

This article was downloaded by: [James T. Enns]

On: 12 January 2015, At: 08:42

Publisher: Routledge

Informa Ltd Registered in England and Wales Registered Number: 1072954
Registered office: Mortimer House, 37-41 Mortimer Street, London W1T 3JH,
UK



Visual Cognition

Publication details, including instructions for authors and subscription information:

<http://www.tandfonline.com/loi/pvis20>

Separating value from selection frequency in rapid reaching biases to visual targets

Craig S. Chapman^a, Jason P. Gallivan^b & James T. Enns^c

^a Faculty of Physical Education and Recreation, University of Alberta, Edmonton, Canada

^b Department of Psychology, Queen's University, Kingston, Canada

^c Department of Psychology, University of British Columbia, Vancouver, Canada

Published online: 06 Dec 2014.



CrossMark

[Click for updates](#)

To cite this article: Craig S. Chapman, Jason P. Gallivan & James T. Enns (2014): Separating value from selection frequency in rapid reaching biases to visual targets, *Visual Cognition*, DOI: [10.1080/13506285.2014.976604](https://doi.org/10.1080/13506285.2014.976604)

To link to this article: <http://dx.doi.org/10.1080/13506285.2014.976604>

PLEASE SCROLL DOWN FOR ARTICLE

Taylor & Francis makes every effort to ensure the accuracy of all the information (the "Content") contained in the publications on our platform. However, Taylor & Francis, our agents, and our licensors make no representations or warranties whatsoever as to the accuracy, completeness, or suitability for any purpose of the Content. Any opinions and views expressed in this publication are the opinions and views of the authors, and are not the views of or endorsed by Taylor & Francis. The accuracy of the Content should not be relied upon and should be independently verified with primary sources of information. Taylor and Francis shall not be liable for any losses, actions, claims, proceedings, demands, costs, expenses, damages, and other liabilities

whatsoever or howsoever caused arising directly or indirectly in connection with, in relation to or arising out of the use of the Content.

This article may be used for research, teaching, and private study purposes. Any substantial or systematic reproduction, redistribution, reselling, loan, sub-licensing, systematic supply, or distribution in any form to anyone is expressly forbidden. Terms & Conditions of access and use can be found at <http://www.tandfonline.com/page/terms-and-conditions>

Separating value from selection frequency in rapid reaching biases to visual targets

Craig S. Chapman¹, Jason P. Gallivan², and
James T. Enns³

¹Faculty of Physical Education and Recreation, University of Alberta,
Edmonton, Canada

²Department of Psychology, Queen's University, Kingston, Canada

³Department of Psychology, University of British Columbia, Vancouver,
Canada

(Received 25 May 2014; accepted 10 October 2014)

Stimuli associated with positive rewards in one task often receive preferential processing in a subsequent task, even when those associations are no longer relevant. Here we use a rapid reaching task to investigate these biases. In Experiment 1 we first replicated the learning procedure of Raymond and O'Brien (2009), for a set of arbitrary shapes that varied in value (positive, negative) and probability (20%, 80%). In a subsequent task, participants rapidly reached toward one of two shapes, except now the previously learned associations were irrelevant. As in the previous studies, we found significant reach biases toward shapes previously associated with a high probable, positive outcome. Unexpectedly, we also found a bias toward shapes previously associated with a low probable, negative outcome. Closer inspection of the learning task revealed a potential second factor that might account for these results; since a low probable negative shape was always paired with a high probable negative shape, it was selected with disproportionate frequency. To assess how selection frequency and reward value might both contribute to reaching biases we performed a second experiment. The results of this experiment at a group level replicated the reach-bias toward positively rewarding stimuli, but also revealed a separate bias toward stimuli that had been more frequently selected. At the level of individual participants, we observed a variety of preference profiles, with some participants biased primarily by reward value, others by frequency, and a few actually biased away from both highly rewarding and high frequency targets. These findings highlight that: (1) rapid reaching provides a sensitive readout of preferential processing; (2) target reward value and target selection frequency are

Please address all correspondence to Craig S. Chapman, Faculty of Physical Education and Recreation, University of Alberta, Edmonton, AB, T6G 2H9, Canada. E-mail: c.s.chapman@ualberta.ca

The authors would like to acknowledge K.L.T. and D.P. for help with data collection and J.L.C. for help during manuscript preparation. This work was supported by an NSERC Discovery Grant to J.T.E., a Killam Post-doctoral fellowship to C.S.C., and a Banting postdoctoral fellowship to J.P.G.

separate sources of bias; and (3) group-level analyses in complex decision-making tasks can obscure important and varied individual differences in preference profiles.

Keywords: Reward; Selection frequency; Rapid reaching; Action; Visual targets; Selective attention.

It is not surprising that we are drawn to objects that confer a reward. After all, we can maximize our expected gains by preferentially processing objects that are known to provide a net benefit. What is surprising, however, is the profound effect that rewards can have even when the objects they are associated with are removed from the rewarding context and their preferential processing becomes detrimental to performance. As an example, in a simple visual search task, if an item that was previously associated with a positive reward is a distractor in the current display, the search is slower (Anderson, Laurent, & Yantis, 2011). This is likely due, in part, to the fact that the eye is captured (Theeuwes & Belopolsky, 2012), or physically drawn toward items that have a history of being positively rewarding (Hickey & van Zoest, 2012), even though that reward contingency is transient and arbitrary. The effects of reward on visual attention to stimuli have been well documented (for reviews see: Anderson, 2013; Chelazzi, Perlatto, Santandrea, & Della Libera, 2013). Here we investigate how arbitrary reward associations affect biases in rapid reaching, allowing us to take advantage of the fact that the internal dynamics of choice can be revealed in the unfolding of a reach trajectory over time (Cisek, 2012; Freeman, Dale, & Farmer, 2011; Resulaj, Kiani, Wolpert, & Shadlen, 2009; Song & Nakayama, 2009; Wolpert & Landy, 2012).

The expected reward value of an object (the product of the gain/loss of a potential outcome and the probability that the outcome will occur), is not the only source of information that can result in its preferential processing. Recent perspectives on selective attention present compelling evidence for multiple routes by which the processing of a given stimulus can be biased (Awh, Belopolsky, & Theeuwes, 2012; Gottlieb, 2012). In examples from our own rapid-reaching work, we have shown that low-level visual properties such as target salience can bias trajectories (Wood et al., 2011), as can high-level associations like symbolic numerical value (Chapman et al., 2014). Given that more than one feature of a stimulus can contribute to its processing, experimentally it becomes problematic if the effects of multiple features are interpreted as arising from only a single source. For example, in one relevant set of background studies (O'Brien & Raymond, 2012; Painter, Kritikos, & Raymond, 2013; Raymond & O'Brien, 2009; Rutherford, O'Brien, & Raymond, 2010), researchers have developed a simple value learning task that establishes associations between arbitrary stimuli and a given expected reward value. The effects of these associations are then probed in a second unrelated and unrewarded task.

Despite all these experiments using the identical value learning paradigm, the conclusions have been as varied as the secondary tasks: (1) an attentional blink task (Raymond & O'Brien, 2009) showed processing advantages for positive stimuli during rapid processing and high probability stimuli during slower processing; (2) a cueing task (Rutherford et al., 2010) showed interference, or inhibition of return, induced by high probability positive stimuli; (3) a rapid recognition memory task showed faster recognition for high probability stimuli (O'Brien & Raymond, 2012); and (4) a reach-to-grasp task showed faster reaches toward high probability stimuli, but straighter reaches toward positive stimuli (Painter et al., 2013). While each of these results is noteworthy, taken together it is difficult to reconcile how the assumed manipulation of a stimuli's expected reward value (the learned combination of probability and point value) can generate such a varied collection of findings. We believe that one potential explanation that can account for this wide range of observations is that it is more than just a target's expected reward value that is being manipulated during learning.

An important and often overlooked source of stimulus processing bias is its "selection history" (Awh et al., 2012) or "attention for action" (Gottlieb, 2012). That is, the physical act of choosing or acting on a specific object is enough for it to be preferentially processed in subsequent encounters. For example, in two recent visual search studies (Belopolsky, Schreij, & Theeuwes, 2010; Theeuwes & Van der Burg, 2011) when a target was repeated across consecutive trials, its processing, and search performance, was facilitated—an effect known as intertrial priming (for a recent review, see: Kristjánsson & Campana, 2010). Importantly, in both visual search studies, it was argued that the bias generated by intertrial priming was separate from other sources of bias. Returning to the value learning task used by Raymond and colleagues (O'Brien & Raymond, 2012; Painter et al., 2013; Raymond & O'Brien, 2009; Rutherford et al., 2010), we noticed that there were differences not just in learned reward value, but also, critically, in selection history. That is, stimuli in the value learning task were always presented in pairs that shared the same valence but differed in their probability of payoff: a positive win-pair (e.g., one high, one low probability stimuli), a negative loss-pair and a neutral-pair. These pairings necessarily resulted in unique selection histories for each target. For instance, the low probability negative target was always paired with the high probability negative target and was therefore selected much more frequently.

In our first experiment, we employed the same value learning task as Raymond and colleagues (O'Brien & Raymond, 2012; Painter et al., 2013; Raymond & O'Brien, 2009; Rutherford et al., 2010) to examine, first, if rapid reach trajectories would be sensitive to learned expected reward value and, second, if the use of rapid reaching would allow us to more directly observe evidence for a bias generated by selection history. Our previous work (Chapman et al., 2010a, 2010b, 2014; Gallivan et al., 2011; Milne et al., 2013; Wood et al., 2011) and the work of others (Freeman et al., 2011; McKinstry, Dale, & Spivey,

2008; O’Hora, Dale, Piironen, & Connolly, 2013; Song & Nakayama, 2009; Spivey & Dale, 2006) has shown that reach trajectories (and other continuous measures of behaviour) can provide a highly sensitive read-out of stimulus processing biases. To anticipate our results, we did indeed find that the hand was drawn toward stimuli previously associated with high probability, positive rewards. However, we also found that, in some of the experimental conditions, the hand was drawn toward the low probability, loss stimulus. This suggests that selection history was indeed having an impact separate from expected reward value.

To properly isolate the biases generated from these two unique sources, we conducted a second experiment in which stimuli independently varied in point value (which was the same as expected reward value since probability was held constant at 1.0) and selection frequency. In looking for independent biasing factors we were also able to investigate an underappreciated aspect of stimulus processing: individual differences in the relative importance of each bias signal. That is, it is just as easy to imagine a participant who is more biased by an object’s selection history as it is to imagine a participant who is more biased by its reward value. This important point—that a participant’s preference profile is ultimately the best descriptor of their behaviour—was well demonstrated in recent work showing how apparent differences across disparate decision-making tasks were collapsed when the tasks were equated and individual differences were factored in (Jarvstad, Hahn, Rushton, & Warren, 2013). Thus, the goal of the current study was, first, to determine the extent to which selection and reward history constitute separable sources of bias and, second, to explore whether individuals had different preference profiles in how they weighted each of these sources of bias.

EXPERIMENT 1

In Experiment 1 (E1) we implemented the procedures of the Value Learning (VL) paradigm from Raymond and colleagues (O’Brien & Raymond, 2012; Painter et al., 2013; Raymond & O’Brien, 2009; Rutherford et al., 2010) in a reaching task and then examined how associations between shape and value, learned in this first VL phase, influenced the trajectories of subsequent rapid reaches (RR) performed toward the same shapes in a second phase.

Methods

Participants. Twenty-six right-handed (mean age 20.2 years, seven male) healthy young adults participated in E1 in exchange for course credit or money. In order to isolate only the fastest (i.e., most automatic) rapid reaching responses (see below for trial inclusion criteria), we imposed strict timing restrictions on

the rapid reaching performance (for review see: Gallivan & Chapman, 2014). Data from 5.1% of trials were eliminated due to recording errors and data from three participants was removed for failing to attain accurate performance on at least 50% of rapid reaching trials (e.g., a participant's data was removed if, after all trial exclusions explained below, they had less than 50% of their trials remaining). Accordingly, all analyses were conducted on the good trials from the remaining 23 participants (81.0% of trials from remaining participants were analyzed; see below for trial exclusion rates across all participants).

Apparatus and stimuli. Participants performed reach responses (recorded from a sensor attached to the tip of the right index finger at 200 Hz with OPTOTRAK) from a start button, positioned 10 cm from the front edge of a table at which they sat, to a 42-inch touch screen (HP LD4200tm) placed at a distance of 40 cm from the start button. On each trial, shape pairs consisting of two outline-shapes (denoted by a 4 pixel-wide black-outline on a white background) were shown on the screen centred 9 cm to the left and right of a central fixation cross. There were six possible shapes: a circle, a square, an "X", a six-point star, a cross, and a diamond (see Figure 1D for examples of the stimuli). The size of all shapes was equated such that the area inside the outline was equivalent to the area of the 2 cm diameter circle. Custom Matlab scripts using Psychtoolbox version 3 (Brainard, 1997; Kleiner, Brainard, & Pelli, 2007; Pelli, 1997) controlled all of the stimulus presentation and data collection.

Procedure. Participants were motivated in the task to accumulate points across 10 experimental blocks. These points could be traded-in at the end of the experiment for treats (chips, pop, or candy) that were located on a side table in the testing room, in line with the participants' field of view. The "price" of these treats was determined through pilot testing so that 100 points could be traded-in for a can of pop or a small bag of chips and 200 points could be traded-in for a candy bar. The average participant was able to earn about three treats throughout testing. Participants were made aware of the "price" of each available treat prior to initiating the experiment. However, they were not explicitly told how long the experiment would last, so we are confident that they remained motivated to accumulate points, even if they were approaching the experiment's end.

The session began with a practice Rapid Reaching (RR) block (40 trials) to familiarize participants with the timing of that task. This was followed by the experimental trials: a Value Learning (VL) phase (five blocks of 30 trials each), a Rapid Reaching (RR) phase (five blocks of 40 trials each), one more top-up VL block (in which the previously learned associations were refreshed), and finally five more RR blocks. Total points accumulated across the five VL blocks, the five RR blocks, the top-up VL block, and the final five RR blocks, were then traded-in for treats.

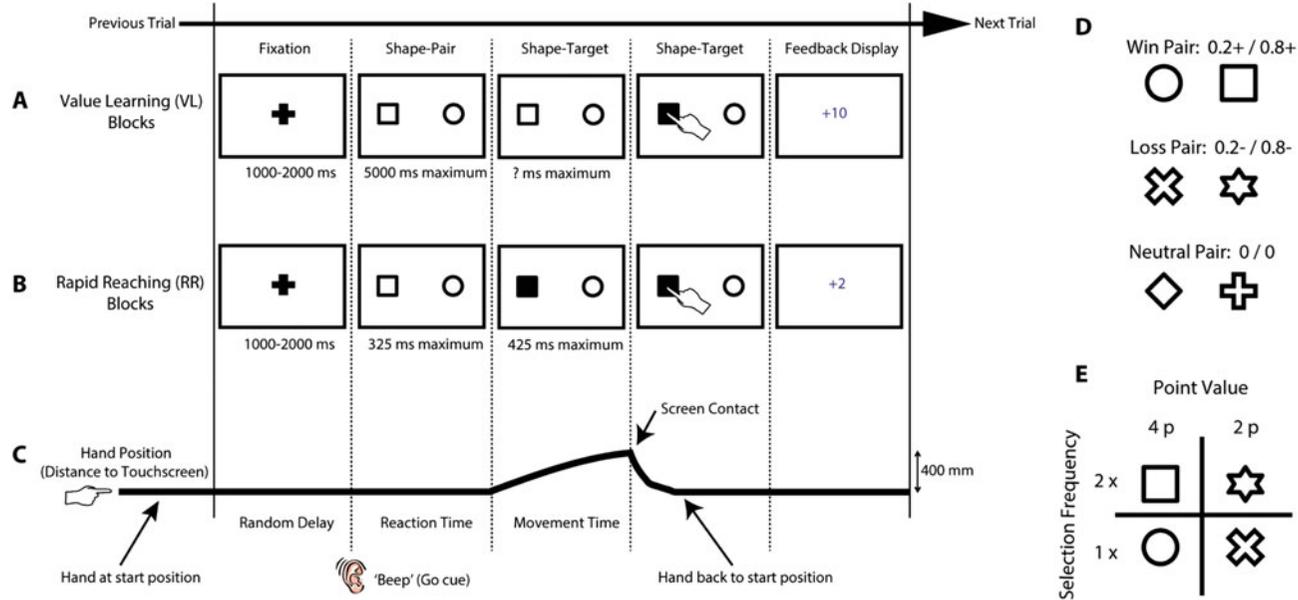


Figure 1. Experimental methods. (A) Representative trial in the Value Learning (VL) phase from E1. (B) Representative trial in the Rapid Reaching (RR) phase from Experiments 1 and 2. (C) Hand position, plotted as a function of events in A and B. C denotes the linkage between hand movement and trial progression. (D) Target shape pairs used in value-learning, following the procedure of Raymond and O'Brien (2009). (E) Target shapes used in Experiment 2. Shapes were randomly assigned either high or low point value and high or low selection frequency, leading to the 2 x 2 design shown. Shapes were presented in all possible pairings (six in total).

Value Learning (VL). Participants were instructed to take their time when making shape selections in the VL phase and were informed that the shapes would have the same probability and point value structure throughout the VL task. Prior to starting the VL task, participants were not given any information about the specific associations between shapes and their probability and point value, but instead were told to focus on maximizing their point score by learning which shapes yielded the highest gain. Each VL trial began when participants brought their hand within 2 cm of the start position (indicated with a sticker on the table top). This triggered the presentation of a small black fixation dot at the centre of the touch screen. After a variable delay (1–2 s), the shape pair appeared on the screen. Concurrent with this event, the fixation dot was removed and an auditory “beep” cue signalled participants to initiate their response. Trials with an anticipatory response (< 100 ms from beep, $< 1\%$ of trials) or a significantly delayed response (no reach initiated within 5 s of the beep, $< 1\%$ of trials) were terminated, an error message was displayed on the screen, and the resulting data from the trial was not analyzed. Trials ended when participants selected one of the two shapes (attempting to maximize the points accrued) by touching within 3 cm of its position. Once contacted, the selected shape was then filled-in (i.e., the outlined shape was filled with black) to indicate selection to the participant. At the end of each trial, participants saw their current accumulated score as well as the points accrued as a result of that single trial (see Figure 1A). For instance, if the selected shape resulted in a gain participants saw the value “+ 10” (appear in blue) which was added to the accumulated score. For a loss, the value “– 10” appeared on the screen (in red) and was deducted from the accumulated score. For a trial in which no points were awarded, “0” was shown on the screen and the accumulated score remained unchanged.

For each participant, each of the six shapes was randomly assigned one of six values following the structure of the value learning task in Raymond and O’Brien (2009). The values were 0.8+ (80% chance of receiving +10 points), 0.2+ (20% chance of receiving +10 points), 0.2– (20% chance of receiving –10 points), 0.8– (80% chance of receiving –10 points) and two shapes had neutral values (100% chance of receiving 0 points). As in Raymond and O’Brien (2009) we presented shapes in only three pairs (see Figure 1C): a win pair (0.8+ vs. 0.2+), a loss pair (0.2– vs. 0.8–), and a neutral pair (0 vs. 0). Each pair was counterbalanced for side of presentation (i.e., in the win pair, the 0.8+ shape appeared an equal number of times on the left and right). Each shape-pair was presented 10 times (five presentations each for a pair and its mirror image, in a randomized order) in each of the six VL blocks.

Rapid Reaching (RR). Participants were encouraged to make fast reaches in the RR phase since the only way to earn points was to complete a rapid reach toward the cued-target within the allotted time (each successful trial resulted in +2 points). The sequence of events for RR trials was identical to VL trials up to

and including the beep. After the beep in an RR trial, however, participants were instructed to initiate a reach movement towards the screen as quickly as possible. Reflecting this demand, trials with anticipatory responses (< 100 ms, “Too Early”, $< 1\%$ of trials) or responses slower than 325 ms (“Time Out”, 8.6% of trials) were terminated, error feedback was given, and the trials were not analyzed. Upon movement onset (determined by velocities that exceeded 20 mm/s with an acceleration of 20 mm/s² over four sampled time points) one of the two shapes in the shape-pair was filled-in black indicating that it was the cued target for that trial (see [Figure 1B](#), and Gallivan & Chapman, 2014, for detailed methods review). Our previous work has shown that by delaying target cuing until the moment the reach is initiated, movement planning biases can be directly observed in the early stages of the subsequent reach trajectory, providing a behavioural “read-out” of evolving cognitive processes (Chapman et al., 2010a, 2010b, 2014; Gallivan et al., 2011; Milne et al., 2013; Wood et al., 2011). The demanding timing constraints were also extended to the movement durations: participants had only 425 ms to complete the reach and touch the selected target once the reach had been initiated (if not, a “Too Slow” error was displayed). Reaches that contacted the display more than 3 cm offset from the cued target resulted in a “Miss” error being displayed and were not analyzed (6.8% of trials). We also removed trials in which the reach movement duration was excessively slow (> 850 ms, 1.7% of trials) and where the movement times were more than two standard deviations above each individual participant’s mean movement time (2.3% of trials). Note, this meant that we included in our analysis some trials in which the participant received a “Too Slow” error at completion of the trial.

Critically, unlike the VL task, shapes in the RR task did not predict reward value. In fact, participants were explicitly instructed to ignore target shape and instead focus on speed, since rapid and accurate reaches were the only way to maximize point gain. Accordingly, participants were penalized (-2 points) any time they made an anticipatory response, reacted too slowly, or missed the target. As we have shown previously (Gallivan & Chapman, 2014), the movement time criterion is by far the most difficult for participants to meet. As such, we elected to have reaches executed too slowly result in no points being awarded: giving movement time penalties could have resulted in a net negative score and a poor motivational state. As in the VL phase, after each RR trial the total accumulated score was shown, as was the score for the current trial (“+2” in blue, “0” in black or “-2” in red), which then adjusted the accumulated score.

In addition to the three shape-pairs shown during the VL phase, in the RR phase we also presented the two pairs that held the probability constant while valence changed. That is, in addition to the win, loss, and neutral pairs there was now also a high probability pair (0.8+ vs. 0.8-) and a low probability pair (0.2+ vs. 0.2-). Combined with the mirror reversals of each pair this led to 10 total possible display pairs. Since a reach toward each of these display pairs could

have the final target cued on the left or right, this resulted in 20 experimental conditions, each of which was repeated twice (in a randomized order) in each of the 10 experimental RR blocks.

Results

Value Learning (VL). Replicating past research with this learning task, participants' learning was proportional to the value and probability of reward associated with each shape pair. **Figure 2A** shows the probability that participants made the optimal choice in the VL phase (i.e., selected the shape with higher expected value), plotted according to experimental block, for each of the win, loss, and neutral pairs. For the win and loss pairs, optimal choice behaviour was quite good ($> 80\%$ for the last two blocks) and significantly above chance (see **Figure 2A**, note that 95% confidence intervals [CIs] are well above 50%), while for the neutral pair, choice performance hovered around chance (50%). The results of a Greenhouse-Geisser (GG) corrected 3×6 (pair-type \times block) repeated measures analysis of variance (RM-ANOVA) revealed significant main effects of both pair-type ($F(1.36, 29.97) = 40.69, p < .001$) and block ($F(3.60, 79.08) = 10.84, p < .001$) but no interaction between the factors ($p = .31$). Bonferroni corrected (to a family-wise $\alpha = 0.05$) post-hoc tests showed that all pair-types were significantly different from each other, with the most optimal choices being observed for the win pair (89.0%), followed by the loss pair (70.9%), and finally the neutral pair (51.8%). We followed up this significant block effect by testing for a linear trend across the six blocks for each of the three pair types. Here, both the win pair ($F(1, 22) = 24.30, p < .001$) and the loss pair ($F(1, 22) = 41.53, p < .001$) showed a significant increase in optimal choices across blocks (i.e., a learning effect) whereas the neutral pair did not ($p = .49$).

Rapid Reaching (RR). The reaching trajectories showed two main findings. First, shape pairs associated with high value and high probability outcomes resulted in significant positive biases in reaching. Second, the shape pair associated with a low probability win resulted unexpectedly in a significant negative bias. These observations were supported by the following analyses.

Reach trajectories were first normalized (across reach distance, see Chapman et al., 2010a; Gallivan & Chapman, 2014, for details) and averaged for each shape-pair and shape-target-location (i.e., cued-left or cued-right). **Figures 2B** and **2C** show the group average reach trajectories toward targets that were cued on the right for the low probability pair (0.2+ vs. 0.2-, **Figure 2B**) and the win pair (0.8+ vs. 0.2+, **Figure 2C**) when the optimal shape (0.2+ and 0.8+, respectively) was on the left (solid lines) or right (dashed lines). We index the degree of bias as the area between the solid and dashed curves (grey area in **Figures 2B** and **2C**) and present these areas in **Figure 2D**, in which positive values indicate a bias toward the optimal shape (i.e., optimal in the VL but not

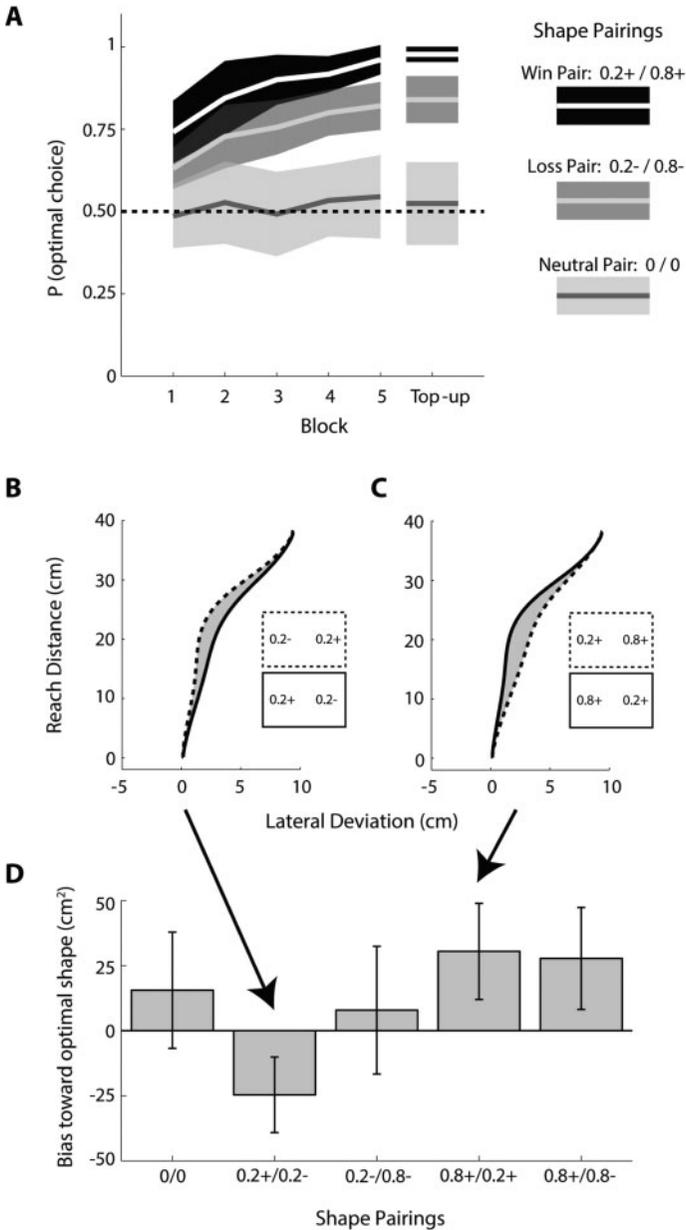


Figure 2. Experiment 1 results. (A) Mean proportion of optimal choices across the three value learning choice pairs (see legend in figure) in the Value Learning phase. The area around each curve shows the 95% CI about the respective block. (B) Reach trajectories made toward the right target in low-probability pair trials: 0.2+ vs. 0.2- trials (solid line) and 0.2- vs. 0.2+ trials (dashed line). Note that, contrary to the expected values, participants were significantly biased toward the negative (0.2-) shape. (C) Reach trajectories made toward the right target on win-pair trials: 0.8+ vs. 0.2+ trials (solid line) and 0.2+ vs. 0.8+ trials (dashed line). Consistent with predictions, on these trials participants were significantly biased toward the high probability positive shape. (D) Areas between trajectories for shape-pairs, plotted as a function of the trajectory bias toward the optimal shape. Arrows from B and C link the trajectory effects to the corresponding shape pairing biases shown in D. Error bars denote the 95% CI.

RR task). We also show the 95% confidence intervals around these bias values, which denote a significant bias anywhere the CIs *do not* include zero.

These graphical results measuring the effect of each bias relative to zero show that two pairs (win-pair: 0.8+ vs. 0.2+, and high-probability pair: 0.8+ vs. 0.8-) resulted in the predicted significant positive bias toward the optimal shape (0.8+ in both pairs) while one shape pair (low-probability pair: 0.2+ vs. 0.2-) resulted in a negative bias, or a bias toward the non-optimal shape (0.2-). The other pairs (loss-pair and neutral-pair) did not result in a significant reach bias. A significant difference across the five pairs was also revealed in a GG-corrected RM-ANOVA ($F(3.16, 69.42) = 5.44, p < .005$). Bonferroni-corrected (to a family-wise $\alpha = 0.05$) post-hoc comparisons revealed that this effect was driven by the negative bias in the low-probability pair being significantly different than all other pairs, except the loss pair. It is interesting that the loss pair did not result in a significant bias toward the 0.2- shape when that same shape did attract the hand in the low-probability pair. We speculate that this null effect might have something to do with the preferential processing of high probability stimuli found in other studies (O'Brien & Raymond, 2012; Painter et al., 2013; Raymond & O'Brien, 2009; Rutherford et al., 2010). While we find no direct evidence for a probability preference (that is, no main effect of a bias toward the 0.8 shapes regardless of valence) it is likely that probability alone does function as yet another route by which a stimulus's processing can be affected.

EXPERIMENT 2

The results of Experiment 1 demonstrated that learned reward contingencies transferred across tasks to influence rapid reach trajectories. As predicted, we found that the hand was drawn toward high-probability rewarding targets (e.g., the 0.8+ target). Interestingly, we also found that in the shape pair comparing the low-probability win target versus the low probability loss target (e.g., 0.2+ vs. 0.2-), the hand was significantly drawn toward the non-optimal loss target. This finding is completely unexpected based on the work that has used the same value learning task (O'Brien & Raymond, 2012; Painter et al., 2013; Raymond & O'Brien, 2009; Rutherford et al., 2010). As described in the introduction, one potential explanation for this finding was that the 0.2+ and 0.2- shapes differ in their selection history. As a preliminary test of this explanation, we asked whether individual differences in the selection history between the 0.2+ and 0.2- shapes in the E1 VL task were significantly related to individual trajectory biases in the E1 RR task. Specifically, we predicted that individuals who chose the 0.2- shape more frequently compared to the 0.2+ shape in the VL task (i.e., the difference in choice percentage) would be more biased toward its location when these two shapes appeared as a pair in the RR

task. The results of this analysis are depicted in Figure 3 and show that, as predicted, a significant relationship does exist ($r = -0.35, p < .05$).

To investigate more fully whether selection frequency operates as a separate source of bias from point value in guiding rapid reach trajectories, we performed Experiment 2. The task in this experiment was designed so that the learning of shape value was concurrent with the requirements to reach rapidly toward one of two shapes, randomly presented on each trial. For each participant, each of four shapes was assigned to one of four combinations of points and frequency of selection in a 2×2 design (see Figure 1E). That is, point values assigned to target shapes could be high (4 points) or low (2 points) and the frequency with which shapes were cued could be high (selected twice as often, or 2x) or low (selected half as often, or 1x).

Methods

Participants. A new group of 34 right-handed (mean age 20.3 years, 13 male) healthy young adults participated in exchange for course credit or money. As in the first experiment, data from three participants was removed for failing to attain accurate performance on at least 50% of rapid reaching trials and 1.8% of trials were eliminated due to recording errors. This left 31 participants whose

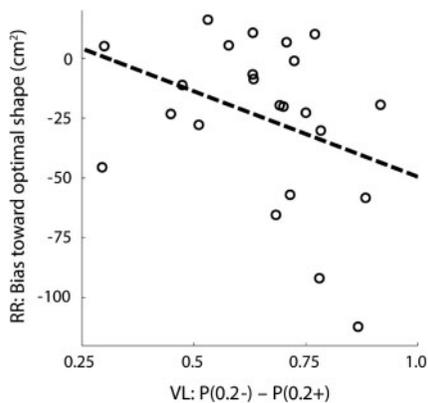


Figure 3. Testing selection frequency hypothesis in Experiment 1. Each circle represents one individual in E1, where the difference in choice frequency between the 0.2- shape (better choice in the loss pair) and the 0.2+ shape (worse choice in the win pair) in the VL task (x-axis) is plotted against the bias toward the optimal choice in the low-probability pair (0.2+/0.2-) in the RR task (y-axis). Note, for RR bias we plot a bias toward the 0.2- shape as negative, since 0.2- is the non-optimal shape (matches Figure 2D). The dashed black line indicates the best fitting straight line and indicates there is a significant correlation between the VL choice frequency difference and the subsequent RR bias: the more a participant chose the 0.2- shape relative to the 0.2+ shape during VL, the more they were biased toward the 0.2- shape during RR.

good trial data was analyzed (86.7% of trials from remaining participants were analyzed; see below for trial exclusion rates across all participants).

Apparatus and stimuli. The experimental setup was identical to E1 except that there were only four possible shapes presented: a circle, a square, a six-point star, or an “X” (see [Figure 1E](#)).

Procedure. Participants in Experiment 2 aimed to accumulate points across nine experimental blocks of Rapid Reaching (64 trials per block) following one practice block (the points were not accumulated and the data were not analyzed from this practice block). All points accumulated across the nine experimental blocks could be traded-in for treats as in E1.

The presentation and timing of trials in Experiment 2 was identical to the RR trials in Experiment 1. That is, participants had to initiate a reach toward a display containing two shapes within 325 ms and, upon movement onset, one of the two shapes in the shape-pair was filled-in black indicating that it was the cued target for that trial (see [Figure 1B](#)). To disentangle the effects of point value from selection frequency, each of the shapes was randomly assigned one of four combinations of points and frequency (see above) and all possible pairs of the four shapes were presented, leading to six different shape-pairs. In addition, for completeness, displays containing identical shapes were also presented (i.e., trials in which a square would appear on both the left and right sides of the display), but data from these pairs was not analyzed. Combined with the mirror reversals of the six analyzed, non-identical pairs (i.e., 12 display pairs), this led to 16 total display pairs. Since a reach toward each of these display pairs could have either shape cued as the final target (i.e., the left or right shape for any pair), this resulted in 32 total experimental conditions. Conditions in which the final target was a high-frequency shape were repeated 24 times across the nine experimental blocks whereas conditions in which the final target was a low-frequency shape were repeated 12 times across the nine experimental blocks. This resulted in a total of 576 trials across the nine experimental blocks. The points values associated with a shape (4 or 2 points) were displayed on the screen to participants only if they successfully completed the reach toward its location within the timing and accuracy constraints. Reaches not meeting these criteria (2.3% “Too Early”, 8.5% “Time Out”, 4.6% “Miss”) received the same error feedback and point penalties as described in the RR task of Experiment 1. In Experiment 2, we only eliminated “Too Slow” trials that exceeded 850 ms (1.3%).

Despite its sensitivity in revealing biases (Gallivan & Chapman, 2014), it is important to note a limitation of the cued-choice, rapid reaching task that we implemented. By enforcing that every shape appear in a pair with every other shape in our task, the manipulation of selection history necessarily confounds the number of times a given shape was acted on with the probability that a given

shape would be cued as the target on every trial. That is, high-frequency shapes were acted on more times and were also more likely to be cued as the final target on any trial. Since the probability of target selection has been shown previously to result in reach bias (Chapman et al., 2010a; Hudson, Maloney, & Landy, 2007) this means that any effects we find due to selection history are likely to be generated from a composite of factors. Notably, however, this restriction does not in any way limit our ability to search for separable influences from reward and selection history, it only makes our interpretation of selection history results more complex.

Results

Reach trajectories. At a group level, the results showed that reaching trajectories were biased by both reward value and by selection frequency. Our analysis of rapid reach data toward the six critical types of shape-pair trials was based on a comparison of three models of reaching bias. These models are illustrated in Figure 4A. The Points model (Figure 4A, top) assumes that only target value contributes to reaching biases, such that any pair containing a point difference should result in a significant bias toward the high-point shape. The Frequency model (Figure 4A, middle) predicts that the reaching bias should be driven exclusively by a difference in selection frequency between the two shapes. Thus, pairs containing a high-frequency and low-frequency shape should result in a significant bias toward the high-frequency shape. Finally, the Points plus Frequency model (Figure 4A, bottom) predicts that the reaching bias should be driven by the combined influence of points and frequency. Under this model, the one shape pair that differs along both dimensions will have the largest bias toward the optimal shape (leftmost bar in bottom panel of Figure 4A), whereas the one shape pair that has points and frequency working in opposing directions will result in no significant bias (rightmost bar in bottom panel of Figure 4A). This model further predicts that the remaining pairs that differ in only one of the two dimensions will have a small bias toward the optimal shape.

After conducting an omnibus GG-corrected RM-ANOVA testing for the effect of the six shape pairs on reaching area, we assessed the fit of each of these models by conducting three planned linear contrasts. This was done using weights that corresponded to the models depicted in Figure 4A. For the omnibus RM-ANOVA the main effect of pair on reach area trended towards statistical significance ($F(2.79,83.73) = 2.28, p = .09$), and was largely driven by a significant bias (see 95% CIs on first bar in Figure 4D) toward the high point, high frequency shape when paired with the low point, low frequency shape. All other pairs failed to generate a significant bias. Of note, we also looked to see how the associations between a given shape and its reward or frequency developed across time. To do so, we conducted the same RM-ANOVA with reach pair as a factor, but included time as a second factor with two levels,

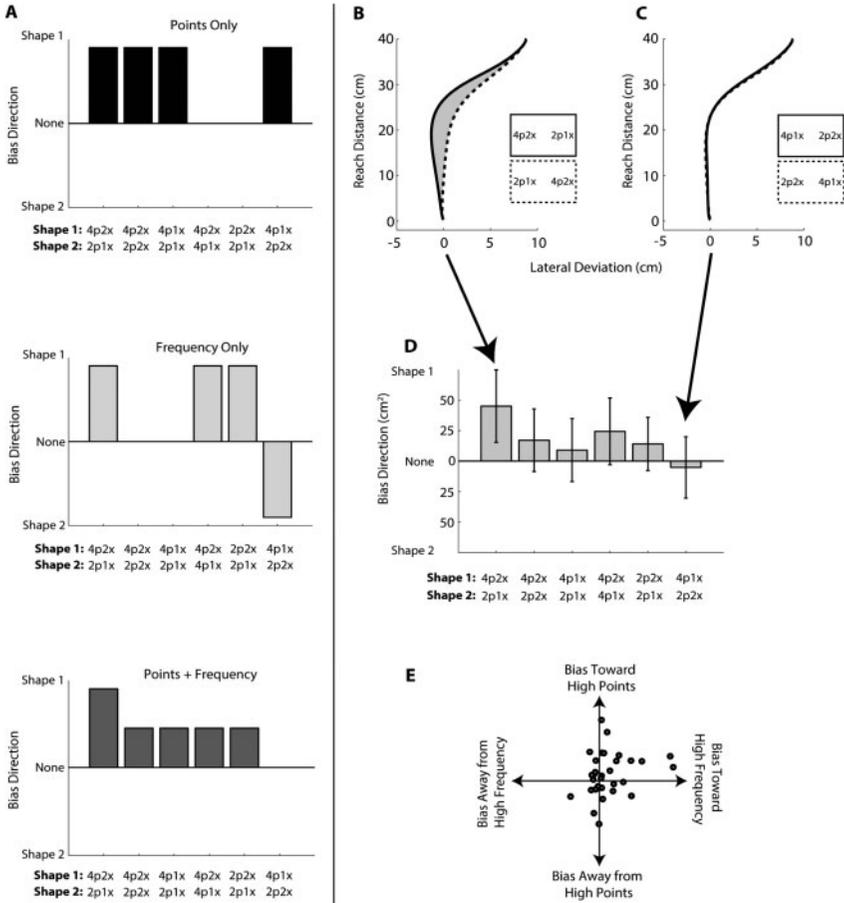


Figure 4. Experiment 2 results. (A) Different models tested with the Experiment 2 data. At top, the “Points Only” model predicts a bias toward a high point shape any time it is paired with a low point shape. At middle, the “Frequency Only” model predicts a bias toward a high frequency shape any time it is paired with a low frequency shape. At bottom, the “Points + Frequency” model is a simple addition of the Points Only and the Frequency Only models and predicts the strongest effect when the high point, high frequency shape is paired with a low point, low frequency shape (bar at far left) and a cancellation effect when the points and frequency values are in direct conflict (bar at far right). (B) Reach trajectories made toward the right target on trials in which the high point, high frequency shape (4p2x) was presented with the low point, low frequency shape (2p1x), when the 4p2x shape was presented on the left (dashed line) or right (solid line). Note that, as predicted by all three models in A, participants show a significant bias toward the high probability, high frequency shape. (C) Reach trajectories made toward the right target on trials in which the high point, low frequency shape (4p1x) was presented with the low point, high frequency shape (2p2x), when the 2p2x shape was presented on the left (dashed line) or right (solid line). Note that, as exclusively predicted by the Points + Frequency model in A, the opposing effects of points and frequency cancel on these trials and participants show no trajectory bias. (D) Areas between trajectories for all possible shape pairings (six in total), plotted as a function of the trajectory bias toward Shape 1 versus Shape 2. Arrows from B and C link the trajectory effects to the corresponding shape pairing biases shown in D. Error bars denote the 95% CI. While only the 4p2x vs. 2p1x shape pairing generates a significant bias (i.e., greater than zero), the remaining results are best accounted for by the Points + Frequency model. (E) Single-subject preference profiles estimated along two dimensions: Bias_H toward and away from high point shapes (y-axis) and Bias_F toward and away from high frequency shapes (x-axis). The wide range of individual preference profiles—from some participants who are biased almost exclusively toward high point or high frequency shapes to others who appear to be biased away from both high point and high frequency shapes—accounts for the variability observed in the group mean trajectories (shown in D).

grouping the first four blocks (including the practice block, which, other than here was not analyzed) into an Early bin and the last four blocks (6–10) into a Late bin. While there was an expected main effect of time with more bias developing as the learned associations got stronger (Early bias, $M = 9.73 \text{ cm}^2$, $SD = 8.29 \text{ cm}^2$; Late Bias, $M = 22.62 \text{ cm}^2$, $SD = 9.01 \text{ cm}^2$; $F(1,30) = 7.48$, $p < .01$) there was no interaction between time and pair ($F < 1$, $p > .5$). Importantly, this indicates that while participants did become more biased by the shapes in a pair as they learned the associations, this learning did not differentially affect the points versus frequency biases observed. However, as explained above, we had *a-priori* reasons to conduct linear contrasts testing each of the models depicted in Figure 4A. The linear contrast testing the Points model was not significant ($F(1,150) = 0.08$, $p = .78$, $\eta^2 = 0.007$) whereas the contrasts testing both the Frequency ($F(1,150) = 6.63$, $p < .05$, $\eta^2 = 0.58$) and Points + Frequency ($F(1,150) = 10.22$, $p < .005$, $\eta^2 = 0.898$) models were significant. Based on the fact that the Points + Frequency model was able to explain 30% more variance than the Frequency alone model, at a group level, our data is clearly best accounted by both points and frequency having independent and additive effects on reach biasing.

Individual preference profiles. Analyzing the group level effects indicates that both a target's associated point value and its selection frequency separately contribute to the reach biasing observed toward its location; we were also interested in how well this group level analysis actually captured the performance of each individual. That is, what appears to be two separable sources of bias at the group level might actually result from taking the average of two separate groups, each of which is influenced by only one biasing factor. Here we investigated whether individuals in Experiment 2 weighted each biasing factor approximately equally—as suggested by the group data—or, alternatively, whether individuals showed subject-specific idiosyncrasies in how they weighted both factors. That is, were the reaches of some individuals more guided by points while others were more guided by frequency? The evidence indicates that there was considerable variability in the biases of individual participants and, as a result, this suggests that the group level analysis does a poor job of characterizing the complexity of bias that occurs at the level of each individual.

To examine this question, we first generated an estimate of each individual's preference profile by fitting their reach area data, across the six pairs in Experiment 2, with the following linear model:

$$Pair * Bias = ReachArea \quad (1)$$

Here, *Pair* is a 6×2 matrix representing, for each of the six pairs (as ordered in Figures 4A and 4D), in column 1 how that pair differs in points (from Figure 4, Shape 1 vs. Shape 2 with: 1 = a pair with high points vs. low points; 0 = no

difference in points) and in column 2 how that pair differs in frequency (from Figure 4, Shape 1 vs. Shape 2 with: 1 = high frequency vs. low frequency; 0 = no difference in frequency; -1 = low frequency vs. high frequency):

$$Pair = \begin{bmatrix} 1 & 1 \\ 1 & 0 \\ 1 & 0 \\ 0 & 1 \\ 0 & 1 \\ 1 & -1 \end{bmatrix}$$

ReachArea in Equation 1 refers to the reach area data for a single participant across the six pairs. To allow for a comparison of preference profiles across participants, we normalized reach area across our participants by dividing each single participant's reach area data by the largest single reach area value found across all participants and all possible pairs. In effect, this normalization sets the maximum reach area we observed in the entire study to a value of 1, and all other values are relative to, and less than this value. We then solve for *Bias*, a 2×1 matrix, in which the first value is $Bias_P$, the bias toward points, and the second value is $Bias_F$, the bias toward frequency. For these bias values, positive values indicate a trajectory bias toward high point or high frequency targets while negative values indicate a bias away from high point or high frequency targets. We refer to each participant's set of bias values as their "preference profile".

The results of this analysis are plotted in Figure 4E, with individual subject's $Bias_P$ plotted on the y-axis against their $Bias_F$ values on the x-axis. This analysis shows that, in this experiment, there is quite a large range of preference profiles; ranging from participants whose reaches are being driven almost exclusively toward high frequency shapes, to those whose reaches are being driven almost exclusively toward high point shapes, to some individuals whose reaches show little to no effect of bias in accordance with these factors and, finally, to even a handful of participants who seem to actively move away from both high frequency and high point shapes. In general, it appears as if there is more spread along each bias axis than along the main diagonal, indicating that participants in Experiment 2 were more likely to be biased by a single factor (e.g., points or frequency) than by a combination of both. However, given the huge variety in preference profiles we observed, it is likely as inaccurate to conclude that all participants had a single preference profile with an equal weighting on points and frequency (as inferred from the group level analysis) as it is to conclude that there are only two preference profiles exhibited in our participant population, one exclusively weighting points, and the other exclusively weighting frequency. Rather, we interpret these results as evidence of the more general observation that: (1) group level analyses likely obscure important details regarding individual level biases, and (2) each individual's preference profile is a complex combination of multiple factors, each with its own relative weighting.

DISCUSSION

In Experiment 1 we found that rapid reach trajectories toward arbitrary shapes in a forced-choice pointing task were influenced by previously learned reward associations (see [Figure 2](#)). Specifically, participants' hands were physically drawn toward a shape that had previously been associated with a high probability of a positive outcome, even though that shape's reward history was no longer task relevant. This finding agrees with recent work showing that visual processing can be strongly biased by even transient or arbitrary reward associations (for reviews see: Anderson, 2013; Chelazzi et al., 2013), especially those that are positive and highly probable. Moreover, and more generally, this finding also further exemplifies the exquisite sensitivity of rapid motor behaviour, and reach actions in particular, as providing an observable index of the internal cognitive factors that bias stimulus processing (Chapman et al., 2010a; Cisek, 2012; Freeman et al., 2011; Song & Nakayama, 2009; Wolpert & Landy, 2012).

We attribute the sensitivity of the rapid reaching measure to our ability to detect the Experiment 1 result that reaches toward low probability shapes (0.2+ vs. 0.2-) are biased toward the non-optimal, low-probability negative shape. As described, the value learning procedure we adopted from Raymond and colleagues (O'Brien & Raymond, 2012; Painter et al., 2013; Raymond & O'Brien, 2009; Rutherford et al., 2010) resulted in the shape pairs differing along a dimension other than valence and probability of reward; they also differed in the number of times they had been selected as the target for action during learning. Since objects in the VL task were only presented in same-valence pairings, across the entire VL task the low-probable negative shape (selected 76.2% of the time in the loss-pair) was selected almost 7x as often as the low-probable positive shape (selected 11.0% of the time in the win-pair) and as [Figure 3](#) shows, individual differences in this selection frequency discrepancy were related to the strength of the subsequent rapid reach biases.

This reappraisal of the learning procedure prompted us to test whether selection frequency acts as a biasing factor independent from point value. In Experiment 2, participants performed a rapid-reach forced-choice task involving shapes that varied independently in both point value and selection frequency. The best account of the observed reach bias data, at a group level, was provided by a model specifying independent and additive biases toward both high value and high frequency shapes (see [Figures 4A–D](#)).

Taken together, the results from Experiments 1 and 2 suggest that an object's selection history (i.e., how often it is the target for action) is as important as its reward history (i.e., how often it results in a positive outcome) in determining whether or not it will be preferentially processed. This finding raises the important methodological point that many tasks exploring reward history have

also unknowingly been exploring selection history. That is, in tasks in which reward associations are learned via free choice (as in O'Brien & Raymond, 2012; Painter et al., 2013; Raymond & O'Brien, 2009; Rutherford et al., 2010) participants are motivated to choose the more rewarding option more often. As a result, subsequent learned association effects are likely due to both of these separate factors, but may only be attributed to one (e.g., reward). As noted earlier, it is important to realize that in our tasks, especially in Experiment 2, biases due to selection history (frequency) are likely themselves composed of multiple biasing factors. As an example, the number of times participants interacted with a shape in E2 was also related to the likelihood that shape would be cued as a target. Since the probability of target cueing has previously been shown to influence reaching (Chapman et al., 2010a; Hudson et al., 2007), this likely accounts for some of what we are referring to as a bias due to selection history. More broadly, this raises the important theoretical question of whether or not physically acting on an object is a different source of bias than the *likelihood* of acting on that object. Notably, while this does not detract from our overall finding that selection history and reward result in separate sources of bias, it does suggest that an important avenue for future investigations will be determining what specific factors are responsible for driving the selection history effect.

As important as it is to acknowledge that reward has, in previous work, been confounded with selection history, it is equally important to acknowledge that group level statistics are inadequate to capture how these different biasing factors contribute to each individual's decision process. That is, while the results of the group modelling in Experiment 2 suggested relatively equal, additive contributions of biases towards each point and frequency, the individual preference profiles (mapped in Figure 4E) showed a much different, more nuanced picture. Different from our group level conclusions, only a handful of participants appeared to weight points and frequency equally, whereas others, especially those with the largest amount of total bias, were strongly biased by only one factor. Importantly, it was not simply the case that one biasing factor consistently dominated the other, but rather, that some participants preferentially weighted point value while others preferentially weighted selection frequency.

These results provide a somewhat cautionary tale—one also recently echoed in the domain of decision making (Jarvstad et al., 2013)—that understanding a behavioural phenomenon requires understanding the individual and his/her context. We believe—despite our group level effects—that it would be the wrong interpretation of our Experiment 2 findings to say that point-value and selection-frequency bias rapid reach trajectories equally. Rather, what seems more accurate is to say that point-value and selection-frequency are independent sources of bias, and that their contribution to any specific individual's performance in the rapid reaching task is set by how that individual weights each bias source. Unfortunately, the current data does not allow us to make any inferences on what might be shaping each individual's preference profile. That

is, we cannot say why one individual was only biased by points, another only by frequency, and a third by a combination of both. However, these findings do suggest that a rich area for further investigation is to map an individual's preference profile in a given decision space while also collecting other measures of differentiation (e.g., personality scales) to explore the underlying factors which may result in a given preference profile.

This leads to the more general point that these experiments, like others (for reviews, see: Awh et al., 2012; Gottlieb, 2012), indicate that biased processing (or selective attention) is the outcome of a complex dynamical system with potentially multiple pathways. This is congruent with theoretical frameworks postulating that neural activity in a unified topographic map of objects charts the relative behavioural relevance assigned to each object (Awh et al., 2012; Bisley & Goldberg, 2010; Fecteau & Munoz, 2006; Gottlieb, 2012), and the allocation of attention and executed actions are direct read-outs of this topographic map (Baldauf & Deubel, 2010; Cisek & Kalaska, 2010). Within such frameworks, different biasing factors can be conceptualized as separate inputs into a location-relevance map, each contributing to the net activity and, thus, guiding the resulting behaviour. For our experiments, we can therefore understand point value and selection frequency (and perhaps its composite components) as separate sources of bias. As described above, the specific factors of reward history and selection history often work in the same direction, thereby having the same effect on behaviour. However, in some cases, as in the low probable negative shape in Experiment 1, and in the best fitting model in Experiment 2, the factors of reward and selection history can work against each other. Of course reward and selection history are but two of what may be a myriad of biasing factors: from low level stimulus salience, to top-down attentional sets, to the strength of remembered association(s), to the repetition of an action to a physical location in space; all of these factors and more will each contribute to the overall activity in the relevance map.

In many ways, this conceptualization mirrors the recent conclusions from several different task specific research domains. For example, in a recent review on repetition priming effects (which can be conceived of as a special case of selection history), Kristjánsson and Campana (2010) outline the disparate views on the origins of this effect, which range from facilitated perceptual processing, to biased attention, and to changes in response selection. As in our account of reaching biases, they discuss the varied entrenched ideas about *what* specifically is being primed, from the physical location of a stimulus, to its simple features, to the entire object, and to a complex configuration of the object in a specific context. Inasmuch as any of these concepts are separable, the answer to the question above, as derived from their research, is "all of the above". As Kristjánsson and Campana (2010) write: "... findings on priming suggest that it may be severely problematic to think of behavioural priming effects as reflecting the operation of a single perceptual process exerting its effects at a single level of

the perceptual hierarchy” (p. 10). Rather, as they argue: “... priming can be based on multiple representations, from more specific ones occurring in low-level visual areas to more general, feature-invariant representations occurring at higher levels of the visual hierarchy. The specific level of representation triggering the priming effect might then, in the end, depend on the specific stimulus, task and context in each particular case” (p. 15). Recasting their arguments into the language we have used here, the specific biasing factors resulting in an observed object preference likely originate from unique sources (e.g., low-level vision or high-level vision) and combine in unique ways for each stimulus, task and context.

If we understand preferential processing as the dynamic contribution to an object’s representation from multiple independent bias factors, then we must also assume that for each individual person, and/or in each behavioural context, the way these factors combine may vary. Just as one individual might have a preference to wear a comfortable old shirt, while another might prefer a newer shirt that is smarter looking, it is easy to envision that, at the lowest levels of decision making, each individual has their own preference profile that structures how a set of biasing factors are combined for a given task. Similarly, the same individual may also weight these two factors differently, depending on the occasion they are dressing for (e.g., recreation versus a wedding). Indeed, when generating separate estimates for each participant’s strongest bias in Experiment 2, either toward points and/or toward frequency, we observe a large range of preference profiles. Not only did participants vary in their total amount of bias (that is, some participants showed larger trajectory deviations than others), but, more importantly, they also varied in the relative strength of the contribution from each biasing factor.

Some objects will always be more likely than other objects to attract our attention and be selected as the targets of action. The important point is that this can occur for any number of reasons; these more attractive objects may be more visually salient, physically closer to the body, or fit our end use goals better than other objects around them. Here, we have shown that a rapid reaching task provides a unique window onto this process of biased processing, demonstrating that both an object’s learned reward value and its history of being selected can attract the hand towards its location. If these findings are representative, then they speak to the multiplicity of biasing factors that must be at play in dynamically adjusting the relative activations elicited by objects as they compete in parallel for our attention and, ultimately, our motor behaviour. At the same time, these findings reveal that while, on average, participants may be biased towards high point-value objects and frequently selected objects, the behaviour of the group may mask unique preference profiles at the level of each individual. Thus, it is just as important to acknowledge the multiple routes by which an object’s processing is biased as it is to recognize that every individual will combine these multiple bias factors in different ways. These findings suggest that

only by embracing both the complexity of the decision space, and the uniqueness of each individual in a given decision-oriented context, can we understand the influence of cognitive factors like reward and choice.

REFERENCES

- Anderson, B. A. (2013). A value-driven mechanism of attentional selection. *Journal of Vision*, *13*(3):7, 1–16. doi:[10.1167/13.3.7](https://doi.org/10.1167/13.3.7)
- Anderson, B. A., Laurent, P. A., & Yantis, S. (2011). Value-driven attentional capture. *Proceedings of the National Academy of Sciences of the United States of America*, *108*(25), 10367–10371. doi:[10.1073/pnas.1104047108](https://doi.org/10.1073/pnas.1104047108)
- Awh, E., Belopolsky, A. V., & Theeuwes, J. (2012). Top-down versus bottom-up attentional control: A failed theoretical dichotomy. *Trends in Cognitive Sciences*, *16*(8), 437–443.
- Baldauf, D., & Deubel, H. (2010). Attentional landscapes in reaching and grasping. *Vision Research*, *50*(11), 999–1013. doi:[10.1016/j.visres.2010.02.008](https://doi.org/10.1016/j.visres.2010.02.008)
- Belopolsky, A. V., Schreij, D., & Theeuwes, J. (2010). What is top-down about contingent capture? *Attention, Perception, & Psychophysics*, *72*(2), 326–341.
- Bisley, J. W., & Goldberg, M. E. (2010). Attention, intention, and priority in the parietal lobe. *Annual Review of Neuroscience*, *33*, 1–21. doi:[10.1146/annurev-neuro-060909-152823](https://doi.org/10.1146/annurev-neuro-060909-152823)
- Brainard, D. H. (1997). The psychophysics toolbox. *Spatial Vision*, *10*, 433–436.
- Chapman, C. S., Gallivan, J. P., Wood, D. K., Milne, J. L., Ansari, D., Culham, J. C., & Goodale, M. A. (2014). Counting on the motor system: Rapid action planning reveals the format- and magnitude-dependent extraction of numerical quantity. *Journal of Vision*, *14*(3), 30. doi:[10.1167/14.3.30](https://doi.org/10.1167/14.3.30)
- 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. doi:[10.1016/j.cognition.2010.04.008](https://doi.org/10.1016/j.cognition.2010.04.008)
- 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. doi:[10.1016/j.bbr.2010.05.012](https://doi.org/10.1016/j.bbr.2010.05.012)
- Chelazzi, L., Perlato, A., Santandrea, E., & Della Libera, C. (2013). Rewards teach visual selective attention. *Vision Research*, *85*, 58–72.
- Cisek, P. (2012). Making decisions through a distributed consensus. *Current Opinion in Neurobiology*, *22*(6), 927–936. doi:[10.1016/j.conb.2012.05.007](https://doi.org/10.1016/j.conb.2012.05.007)
- Cisek, P., & Kalaska, J. F. (2010). Neural mechanisms for interacting with a world full of action choices. *Annual Review of Neuroscience*, *33*, 269–298. doi:[10.1146/annurev.neuro.051508.135409](https://doi.org/10.1146/annurev.neuro.051508.135409)
- Fecteau, J. H., & Munoz, D. P. (2006). Saliency, relevance, and firing: A priority map for target selection. *Trends in Cognitive Sciences*, *10*(8), 382–390. doi:[10.1016/j.tics.2006.06.011](https://doi.org/10.1016/j.tics.2006.06.011)
- Freeman, J. B., Dale, R., & Farmer, T. A. (2011). Hand in motion reveals mind in motion. *Frontiers in Psychology*, *2*, 59. doi:[10.3389/fpsyg.2011.00059](https://doi.org/10.3389/fpsyg.2011.00059)
- 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. doi:[10.3389/fnins.2014.00215](https://doi.org/10.3389/fnins.2014.00215)
- 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. doi:[10.1177/0956797611408733](https://doi.org/10.1177/0956797611408733)
- Gottlieb, J. (2012). Attention, learning, and the value of information. *Neuron*, *76*(2), 281–295.

- Hickey, C., & van Zoest, W. (2012). Reward creates oculomotor salience. *Current Biology*, 22(7), R219–R220.
- Hudson, T. E., Maloney, L. T., & Landy, M. S. (2007). Movement planning with probabilistic target information. *Journal of Neurophysiology*, 98(5), 3034–3046. doi:10.1152/jn.00858.2007
- Jarvstad, A., Hahn, U., Rushton, S. K., & Warren, P. A. (2013). Perceptuo-motor, cognitive, and description-based decision-making seem equally good. *Proceedings of the National Academy of Sciences of the United States of America*, 110(40), 16271–16276. doi:10.1073/pnas.1300239110
- Kleiner, M., Brainard, D., & Pelli, D. (2007). What's new in psychtoolbox-3. *Perception*, 36(ECVP Abstract Supplement), 36(14), 1–1.
- Kristjánsson, Á., & Campana, G. (2010). Where perception meets memory: A review of repetition priming in visual search tasks. *Attention, Perception, & Psychophysics*, 72(1), 5–18.
- McKinstry, C., Dale, R., & Spivey, M. J. (2008). Action dynamics reveal parallel competition in decision making. *Psychological Science*, 19(1), 22–24. doi:10.1111/j.1467-9280.2008.02041.x
- Milne, J. L., Chapman, C. S., Gallivan, J. P., Wood, D. K., Culham, J. C., & Goodale, M. A. (2013). Connecting the dots: Object connectedness deceives perception but not movement planning. *Psychological Science*, 24(8), 1456–1465. doi:10.1177/0956797612473485
- O'Brien, J. L., & Raymond, J. E. (2012). Learned predictiveness speeds visual processing. *Psychological Science*, 23(4), 359–363. doi:10.1177/0956797611429800
- O'Hara, D., Dale, R., Piironen, P. T., & Connolly, F. (2013). Local dynamics in decision making: The evolution of preference within and across decisions. *Scientific Reports*, 3, 2210. doi:10.1038/srep02210
- Painter, D. R., Kritikos, A., & Raymond, J. E. (2013). Value learning modulates goal-directed actions. *The Quarterly Journal of Experimental Psychology*, 67(6), 1166–1175. doi:10.1080/17470218.2013.848913
- Pelli, D. G. (1997). The VideoToolbox software for visual psychophysics: Transforming numbers into movies. *Spatial Vision*, 10, 437–442.
- Raymond, J. E., & O'Brien, J. L. (2009). Selective visual attention and motivation: The consequences of value learning in an attentional blink task. *Psychological Science*, 20(8), 981–988. doi:10.1111/j.1467-9280.2009.02391.x
- Resulaj, A., Kiani, R., Wolpert, D. M., & Shadlen, M. N. (2009). Changes of mind in decision-making. *Nature*, 461(7261), 263–266.
- Rutherford, H. J., O'Brien, J. L., & Raymond, J. E. (2010). Value associations of irrelevant stimuli modify rapid visual orienting. *Psychonomic Bulletin & Review*, 17(4), 536–542.
- Song, J. H., & Nakayama, K. (2009). Hidden cognitive states revealed in choice reaching tasks. *Trends in Cognitive Sciences*, 13(8), 360–366. doi:10.1016/j.tics.2009.04.009
- Spivey, M. J., & Dale, R. (2006). Continuous dynamics in real-time cognition. *Current Directions in Psychological Science*, 15(5), 207–211. doi:10.1111/j.1467-8721.2006.00437.x
- Theeuwes, J., & Belopolsky, A. V. (2012). Reward grabs the eye: Oculomotor capture by rewarding stimuli. *Vision Research*, 74, 80–85.
- Theeuwes, J., & Van der Burg, E. (2011). On the limits of top-down control of visual selection. *Attention, Perception, & Psychophysics*, 73(7), 2092–2103.
- Wolpert, D. M., & Landy, M. S. (2012). Motor control is decision-making. *Current Opinion in Neurobiology*, 22(6), 996–1003.
- Wood, D. K., Gallivan, J. P., Chapman, C. S., Milne, J. L., Culham, J. C., & Goodale, M. A. (2011). Visual salience dominates early visuomotor competition in reaching behavior. *Journal of Vision*, 11(10):16, 1–11. doi:10.1167/11.10.16