm, direction, and extent. *Journal of Experimental Psychology: General, 109*(4), 444-474.

Rosenbaum, D. A., & Gaydos, M. J. (2008). A method for obtaining psychophysical estimates of movement costs. *Journal of Motor Behavior, 40*(1), 11-17.

Rosenbaum, D. A., & Gregory, R. W. (2002). Development of a method for measuring movement-related effort. *Experimental Brain Research, 142*(3), 365-373.

Soechting, J., & Lacquaniti, F. (1983). Modification of trajectory of a pointing movement in response to a change in target location. *Journal of Neurophysiology, 49*(2), 548-564.

Song, J.-H., & Nakayama, K. (2008). Target selection in visual search as revealed by movement trajectories. *Vision Research, 48*(7), 853-861.

Song, J.-H., & Nakayama, K. (2009). Hidden cognitive states revealed in choice reaching tasks. *Trends in Cognitive Sciences, 13*(8), 360-366.

Stewart, B. M., Baugh, L. A., Gallivan, J. P., & Flanagan, J. R. (2013). Simultaneous encoding of the direction and orientation of potential targets during reach planning: evidence of multiple competing reach plans. *Journal of Neurophysiology, 110*, 807-813.

Stewart, B. M., Gallivan, J. P., Baugh, L. A., & Flanagan, J. R. (2014). Motor, not visual, encoding of potential reach targets. *Current Biology, 24*(19), R953-R954.

Trommershäuser, J., Maloney, L. T., & Landy, M. S. (2003a). Statistical decision theory and the selection of rapid, goal-directed movements. *JOSA A, 20*(7), 1419-1433.

Trommershäuser, J., Maloney, L. T., & Landy, M. S. (2003b). Statistical decision theory and trade-offs in the control of motor response. *Spatial Vision, 16*(3), 255-275.

Trommershäuser, J., Maloney, L. T., & Landy, M. S. (2008). Decision making, movement planning and statistical decision theory. *Trends in Cognitive Sciences, 12*(8), 291-297.

Van der Stigchel, S., Meeter, M., & Theeuwes, J. (2006). Eye movement trajectories and what they tell us. *Neuroscience & Biobehavioral Reviews, 30*(5), 666-679.

Walker, R., McSorley, E., & Haggard, P. (2006). The control of saccade trajectories: Direction of curvature depends on prior knowledge of target location and saccade latency. *Perception & Psychophysics, 68*(1), 129-138.

Wispinski, N. J., Gallivan, J. P., & Chapman, C. S. (2018). Models, movements, and minds: bridging the gap between decision making and action. *Annals of the New York Academy of Sciences*, 1-22. doi:doi: 10.1111/nyas.13973

Wolpert, D. M., & Landy, M. S. (2012). Motor control is decision-making. *Current opinion in Neurobiology, 22*(6), 996-1003.

Wong, A. L., & Haith, A. M. (2017). Motor planning flexibly optimizes performance under uncertainty about task goals. *Nature communications, 8*, 14624.

Wu, S.-W., Trommershäuser, J., Maloney, L. T., & Landy, M. S. (2006). Limits to human movement planning in tasks with asymmetric gain landscapes. *Journal of Vision, 6*(1), 5-5.