

See discussions, stats, and author profiles for this publication at: <http://www.researchgate.net/publication/279066754>

The Snooze of Lose: Rapid Reaching Reveals That Losses Are Processed More Slowly Than Gains.

ARTICLE *in* JOURNAL OF EXPERIMENTAL PSYCHOLOGY GENERAL · JUNE 2015

Impact Factor: 5.5 · DOI: 10.1037/xge0000085 · Source: PubMed

DOWNLOADS

51

VIEWS

44

5 AUTHORS, INCLUDING:



Craig S Chapman

University of Alberta

37 PUBLICATIONS 272 CITATIONS

SEE PROFILE



Jason P Gallivan

Queen's University

35 PUBLICATIONS 486 CITATIONS

SEE PROFILE



Nathan Wispinski

University of Alberta

2 PUBLICATIONS 0 CITATIONS

SEE PROFILE



James Enns

University of British Columbia - Vancouver

252 PUBLICATIONS 6,716 CITATIONS

SEE PROFILE

Journal of Experimental Psychology: General

The Snooze of Lose: Rapid Reaching Reveals That Losses Are Processed More Slowly Than Gains

Craig S. Chapman, Jason P. Gallivan, Jeremy D. Wong, Nathan J. Wispinski, and James T. Enns

Online First Publication, June 22, 2015. <http://dx.doi.org/10.1037/xge0000085>

CITATION

Chapman, C. S., Gallivan, J. P., Wong, J. D., Wispinski, N. J., & Enns, J. T. (2015, June 22). The Snooze of Lose: Rapid Reaching Reveals That Losses Are Processed More Slowly Than Gains. *Journal of Experimental Psychology: General*. Advance online publication. <http://dx.doi.org/10.1037/xge0000085>

The Snooze of Lose: Rapid Reaching Reveals That Losses Are Processed More Slowly Than Gains

Craig S. Chapman
University of Alberta

Jason P. Gallivan
Queen's University

Jeremy D. Wong
Simon Fraser University

Nathan J. Wispinski
University of Alberta and University of British Columbia

James T. Enns
University of British Columbia

Decision making revolves around weighing potential gains and losses. Research in economic decision making has emphasized that humans exercise disproportionate caution when making explicit choices involving loss. By comparison, research in perceptual decision making has revealed a processing advantage for targets associated with potential gain, though the effects of loss have been explored less systematically. Here, we use a rapid reaching task to measure the relative sensitivity (Experiment 1) and the time course (Experiments 2 and 3) of rapid actions with regard to the reward valence and probability of targets. We show that targets linked to a high probability of gain influence actions about 100 ms earlier than targets associated with equivalent probability and value of loss. These findings are well accounted for by a model of stimulus response in which reward modulates the late, postpeak phase of the activity. We interpret our results within a neural framework of biased competition that is resolved in spatial maps of behavioral relevance. As implied by our model, all visual stimuli initially receive positive activation. Gain stimuli can build off of this initial activation when selected as a target, whereas loss stimuli have to overcome this initial activation in order to be avoided, accounting for the observed delay between valences. Our results bring clarity to the perceptual effects of losses versus gains and highlight the importance of considering the timeline of different biasing factors that influence decisions.

Keywords: decision making, reward, gain, loss, rapid reaching, stimulus response, relevance map

Supplemental materials: <http://dx.doi.org/10.1037/xge0000085.supp>

Making a decision between alternatives seems like a simple task. One only needs to list the potential gains and losses associated with each possible choice (i.e., a pros and cons list), and then assign each outcome a weighting based on its importance or desirability. From here, it is just a matter of simple arithmetic to deduce which outcome is the best (i.e., most gain and/or the least loss). However, applying this simple formulation to real decision making is incredibly complex. Not only can decisions range from

the trivial (“What shirt should I wear?”) to the significant (“Should I accept this job offer?”), but they can also bridge multiple domains (e.g., from motor to economic decisions) and operate across a huge range of time intervals (e.g., from milliseconds to years).

The goal of the present study was to directly test how recently acquired gain and loss associations affect the time course of rapid visuomotor choices. In pursuit of this objective, we necessarily integrate previous work from at least three decision science domains: economic decision making (EDM), perceptual decision making (PDM), and rapid visuomotor decision making (RVDM). Here, stemming from the seminal work of Kahneman and Tversky (1979), we use EDM to refer to problems that require participants to make explicit choices between options that vary in both the probability and amount of a payoff. For example, participants in an EDM task might be asked to report, via a questionnaire, their choice between two gambles: (a) a 50% chance of losing 25 dollars, and a 50% chance of winning 25 dollars; and (b) a 100% chance of winning 0 dollars. We use PDM to refer to the implicit (or nonconscious) competition and/or decisions between stimuli (and their associated actions) based on noisy sensory inputs. In PDM tasks, a participant might be asked to make one of two reach movements in the direction that is consistent with the overall

Craig S. Chapman, Faculty of Physical Education and Recreation, University of Alberta; Jason P. Gallivan, Department of Psychology, Centre for Neuroscience Studies, Queen's University; Jeremy D. Wong, Department of Biomedical Physiology and Kinesiology, Simon Fraser University; Nathan J. Wispinski, Faculty of Physical Education and Recreation, University of Alberta, and Department of Psychology, University of British Columbia; James T. Enns, Department of Psychology, University of British Columbia.

Correspondence concerning this article should be addressed to Craig S. Chapman, Faculty of Physical Education and Recreation, University of Alberta, Edmonton, AB, Canada T6G 2H9. E-mail: c.s.chapman@ualberta.ca

motion of a set of moving dots, some of which are moving randomly (e.g., Resulaj, Kiani, Wolpert, & Shadlen, 2009). Finally, we use RVDM to refer to tasks that require participants to react and move very quickly to touch presented visual choice options. As an example, participants might have less than 1 s to reach and accurately hit a positive target area while avoiding a negative target area (Trommershäuser, Maloney, & Landy, 2008).

Of particular interest in the current study is whether or not gains and losses are weighted equally in RVDM, and how gain and loss reward modulations evolve across time. Interestingly, each of the three decision domains outlined earlier has a different prediction regarding the effects of gains and losses. A hallmark of EDM tasks is the finding that participants are loss averse—that is, losses are weighed more heavily than gains (Kahneman & Tversky, 1979; Tversky & Kahneman, 1992). Even though the strength of loss aversion can be reduced when participants receive extensive feedback (for reviews, see Hertwig & Erev, 2009; Rakow & Newell, 2010; Yechiam & Hochman, 2013), it is not eliminated (Yechiam & Hochman, 2014). Conversely, recent work on how learned reward associations affect the perceptual processes in PDM models (Chelazzi et al., 2014; Chelazzi, Perlato, Santandrea, & Della Libera, 2013; Raymond & O'Brien, 2009) suggest that high-gain targets (especially those that pay off reliably) are preferentially processed relative to low-gain or loss targets. Thus, PDM would predict a gain–loss asymmetry that favors gain targets. Finally, recent RVDM work (Trommershäuser, Maloney, & Landy, 2003a, 2003b; Trommershäuser et al., 2008) has shown that participants equally weight loss and gain regions of a display when performing a rapid pointing task. Thus, this RVDM work predicts a relative equivalence of gain and loss targets.

Though some efforts have been made to unite these disparate decision science domains (Gottlieb, 2012; Jarvstad, Hahn, Rushton, & Warren, 2013; Summerfield, Tsetsos, Van Maanen, Dreher, & Ruff, 2012), two major differences are present across domains that make direct comparisons challenging and difficult: (a) loss-related processing has not received the same level of attention in each domain, and (b) the time scales over which the decisions occur within each domain drastically differ.

Regarding loss-related research, only in EDM research has loss been systematically investigated. Here, the results have been predominantly characterized by loss aversion behavior. Comparatively, PDM research has focused almost exclusively on gain-related effects, and shows that stimuli associated with high gain elicit preferential perceptual processing relative to low gain or neutral stimuli, which can result in both more and less efficient behavior (for reviews, see Anderson, Laurent, & Yantis, 2013; Chelazzi et al., 2013). The PDM work that has investigated loss processing has produced mixed results. In three studies by Raymond and colleagues (O'Brien & Raymond, 2012; Raymond & O'Brien, 2009; Rutherford, O'Brien, & Raymond, 2010), participants first learned to associate faces with specific monetary outcomes that varied in both probability (high or low likelihood of receiving a payoff) and payoff amount (gain or loss). Similar to some work on gain effects (Anderson, Laurent, & Yantis, 2011; Della Libera & Chelazzi, 2009), when the participants then performed a subsequent task using these same faces the researchers found that the previously learned association had lingering effects on performance, despite the fact that the former values were not relevant to this second task. Across the three studies, the results

from Raymond and colleagues suggest that either only the reward-related targets, or only the highly probable targets, would be preferentially processed. In a particularly relevant extension of the Raymond work, researchers recently showed in a reach to grasp task (following value learning) that movements were affected by gain and loss associations, moving more efficiently (of special relevance to the current work, in a straighter path) toward gain relative to loss targets (Painter, Kritikos, & Raymond, 2014). As discussed in our motivations section, the current study seeks to extend this work on grasping to examine the specific timeline of gain- and loss-related effects in a rapid pointing task. Finally, RVDM is itself a relatively small field, and as such, there are not many studies specifically comparing the effects of gain and loss. However, as mentioned earlier, the majority of studies seem to suggest that losses and gains are equally weighted (Trommershäuser et al., 2003a, 2003b, 2008), with the caveat that true optimal loss avoidance only comes with experience (Neyedli & Welsh, 2013, 2015).

Regarding the profound differences in time scales across the decision science domains, it is important to recognize that decisions in most EDM tasks operate on the order of many seconds to minutes. Comparatively, decisions in PDM tasks are faster, occurring on the order of 1 to 2 s. Finally, decisions in RVDM tasks are the fastest, with decisions often being initiated within half a second. The total time taken for a decision is naturally important, because it is known that decisions over short time scales can be very different from those arrived at over long deliberation. For example, when considered from the perspective of Kahneman's (2003, 2011) tripartite characterization of decisions ranging from “perception” to “fast thinking” to “slow thinking,” our RVDM task operates on the earliest “perception” responses, and as such, we cannot expect our findings to necessarily hold for PDM or EDM decisions, where more time is made available to deliberate. Here we use the term “perception” loosely, as we believe all of these processes reflect an integrated neural system designed to produce efficient actions. Thus, the boundaries between perception, thinking, and acting are likely artificial, but here provide useful terminology for delineating the different time scales of decision making. Here, we use our RVDM task to provide a comprehensive description of the effects of time on gain and loss processing for the earliest ~500 ms “perception” time window. In doing so, we aim to demonstrate the initial stages of how every decision must ultimately unfold, and thus provide the basis of a temporal framework, which can be extended across the different decision science domains.

Motivation

Inspired by this comparison between decision domains, our first aim for the present study was to test the weightings assigned to gains versus losses in a RVDM task. Previous research has shown straighter reaches toward high-gain targets in a grasping task, indicating preferential processing based on gain magnitude (Painter et al., 2014). This is consistent with previous work using eye movements, showing that when a previously positively rewarded target becomes a distractor, on the subsequent trial, participants' gaze is physically drawn toward the gain distractor (Hickey & van Zoest, 2012). Here, we seek to characterize the effects of gains and losses on rapid reach (RR) trajectories.

Our second aim was to test for potential asymmetries of gain and loss processing over time. Briefly, two choice options (with different learned gain or loss associations) are presented to the participant, who then must rapidly reach to touch either the target that is cued for them (Experiment 1) or the target of their choice (Experiments 2 and 3). We vary the processing time by either giving participants less time to view targets (Experiment 2; in the extreme, targets are not presented until after movement onset) or more time to view the targets (Experiment 3; on half the trials, participants get a preview of the targets before responding). We record their RR trajectories and use deviations away from a straight-hand path as indicative of competitive processing for the two choice options (for review of this methodology, see Gallivan & Chapman, 2014). By studying the evolving output of the perceptual-motor pipeline, one can measure the critical “end state” of the decision-making process (Cisek, 2012; Cisek & Kalaska, 2010; Wolpert & Landy, 2012), with the details of the movement trajectories providing a window into moment-by-moment decision-making dynamics (Dshemuchadse, Scherbaum, & Goshcke, 2013; Freeman, Dale, & Farmer, 2011; Gallivan & Chapman, 2014; McKinstry, Dale, & Spivey, 2008; O’Hora, Dale, Piironen, & Connolly, 2013; Scherbaum, Dshemuchadse, Fischer, & Goshcke, 2010; Song & Nakayama, 2009; Spivey & Dale, 2006).

The third theoretical aim of this study was to elucidate how gain- and loss-processing timelines compare with those of known neural responses. Specifically, neurophysiological recordings (Leathers & Olson, 2012; Pastor-Bernier & Cisek, 2011; Peck, Jangraw, Suzuki, Efem, & Gottlieb, 2009), electroencephalography (Kiss, Driver, & Eimer, 2009), and functional MRI (fMRI; Serences, 2008) show that reward associations modulate early sensory responses to visual stimuli. Thus, we sought to develop a neurally constrained (e.g., Purcell et al., 2010) descriptive model that defines decision making as a process that accumulates sensory evidence toward a decision threshold. In principle, this is an extension of the very successful class of models, known as drift diffusion models (DDMs), that have been shown to accurately predict human behavior (for review, see Bogacz, Brown, Moehlis, Holmes, & Cohen, 2006), and that are also a good candidate for the processes underlying the neural mechanisms of perceptual decisions (e.g., Gold & Shadlen, 2007; Heekeren, Marrett, & Ungerleider, 2008). Here we focus on what reward-related changes might be occurring at the very earliest sensory responses, and whether or not these changes can account for the behavior we observe in our RVDM task.

Experiment 1

In Experiment 1, we adopted the paradigm of Raymond and colleagues (O’Brien & Raymond, 2012; Painter et al., 2014; Raymond & O’Brien, 2009; Rutherford et al., 2010) and tested how associations between shape and value that were learned in the first phase of the study (value learning [VL] phase) affected the trajectory of a subsequent RR toward the same shapes in a second phase (RR phase). Critically, this allowed us to measure the carry-over effects of value on rapid visuomotor processing in a context in which this learned value was task irrelevant (and in fact, counterproductive to the RR). As such, any systematic value-based deviations toward or away from a shape provide strong evidence for

automatic value processing. Moreover, the direction and magnitude of these deviations could be used diagnostically to assess gain and loss processing.

Method

Participants. Thirty-one right-handed (mean age = 22.9 years; 20 female) healthy young adults participated in exchange for course credit or money. In order to isolate only the fastest (i.e., most automatic) RR responses, we imposed strict timing restrictions on the RR performance. Data from four participants were removed for failing to attain accurate performance on at least 50% of RR trials and/or failing to have at least four successful reaches in each of the RR conditions. Accordingly, all analyses were conducted on the remaining 27 participants.

Apparatus and stimuli. Participants performed reach responses (recorded from the tip of the right index finger at 200 Hz with OPTOTRAK) from a start button, placed 10 cm from the front edge of a table at which they sat, to a 40-in. touch screen (with 60 Hz refresh rate) placed 40 cm away from the start button. On each trial, shape pairs consisting of two outline-shapes (4 pixels wide, black outline on a white background) were shown on the screen, centered 9 cm to the left and right of a central fixation cross. There were four possible shapes: a circle, a square, an “X,” or a six-point star (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 (this resulted in a 1.77-cm-wide square, a 1.89-cm-wide “X,” and a 2.51-cm-wide star; all shapes had height equal to their widths). 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 aimed to accumulate points across 10 experimental blocks in Experiment 1. At the end of the experiment, points could be traded in for treats (chips, soda pop, or candy). Following piloting, the “cost” of items was set such that 100 points could be traded for a can of soda pop or a small bag of chips, and 200 points could be traded for a candy bar. These values were set such that the average participant could trade in for about three treats. Blocks were divided into VL blocks (60 trials per block) and RR blocks (64 trials per block). The following block order was used for each participant: subjects first completed a practice RR block to familiarize themselves with the timing of the RR task (points accumulated in the practice block did not count toward the final point total, and data from this block were not analyzed); then three VL blocks; then three RR blocks; then one top-up VL block (in which the previously learned associations were refreshed); and, finally, three more RR blocks. Points accumulated across these final 10 blocks counted toward the final score.

VL. Participants were instructed (a) to take their time initiating their movement and making shape selections in VL blocks, (b) that shapes would always have the same probability and payoff throughout the VL task, and (c) to focus on maximizing their points 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 cross at the center of the touch screen. After a variable delay (1 to 2 s) the shape pair appeared on the screen. At the same

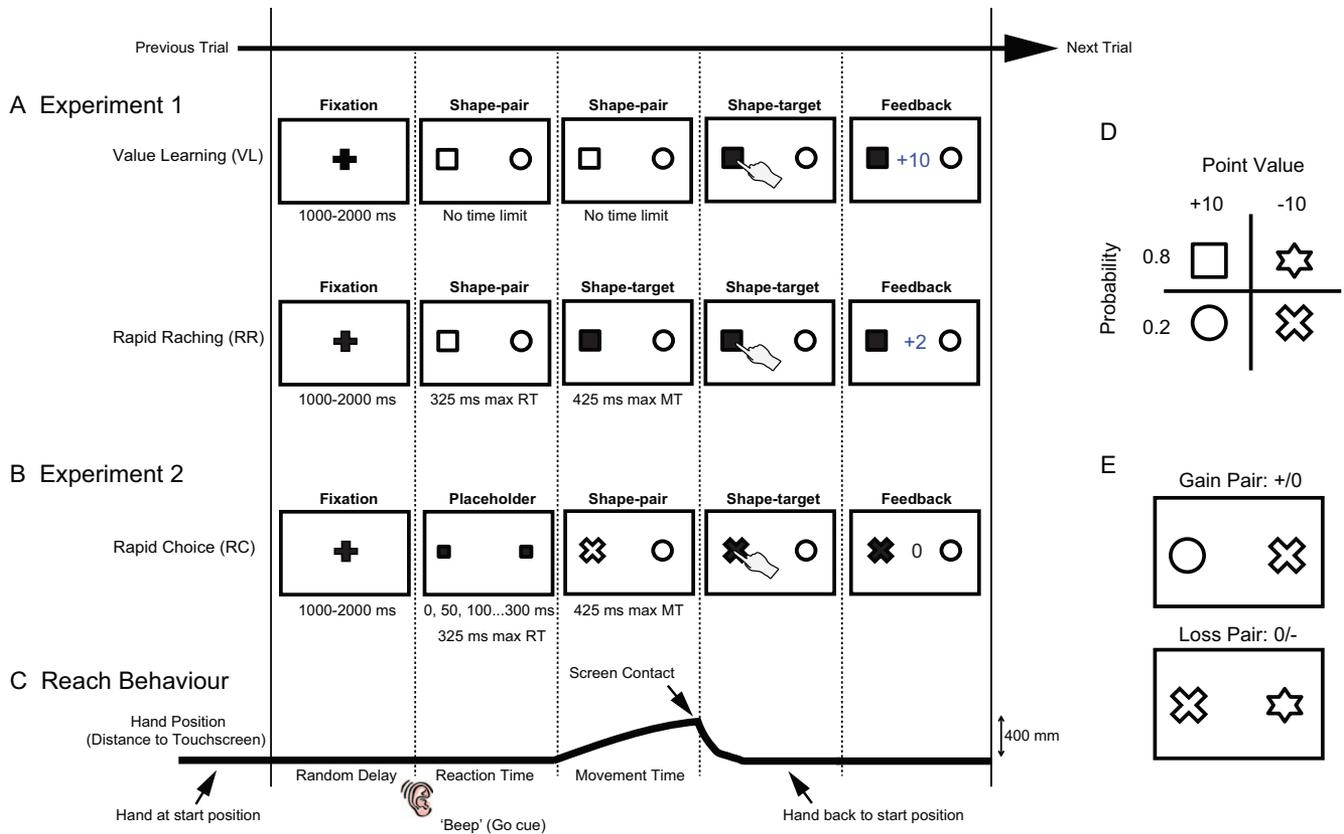


Figure 1. Schematic of (a) the value learning (VL) and rapid reach (RR) tasks from Experiment 1, and (b) the rapid choice (RC) task from Experiment 2, with (c) the linkage between the subject's movement and sequence of events in the trial. All tasks began with the presentation of a fixation cross for a random interval (1,000 ms to 2,000 ms) followed by an auditory "beep" cue that served as the go signal to begin movement. In the VL and RR tasks, fixation was replaced by a shape-pair. In the RC task, fixation was replaced by a placeholder that remained on the screen for a variable amount of time (0 ms to 300 ms, varying in 50-ms increments) before being replaced by the shape-pair. In the RR and RC tasks, participants had to initiate their reach within 325 ms upon hearing the go signal (see Panel c). Upon movement onset in the RR task, the cued target shape filled-in black. In the VL and RC tasks, participants could choose which shape they wanted to touch (and once touched, the selected shape filled-in black). In the RR and RC tasks, participants had 425 ms from movement onset to touch the shape-target (see Panel c). (d) Possible shapes (VL and RR) and values (VL) used in Experiment 1. (e) The two pair-types used in Experiment 2 (RC). For all tasks, values were randomly assigned to shapes for each participant. See the online article for the color version of this figure.

time, the fixation cross was removed and an auditory "beep" cue was played, signaling participants to initiate their response. Trials with an anticipatory response (<100 ms from beep) or a significantly delayed response (no reach initiated within 5 s of the beep) were terminated, an error message was displayed on the screen, and the resulting data from the trial were not analyzed. Trials ended when participants selected one of the two shapes (attempting to maximize the points accrued) by touching within 3 cm of it. The selected shape was then filled-in black to indicate selection to the participant (see Figure 1a, VL). At the end of every trial, participants saw their current accumulated score as well as the points accrued as a result of that single trial. For instance, if the selected shape resulted in a gain, participants saw the value "+10" (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 four shapes was randomly assigned one of four values along two orthogonal dimensions: probability and points (see Figure 1d). 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) and 0.8- (80% chance of receiving -10 points), matching the values used in previous work that examined the effects of learned value (e.g., Raymond & O'Brien, 2009). For trials in which a shape did not pay out (e.g., 80% of the time that a 0.2+ target did not provide points), participants received zero points. All possible shape-pairs were presented in the learning task (for a total of six pairs) and were counterbalanced for side of presentation (e.g., both 0.8+ vs. 0.2- and 0.2- vs. 0.8+ displays appeared an equal number of times). Each shape-pair was presented 10 times (five

presentations each for a pair and its mirror image, in a random order) in each of the four VL blocks.

RR. Participants were instructed to make fast reaches in RR blocks, as the only way to earn points was to complete a RR toward the cued-target within the allotted time (each successful trial resulted in +2 points). The structure of RR trials was identical to VL trials up to and including the beep. After the beep, however, in an RR trial, participants were instructed to initiate a reach movement toward the screen as quickly as possible. Reflecting this demand, trials with anticipatory responses (<100 ms) or responses slower than 325 ms were terminated (error feedback was given and trials were not analyzed). Upon movement onset (determined by velocities that exceeded 20 mm/s with an acceleration of 20 mm/s² over four 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 1a, RR). Our previous work has shown that by delaying target cuing until the reach is initiated, movement planning biases can be directly observed in the early stages of the subsequent reach trajectory, providing a behavioral “read-out” of evolving cognitive processes (Chapman, Gallivan, & Enns, 2015; Chapman et al., 2010a, 2010b; Chapman et al., 2014; Gallivan & Chapman, 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 or an error was displayed. Reaches that contacted the display more than 3 cm offset from the cued target resulted in an error and were not analyzed.

Critically, unlike the VL task, shapes in the RR task did not indicate value. In fact, participants were explicitly told to ignore shape and instead focus on speed, as 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. If they executed the reach too slowly (the most common type of error), they were awarded no points. As in the VL blocks, 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 six shape-pairs shown during VL blocks, RR blocks also contained displays with the four shape-pairs of identical shapes (i.e., trials in which a square would appear on both the left and right sides of the display; data from these pairs were not analyzed). Combined with the mirror reversals of the six nonidentical pairs (i.e., 12 display pairs), this led to 16 total display pairs. Because a reach toward each of these display pairs could have the final target cued on the left or right, this resulted in 32 conditions, each of which was repeated twice (in a random order) in each of the six experimental RR blocks.

Data analysis.

Reaction times (RTs). For a complete analysis of RTs for all experiments, please refer to the online supplemental materials and Supplemental Figure S1.

VL. We analyzed the choice behavior of participants over the final two VL blocks of the experiment. These were the VL blocks that appeared immediately prior to and between the nonpractice RR blocks. We calculated the probability of an optimal choice—P(optimal choice)—as the number of times participants chose the shape in a pair with the higher expected value (Probability ×

Points) relative to the number of times a given pair was presented. We conducted a Greenhouse-Geisser (GG) corrected repeated measures analysis of variance (RM-ANOVA) on P(optimal choice), with Shape-pair as the factor (six shape-pairs). We used Bonferroni corrected post hoc pairwise comparisons to assess differences between individual shape-pair means.

RR. In addition to the removed trials described earlier, we also removed trials for which the reach movement duration was excessively slow (>850 ms), and then removed trials with movement times that were more than two standard deviations above each individual participant’s mean. This meant that we included in our analysis some trials in which the participant saw a “Too Slow” error.

Reach trajectories were normalized (across reach distance; see Chapman et al., 2010a, and Gallivan & Chapman, 2014, for details) and averaged for each nonidentical shape-pair and shape-target-location (i.e., cued-left or cued-right). To provide a sensitive measure of any reach trajectory biases induced by a shape-pair, we calculated the area between trajectories toward a given pair and its mirror image for a common shape-target cued location (see Figure 2b, c). If there is a bias toward one of the two shapes (as was seen with eye movements; Hickey & van Zoest, 2012), then these two trajectories will diverge and the area between them will be larger than zero. However, if there is no bias, then the trajectories will not diverge and the area between them will be close to zero (i.e., the trajectories will be largely overlapping). Average area values across shape-target-location (left and right) for the six shape pairs (identical shape pairs were not analyzed) were then tested using a GG corrected RM-ANOVA.

Results

VL. Figure 2a shows the probability that participants made the optimal choice (chose the higher value stimulus) in the final two blocks of the VL phase of the experiment. For all shape-pairs, optimal choice behavior was quite good (>80%) and significantly above chance (see Figure 2a; note that 95% confidence intervals [CIs] are well above 50%). However, we found significant differences in the P(optimal choice) behavior across shape-pairs, $F(12.94, 65.93) = 19.08, p < .01$. Further inspection revealed that this result was largely driven by lower P(optimal choice) for shape-pairs in which both shapes had low probability (0.2+/0.2−) or both shapes had negative value (0.2−/0.8−) compared with pairs that had the high-probability, high-value shape (0.8+) or pairs that had mixed valence (i.e., 0.2+/0.8−), in which the resulting choice behavior was close to, or at, ceiling. Specific results of the pairwise comparisons are also shown in Figure 2a.

RR. Figure 2b and 2c show the average reach trajectories toward targets on the left for shape pairs with the same expected value difference (6 points) when the optimal shape was on the left (solid lines) or right (dashed lines). Figure 2b is for the pair with two negative shapes, and Figure 2c is for the pair with two positive shapes. We index the degree of bias as the area between the solid and dashed curves (gray area in Figure 2b and 2c) and present these areas in Figure 2d. These results show that only one shape had a significant bias on reach trajectories. The data showed significant area differences across Shape-pair, $F(2.36, 61.46) = 3.65, p < .05$, driven by trajectory biases toward the shape that had previously been associated with high probability and high points

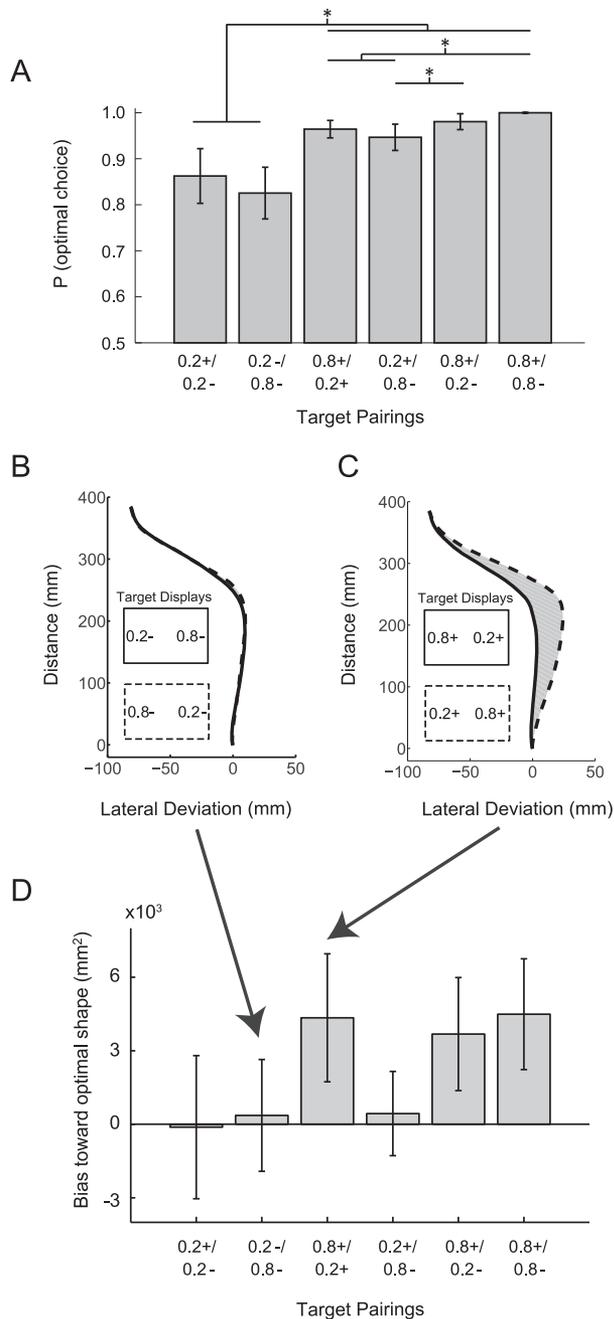


Figure 2. Results from the value learning (VL; a) and rapid reaching (RR; b through d) tasks in Experiment 1. (a) Percent optimal choice by shape-pair for the final two VL blocks. Error bars represent 95% confidence intervals (CIs) for that shape-pair. Lines and asterisks above bar plot show where pairwise comparisons were significant (Bonferroni corrected $p < .05$). (b, c) Average RR trajectories toward the left shape for the 0.2-/0.8- shape-pair (b) and the 0.8+/0.2+ shape-pair (c) when the learned-optimal shape was on the left (solid line) or right (dashed line). Note that even though the expected value difference for the pairs in (b) and (c) is identical (i.e., 6 points), the trajectories clearly are not. Dark gray shading depicts the area between the two trajectory traces. (d) Total area between the curves for each shape-pair, as depicted in (b) and (c). Error bars represent 95% CIs for that shape-pair.

(0.8+), regardless of what pair it was in.¹ Recall that in the RR task, these associations are no longer relevant, and points were only earned for successful completion of the movement, independent of shape. The bias toward the previously associated high-reward, high-points shape is evident in that the three pairs that include the 0.8+ shape all have area-indexed biases that are significantly greater than zero (see 95% CIs in Figure 2d), whereas pairs that do not include the 0.8+ shape are clearly not different than zero. We confirmed this difference statistically using a post hoc comparison showing that the mean area of the three shape-pairs (4172.30 mm²) containing the 0.8+ shape was significantly larger than the mean area of the three shape-pairs (226.66 mm²) that did not contain the 0.8+ shape, $t(26) = 3.20$, $p < .01$.

Discussion

Experiment 1 followed the research protocol of Raymond and colleagues (O'Brien & Raymond, 2012; Raymond & O'Brien, 2009; Rutherford et al., 2010), who studied the role of reward in perception (and recently, grasping actions; see Painter et al., 2014). Participants first completed a VL task in order to establish an association between shapes and values. In the subsequent rapid-reaching task, in which these associations were now task irrelevant, only the high-gain shapes (0.8+) resulted in initial reach biases. This finding held even though the value of all shapes was learned successfully (80% or better choice performance for all shapes). This finding extends the work of Raymond and colleagues in showing that learned value not only has a profound effect on future perceptual processing, but that these associations influence immediate pointing actions based on those perceptions. This result also clarifies that under significantly limited processing times, it is only the confluence of high gain and high probability (i.e., the best option) that results in preferential processing.

One interpretation of this result is that under the time constraints of Experiment 1, only the high-gain reward association had enough time (and generated a strong enough signal) to modulate the stimulus response toward that shape. As a preliminary test of this delayed loss-processing hypothesis, we examined the reach trajectory behavior from the Experiment 1 VL phase on gain-only (0.8+/0.2+) and loss-only (0.2-/0.8-) trials. Note that although no time constraint was provided in VL trials, we can look at the effects of time by using the natural variability in participant RTs. Our prediction is that when RTs are short, decision processing on gain, but not loss, trials will have been completed, resulting in larger differences in behavior when comparing these trial types. Conversely, on trials with slower RTs, we predict that the difference between gain and loss trials will be reduced, because the additional time afforded by a longer RT will allow for more decision processing to occur on the loss trials, allowing it to catch up to, and thereby decrease, the gain processing advantage. To test this, for each participant, and separately for the gain and loss pairs, we sorted the VL trials into those with the fastest one third of RTs

¹ Note that the area between curves measure does not allow for direction to be inferred, only to index that a bias does exist. We ask the reader to refer to the online supplemental materials (and Supplemental Figure S2) for a complementary trajectory analysis that shows that, indeed, the hand moves toward the 0.8+ shape.

and those with the slowest one third of RTs (note that three participants were removed from this analysis due to a lack of optimal reaches). We then calculated the reach area on choices toward the optimal choice for both the gain-pair (optimal choice is 0.8+) and loss-pair (optimal choice is 0.2-). This reach area calculation differs from that in the Experiment 1 RR and presented in Figure 2, and instead looks at reaches toward both the left and right endpoints (as in Experiment 2). The resulting reach area measure is larger when a reach path is straighter, and smaller when the reach path is less straight (because ambivalent reaches, made with less decision information, are initially directed more toward the center of the screen). Here, we use area magnitude as an index of decision processing, with larger areas reflecting decisions based on processing that is more complete than smaller areas. As depicted in Figure 3, when we took the difference between the gain and loss pair areas, we found our prediction to be borne out by the results: The difference in reach area toward the optimal choice in gain and loss pairs is significantly larger with faster RTs than it is with slower RTs, $t(23) = 2.22, p < .05$.

To more directly test the loss delay hypothesis, we conducted Experiment 2, which employed a rapid choice (RC) task that was a hybrid of the two phases studied in Experiment 1: Now participants could freely choose between shapes, but they did so under a range of different processing times. By parametrically varying how long participants could see the two shapes before they had to choose one, we were able to directly estimate the time course of gain and loss processing. The loss delay hypothesis would predict that with sufficient time, the loss-related shapes, which did not influence RR in Experiment 1, will bias the reach trajectories in Experiment 2.

Value Learning

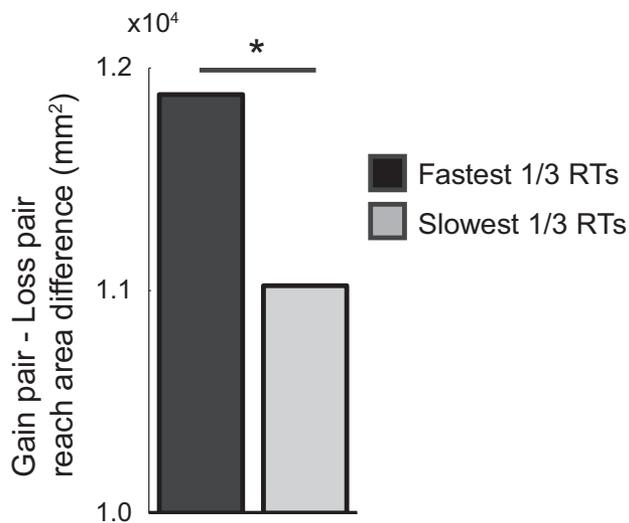


Figure 3. Reach-area differences between gain and loss pairs on Experiment 1 value learning trials for trials with fast reaction times (RTs; black) and slow RTs (gray). Asterisk indicates a significant ($p < .05$) difference, suggesting that with less processing time (faster RTs), the asymmetry between gain and loss processing is magnified relative to trials with more processing time (slower RTs).

Experiment 2

Method

Participants. A new group of 38 right-handed (mean age = 19.9 years; 27 female) healthy young adults participated in exchange for course credit or money. As in the first experiment, data from four participants were removed for failing to attain accurate performance on at least 50% of RR trials and/or failing to have at least four successful reaches in each of the RR conditions, leaving 34 participants whose data were initially analyzed.

Apparatus and stimuli. The experimental setup was identical to Experiment 1, except that there were only three possible shapes: a circle, a square, or a six-point star (see Figure 1e).

Procedure. Participants in Experiment 2 aimed to accumulate points across 10 experimental blocks (56 trials per block) following one practice block (points were not accumulated and data not analyzed from this block). All points accumulated could be traded in for treats, as in Experiment 1.

The timing of trials in Experiment 2 was identical to the temporally demanding RR trials in Experiment 1. That is, participants had to initiate a reach within 325 ms and complete a reach (once initiated) and touch the selected target within 425 ms. The goals of the task for participants in Experiment 2, however, were much closer to the task on the VL trials in Experiment 1; that is, instead of a target being cued at movement onset (as in RR trials in Experiment 1), participants were instructed to choose which of the two shapes they wanted to reach toward (see Figure 1b). Upon touching within 3 cm offset of either shape, it filled-in black to indicate selection. Shapes in Experiment 2 were tied to value, as in the Experiment 1 VL trials. Specifically, for each participant in Experiment 2, each of the three shapes was randomly assigned one of three values: +5 points, 0 points, or -5 points. We did not manipulate probability in Experiment 2, as all shapes gave their assigned payoff 100% of the time. As in Experiment 1, at the end of each trial in Experiment 2 participants saw their total accumulated score: the points they earned on that trial and then the updated accumulated score.

Unique to Experiment 2, we parametrically manipulated the amount of time the shapes were presented on the screen following the reach onset cue (thus manipulating the amount of time that subjects could use the shape information in order to plan their initial arm movements). Specifically, the presentation of shape-pairs was randomly delayed between 0 ms and 300 ms following the auditory beep “go” signal in 50-ms increments (so a total of seven different delays were used²). At 0 ms, the timing was identical to Experiment 1 because the shapes appeared coincidentally with the beep. For all other delay intervals (e.g., 50 ms to 300 ms), a small black placeholder (20 × 20 pixel square) appeared at both of the upcoming shape locations for the specified delay, after which time it was replaced by a shape. As such, participants were given knowledge where the shapes would eventually appear without knowing exactly what they would be. Because RTs across participants were, on the average, close to or less than 250 ms, this

² As we were using a 60-Hz refresh rate, our 50-ms increments are not exact. Instead, we presented the shape-pairs at the first refresh after the desired delay (e.g., 50 ms), so the actual delays could be as much as a full screen refresh later (i.e., additional 16.7 ms).

meant that on the majority of trials with 250-ms and 300-ms delay intervals participants were initiating reaches prior to the shape information being available. This allowed us to test the extreme case in which no sensory evidence for a choice was available during reach planning.

In Experiment 2, we presented only two shape-pairs (see Figure 1e)—a gain pair (+/0; consisting of the +5 and 0 shapes) and a loss pair (0/-; consisting of the 0 and -5 shapes). With the shape-pairs and their mirror images, there were four conditions, each presented 20 times across each of the seven different placeholder durations. The 560 resulting trials were presented in a random order across the 10 experimental blocks.

Data analysis. For this experiment, we performed two separate analyses. In the first analysis, we characterized the choice behavior (similar to the VL blocks in Experiment 1), and in the second analysis, we characterized the reach trajectories (similar to the RR blocks in Experiment 1). For the first analysis, all 34 subjects' data were analyzed. However, for the second analysis, because we wanted a reliable estimate of reach behavior, we eliminated participants who did not perform at least four reaches toward the optimal shape in each time bin and for each shape pair. These selection criteria removed seven participants from the trajectory analysis, leaving 27 participants whose reach trajectories were analyzed.

Choice behavior. As in the VL blocks from Experiment 1, we calculated the P(optimal choice) across each of the four shape-pairs (gain, loss, and mirror images of both) across all seven placeholder delays. These data were then entered into a 2 (Valence: gain or loss) \times 2 (Side of optimal target: left or right) \times 7 (Delay: 50-ms time bins) GG corrected RM-ANOVA. Additionally, to assess the learning of the reward association, we examined how P(optimal choice) changed across the experiment. Here, we collapsed across Target Side and Delay and entered the data into a 2 (gain or loss) \times 11 (Experimental Block) GG-corrected RM-ANOVA. For this analysis, we included the practice block (and, as a result, one additional participant who now met our inclusion criteria), as a large portion of learning happens very early in the experiment and we wanted to capture this initial learning in our analysis.

Reach trajectories. We analyzed only reaches toward the optimal shape of any given pair (positive for +/0 and neutral for 0/-). To index the degree of decision difficulty, we measured the area between trajectories produced for a given pair and its mirror image (see Figure 4c, d). Decisions that were "easier" would result in trajectories straight toward the optimal shape (as in Painter et al., 2014; Resulaj et al., 2009). Thus, for an easy decision, when the optimal shape is on the left, trajectories will go straight to the left, whereas when the optimal shape is on the right (mirror image), trajectories will go straight to the right, resulting in a large area between these trajectories. However, if the decision is more difficult, trajectories will not be as straight, indicative of indecision between the two targets resulting in a smaller area between the two trajectories. The resulting areas were calculated for each pair-valence (gain: +/0 and loss: 0/-) and for each of the seven placeholder delays, and then analyzed using a 2 \times 7 GG corrected RM-ANOVA.

Results

Choice behavior. Figure 4a shows P(optimal choice) as a function of the duration of the placeholders and whether the choice involved a gain pair (red) or a loss pair (blue). Not surprisingly, we

found that with increased delay (and therefore less time to process the shapes), participants made fewer optimal choices, $F(2.28, 75.09) = 59.38, p < .01$. A main effect of Side, $F(1, 33) = 18.45, p < .01$, revealed that participants made more optimal choices on the right than on the left. We also found a Delay \times Side interaction, $F(2.77, 91.55) = 3.14, p < .05$, in which, at the long delays (200+ ms), the Side effect became magnified; that is, participants made more optimal choices when the best target was on the right compared with the left. This, combined with the overall Side effect, replicates our previous work demonstrating that participants favor targets on the right, and under extreme time pressure, will begin to guess in that direction (Chapman et al., 2010a). Most interestingly, we found an effect of Valence, $F(1, 33) = 4.55, p < .05$. Consistently, participants performed more poorly (made fewer optimal choices) on loss relative to gain pair trials.

Figure 4b shows P(optimal choice) as a function of the experimental block and whether the choice involved a gain pair (red) or a loss pair (blue). Here, we found a main effect of Block, $F(5.23, 179.68) = 26.39, p > .001$ and a main effect of Valence, $F(1, 34) = 6.53, p < .05$. The Block effect is driven by performance that gets better across time, characterized by relatively poor performance during the practice block (Block 1) compared with all other blocks (all Bonferroni corrected p 's < 0.05). Similarly, P(optimal choice) on Block 2 is also significantly lower than all blocks that follow it, except for Blocks 3 and 11 (Bonferroni corrected p 's < 0.05). Performance across the last 9 blocks remains statistically stable. As in the previous analysis, the Valence effect is driven by higher P(optimal choice) on gain trials than on loss trials.

Reach trajectories. Figure 4c and 4d show the average reach trajectories toward the optimal shape for gain pair (4c, red) and loss pair (4d, blue) trials when the optimal shape was on the left (solid lines) or right (dashed lines). We calculated the area between these trajectories as an index of the decision difficulty (easier decisions lead to straighter hand paths and larger areas). These area measures are presented in Figure 4e for each placeholder duration and are shown for both the gain pair (red bars) and loss pair (blue bars) trials. As expected, we found a main effect of Delay, $F(1.94, 50.44) = 26.18, p < .01$, in which, as the time for processing shapes diminished, so, too did the area between the trajectories, indicating the decisions were more difficult. Of most interest to this study, we also found a main effect of Valence, $F(1, 26) = 26.39, p < .01$, in which gain pair trials produced larger areas (i.e., easier decisions) than loss pair trials. Finally, there was an interaction between Valence and Delay, $F(2.92, 75.91) = 3.31, p < .05$, again demonstrating the diminishing effect of Valence with longer delays. This was confirmed by conducting paired t tests of gain- versus loss-pairs for each of the seven placeholder durations (corrected for multiple comparisons). For the area measure, only the first four delays (0 ms, $t[26] = 4.87, p < .007$; 50 ms, $t[26] = 4.72, p < .007$; 100 ms, $t[26] = 3.06, p < .007$; 150 ms, $t[26] = 3.09, p < .007$) showed a significant difference between valences.

Discussion

We investigated the time course of gain- and loss-based decision making in Experiment 2. In this experiment, we combined the learning and reaching phases of Experiment 1 into a single hybrid task in which participants were forced to make a rapid choice

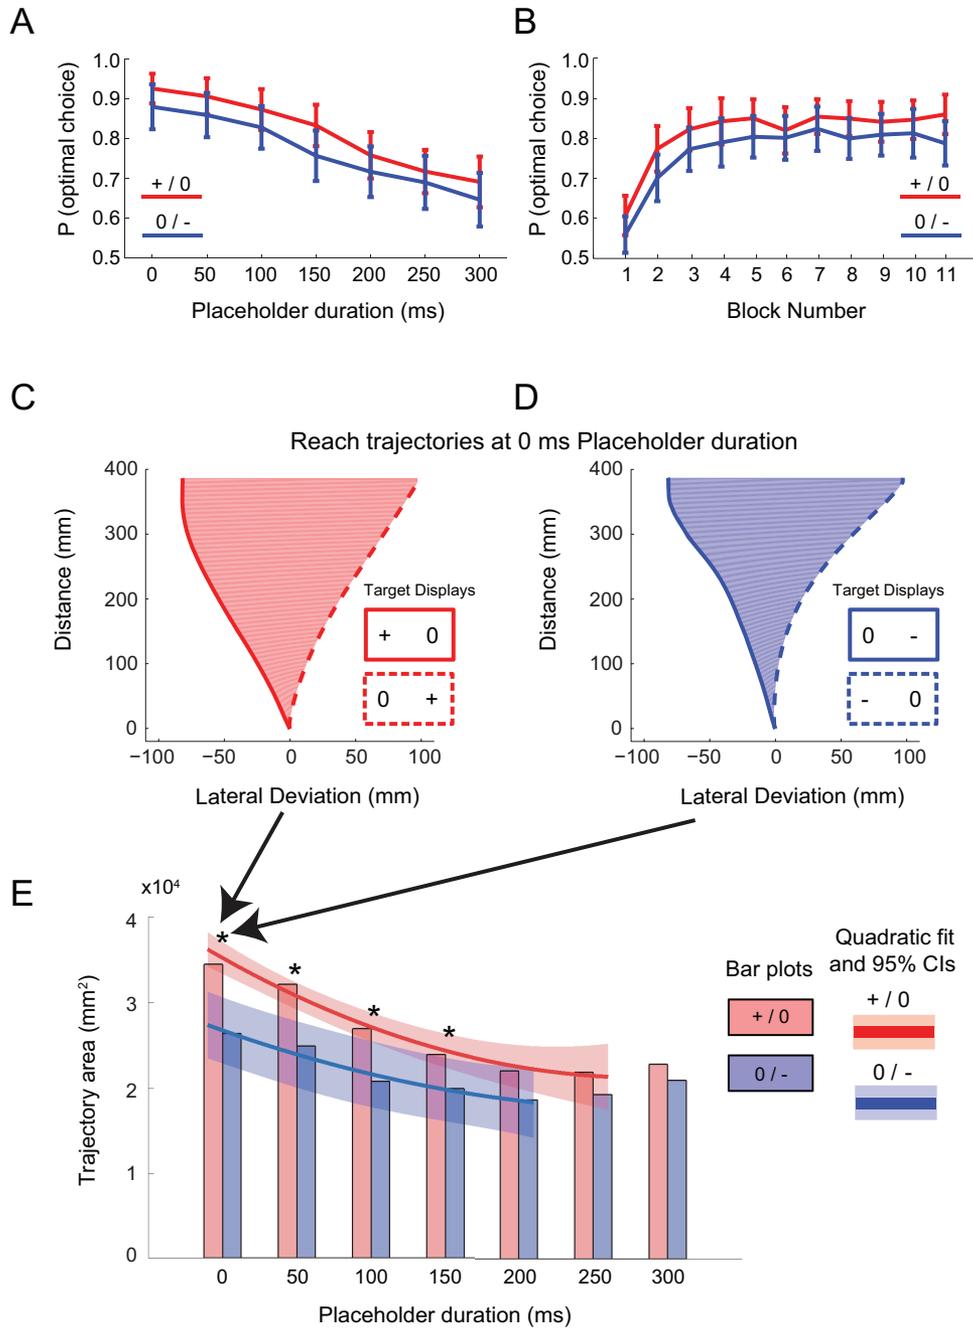


Figure 4. Results from the rapid choice task in Experiment 2. (a, b) Results from the choice behavior showing the P(optimal choice) across all seven placeholder durations (a) and block number (b) for the gain pair (red) and loss pair (blue). Error bars in (a) and (b) represent 95% confidence intervals (CIs) and show that for all time points, blocks and pairs choice behavior was significantly above chance. (c, d) Average trajectories toward the optimal shape at the 0-ms placeholder duration for the gain pair (c) and the loss pair (d) when the optimal shape was on the left (solid line) or right (dashed line). Note that even though the expected value difference for the pairs in (c) and (d) is identical, the trajectories clearly are not. (e) The total area between the curves as depicted in (c) and (d) across all placeholder durations for gain (red) and loss (blue) pairs. Asterisks in (e) show at which time points the gain and loss pairs significantly differed ($p < .007$). Quadratic fits and 95% CIs of the fit are shown for the first five gain-pair data points (solid red line and shading) and the first four loss-pair data points (solid blue line and shading). Fits were restricted to these time points due to evidence for guessing for delay intervals beyond these time points. See the online article for the color version of this figure.

between two shapes that were each associated with a gain (+5), a neutral (0), or a loss (−5) outcome. We investigated the evolution of the gain and loss response by varying the time given to view the shapes prior to response completion. Importantly, unlike the RR task in Experiment 1, participants were free to choose which shape to touch and were motivated through reward to accumulate points. Accordingly, the gain shape was optimal when paired with the neutral shape, whereas the neutral shape was optimal when paired with the loss shape.

The results of Experiment 2 showed that participants were now clearly able to use both gain and loss information. However, they showed that gain information had an earlier effect on movement planning than loss information, as indicated by the more direct hand paths for gain over neutral shapes than for neutral over loss shapes. This time course information in Experiment 2 thus offers an explanation for why only gain information appeared to transfer across tasks in Experiment 1: The slower signals associated with loss shapes did not have the same opportunity to become active in that first experiment.

We used the data in Experiment 2 to derive an estimate of the delay of loss relative to gain information encoding. We first fit quadratics to the gain and loss data from Experiment 2 over time points in which guessing was minimal (guessing behavior was conveyed through diminishing error counts combined with diminishing percent optimal choice; see the online supplemental materials and Supplemental Figure S3 for details). The curves are overlaid on the bar plots in Figure 4e. From these fits, we estimated the temporal shift required to align the loss and gain curves. That is, we calculated what value of the parameter responsible for shifting the quadratic curve along the x -axis (time) would be required to align the specific points from the loss data to the gain curve (and then, from the gain data to the loss curve). The consistency of these values is notable; for the loss data, the first three time points (0 ms, 50 ms, and 100 ms, after which a simple translation is not sufficient) need to be shifted forward 112 ms, 89 ms, and 184 ms, respectively, whereas for the gain data, the first five time points (0 ms, 50 ms, 100 ms, 150 ms, and 200 ms, after which the quadratic is reaching an asymptote) need to be shifted backward 107 ms, 128 ms, 103 ms, 100 ms, and 109 ms, respectively. The consistency of these observations suggests that the encoding of loss information is delayed by approximately 100 ms relative to gain information. We use this 100-ms value as the estimated relative delay of loss relative to gain in the modeling exercise described later.

A plausible alternative explanation for our results is that the effect we are attributing to a loss delay (i.e., the time required to inhibit the initial positive stimulus response) is instead a consequence of the poorer learning of loss reward associations. Recall that in both Experiment 1 (see Figure 2a) and Experiment 2 (see Figure 4a, b), the pairs with loss shapes showed reduced learning relative to the pairs with gain shapes. If the difference in reach area between gain and loss pairs exists because the gain pair was better learned, then reach area should not differ between gain and loss pairs when the amount of learning is equated for both pairs. To test this hypothesis, we compared gain-pair data from the three blocks with the lowest P(optimal choice) to loss-pair data from the three blocks with the highest P(optimal choice). We therefore use the probability that subjects chose the correct stimulus in a given pair as a measure of learning. For this analysis, we also included the

practice block (and, in doing so, rendered one additional participant eligible for inclusion), and as such, the worst-gain blocks were the practice and first two experimental blocks and the best-loss blocks were Blocks 7 through 9. For both P(optimal choice) and reach area, we then compared gain and loss performance using a t test. The results of this analysis are depicted in Figure 5. As intended, choice performance was better on the best-loss trials than it was on the worst-gain trials, $t(34) = 3.12, p < .005$. Remarkably, however, despite a learning advantage for the loss-pair trials in this subset of the Experiment 2 data, the reach area effects still clearly shows that reaches are straighter for the gain-pairs relative to the loss pairs, $t(34) = 6.50, p < .001$. This result shows that differences in reach area between gain and loss pairs are not related to inferior learning of loss-pairs.

Another potential limitation with our Experiment 2 results may be that what we interpret as a loss-delay could actually result from having presented participants with only two pairs. That is, with this reduced decision set, decision making on the gain pair may, in principle, be reduced simply to a go response toward the gain shape, whereas decision making on the loss pair may include the failed search for the go-related gain shape (thereby accounting for the delay) followed by a go response to the neutral shape. To address this potential shortcoming, and to provide additional evidence against the poor loss learning account noted earlier, we conducted a third experiment. In this third experiment, to address any learning-related differences, we included an extended experimental session. This allowed us to have sufficient statistical power across different stages of learning to compare trials in which the choice performance on loss pairs was equal to, or better than, choice performance on gain pairs. This is similar in spirit to the preliminary analysis conducted here (see Figure 5) but allowed for vastly more statistical power. If in this new experiment we still observe persistent effects consistent with loss delay, then we can definitively rule out learning effects as an alternative explanation of our results. To address the problem of having only two pairs, Experiment 3 included five shapes, in all possible pairings, for a total of 10 pairs. If we further show that a loss delay persists for pairs that require more computation than a simply go-for-gain strategy, then this will strengthen our interpretation of delayed loss processing.

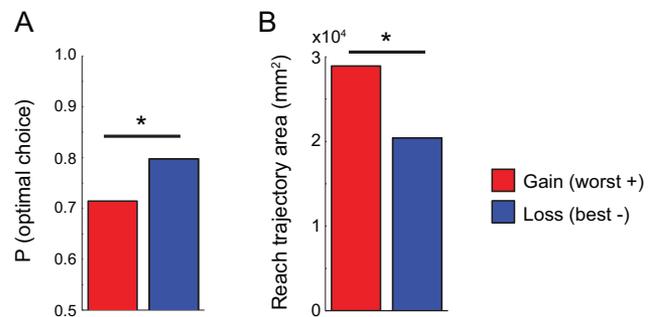


Figure 5. (a) P(optimal choice) and (b) reach area for gain and loss trials selected from the three worst-gain blocks (red) and the three best-loss blocks (blue). Asterisks indicate a significant ($p < .05$) learning advantage for the best-loss trials, but despite this, a significantly straighter hand path for the worst-gain trials. See the online article for the color version of this figure.

Importantly, our third experiment also allows us to test a critical prediction of our delayed loss processing hypothesis: If the processing of loss information is simply delayed by approximately 100 ms, then giving participants extra time to preview the shape stimuli before acting on them should erase any gain–loss asymmetry evident in their reach behavior. Therefore, in a critical new manipulation in Experiment 3, we include both an immediate condition (replicating the timing demands of Experiment 1 RR and Experiment 2, 0-ms placeholder delay) and a delay condition, which gave participants a 500-ms preview of the shape targets before they were cued to by the auditory beep to launch their reach movement. The hypothesis of delayed loss processing predicts that, in the immediate condition, we will replicate the gain loss asymmetry observed in Experiment 2, but that in the delay condition (with the additional 500 ms), the loss processing will “catch up” to the gain processing, resulting in equivalent reach behavior.

Experiment 3

Method

Participants. A new group of 29 right-handed (mean age = 24.3 years; 16 female) healthy young adults participated in exchange for course credit or money. Data from four participants were removed for failing to have least two successful reaches in each of the RC conditions (see RC Procedure); accordingly, all subsequent analyses were conducted on the remaining 25 participants.

Apparatus and stimuli. The experimental setup was identical to Experiment 1 and Experiment 2, except that there were five possible shapes: a diamond (2.51 cm wide), a circle, a square, a six-point star, or an “X” (see Figure 6d).

Procedure. Participants in Experiment 3 aimed to accumulate points across 20 experimental blocks (60 trials per block). The experiment was divided into two, approximately 1-hr, 10-block sessions. All participants completed the sessions back to back (e.g., as though it were a single, 2-hr experiment). Immediately preceding the first experimental block of each session, participants completed 10 practice trials (RC trials; points were not accumulated and data not analyzed from these trials). Unlike Experiment 1 and Experiment 2, points in Experiment 3 directly corresponded to additional cash incentive—every point was worth 1 cent, and we paid out this extra money at the end of the second session (rounded to the nearest \$0.25, with no penalty for negative points). Notably, having a monetary rather than food reward (as in Experiment 1 and Experiment 2) alleviates potential limitations to interpretation due to different incentives based on hunger. On average, participants earned an extra \$5.00.

Blocks were divided into VL blocks and RC blocks. Within each experimental session, participants completed two repetitions of the following block order: one VL block followed by four RC blocks—hereafter, we will refer to one grouping of these 5 total blocks as one “experiment repetition.” Therefore, across the entire experiment, participants completed four experiment repetitions, for a total of four VL blocks and 16 RC blocks (see Figure 6c).

Participants saw five different shapes in Experiment 3, which were each assigned a different combination of probability and points (see Figure 6d): 1.0+ (100% chance of receiving +2 points), 0.5+ (50% chance of receiving +2 points, 50% chance of

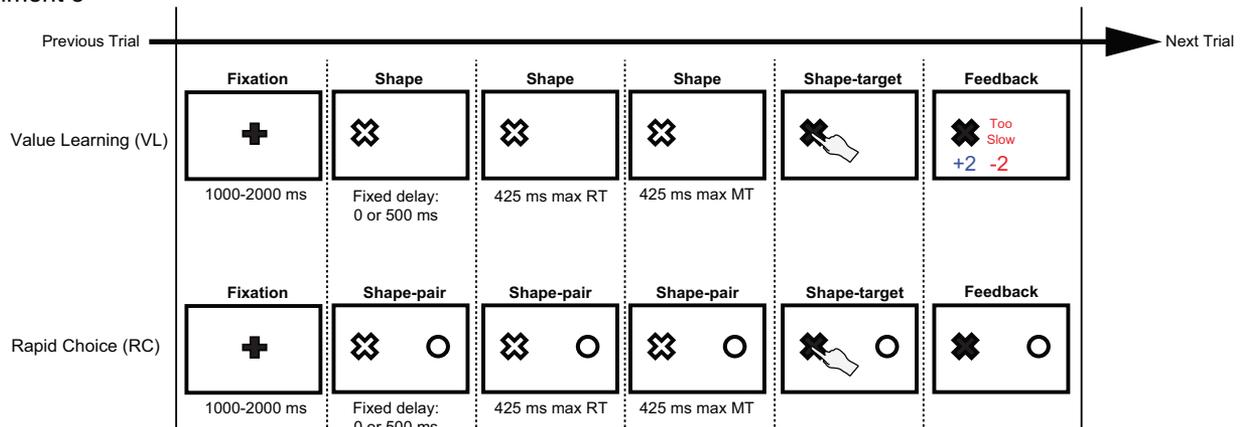
receiving 0 points), 0 (100% chance of receiving 0 points), 0.5– (50% chance of receiving –2 points, 50% chance of receiving 0 points) and 1.0– (100% chance of receiving –2 points).

The timing of VL and RC trials in Experiment 3 was identical, and for half the trials, this timing was the same as Experiment 2 trials that had a 0-ms placeholder delay. We refer to these trials as “immediate” trials because the shape(s) appeared coincident with the auditory cue to move. On the other half of the Experiment 3 trials, we introduced a 500-ms delay between the onset of the shape stimuli and the cue to move. We refer to these as “delay” trials, as participants had this extra 500 ms to process the shapes prior to initiating a movement (see Figure 6a). The specific value of 500 ms was selected for this experiment based on our previous work showing that it is a long enough duration to overcome inherent processing delays associated with salience biases (Wood et al., 2011) and symbolic processing (Chapman et al., 2014). As with all of our experimental conditions, immediate and delay trials were randomly interleaved. Compared with the timing of Experiment 2 trials, the only other change we made was to relax the RT constraint from 325 ms to 425 ms (applied on both immediate and delay trials). After extensive pilot testing, we found that the need to wait for the go cue on the delay trials made participants marginally slower to react on the immediate trials (which were intermixed with the delay trials), and thus they were often receiving “time out” errors with a 325-ms limit. We found that this extra 100 ms to react was sufficient to achieve RT error rates that were comparable with Experiment 2.

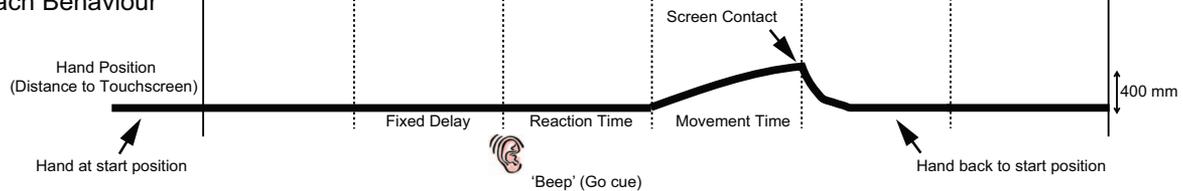
VL. VL and RC trials were identical except for two critical differences. First, on VL trials, only a single target appeared on the screen. Thus, there was no choice to make, simply a requirement to hit the presented target within the timing demands. Second, only on VL trials did participants receive trial-by-trial feedback about the outcome of their movement. This feedback was “full,” in that even if participants made an error, the outcome of hitting the shape presented was still shown. For example, if participants moved too slowly to select the 1.0+ shape, they would have seen the error message “Too Slow” (appearing in the center of the screen) and the penalty (–2) would have been shown under this error message, animated downward and subtracted from their total score, which appeared at the bottom of the screen (see Figure 6a, VL). Concurrent with this error message and score deduction being shown, participants would also have seen a “+2” appear under the shape, indicating the points they would have received if the reach would have been fast enough. On a successful trial, these reward points (e.g., “+2”) would have been animated downward and added to the total score, and no error message or penalty points would have appeared on the screen. We opted for full feedback during VL trials, and no feedback during RC trials, in order to maximize learning and largely restrict it to the VL blocks. Presumably, this created learning states that were relatively stable across a given RC block and that only were updated with each VL block.

During each VL block, each of the five shapes appeared 12 times: 6 times on the left and 6 times on the right, for a total of 60 trials. As mentioned earlier, half of the VL trials were immediate trials, and half were delay trials. In order to guarantee that participants experienced the 0.5+/- shapes as truly having a 50% chance of reward, the reward and feedback accompanying these shapes were drawn randomly, without replacement, from a distribution of six neutral and six positive (or negative) outcomes.

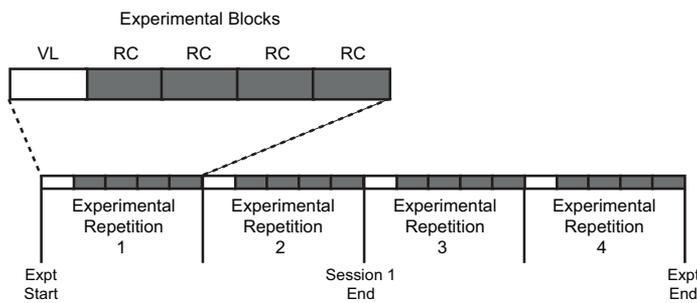
A Experiment 3



B Reach Behaviour



C



D

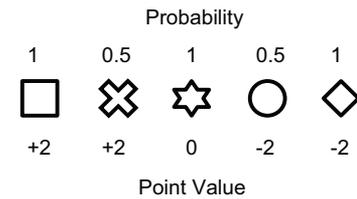


Figure 6. Schematic of (a) the value learning (VL) and rapid choice (RC) tasks from Experiment 3, with (b) the linkage between the subject's movement and sequence of events in the trial. Trial timing was very similar to the RR task depicted in Figure 1, with the addition that on half the trials, participants were given a 500-ms preview (delay) prior to hearing the go cue. VL and RC trials differed in that on VL trials, participants saw only a single shape and received full feedback (i.e., saw what the outcome of the touching the shape would have been even if they received an error), whereas on RC trials, participants made a choice between two shapes and received no feedback. (c) Diagram showing that within each of four experimental repetitions, participants completed one VL block followed by four RC blocks. (d) Possible shapes and values (combination of probability and points) used in Experiment 3. For all tasks, values were randomly assigned to shapes for each participant. See the online article for the color version of this figure.

RC. RC trials differed from VL trials in that there were always two targets on the screen (as in Experiment 2, thus requiring a rapid choice), and feedback was not given every trial, but instead was only given after each block of 60 trials (i.e., participants were shown their current total score). All possible pairs of the five shapes were presented, for a total of 10 pairs (same-shape pairs were not presented). Pairs were counterbalanced for side of presentation and were each presented 12 times (six presentations each for a pair and its mirror image, in a random order) in both the immediate and delay conditions across each of the four experimental repetitions. Together, the factors of Pair ($20 = 10 \text{ pairs} \times 2 \text{ for mirror image}$), Timing (2), and Experimental Repetition (4)

meant there were 160 unique conditions in this experiment. The participants that were removed from analysis failed to have at least two reaches in each of the 160 conditions. All other participants had at least two reaches in these conditions (less than 1% of all remaining conditions, across participants, had only two reaches in them). Thus, during each experimental repetition, there were 240 RC trials, presented randomly across four 60-trial blocks.

Although feedback regarding reward points was not shown during RC trials, participants still received error feedback regarding their reach speed and accuracy (i.e., saw "Too Early," "Time Out," "Too Slow," or "Miss" on the corresponding error trials). Even though participants did not see it on every trial, their total

score was being adjusted based on the shapes they were choosing—for example, they earned (or lost) 2 points every time they selected the 1.0+ (or 1.0−) shape. Unlike the VL trials which guaranteed that participants would experience the 0.5+/- as truly 50% (because the results were drawn without replacement from a pool of experimenter generated options), in the RC trials, the outcome of choosing a 0.5+/- shape was drawn randomly from a uniform distribution.

Data analysis.

VL. We analyzed the RT of participants in the VL blocks of Experiment 3 as an index of preferential processing. Because participants saw only a single target on these trials, an analysis of reach trajectory data would not be meaningful, as no second-choice option was available to bias trajectories. To analyze RT, we conducted two separate GG-corrected RM-ANOVAs. The first was a $5 \times 2 \times 2$ RM-ANOVA examining the effects of Shape (five possible shapes), Timing (immediate or delay trials), and Target Side (left or right). The second was a $5 \times 4 \times 2$ RM-ANOVA examining the effects of Shape (5), Timing (2), and Experimental Repetition (4; for this second ANOVA, one participant did not have data for all cell means and was not included in this analysis). Notably, we did not have enough trials to run this as one large $5 \times 2 \times 2 \times 4$ RM-ANOVA, and thus included the second RM-ANOVA as a test of the effect of learning, without the target side factor (which, as described later, was not significant or involved in any significant interactions). We used Bonferroni corrected post hoc pairwise comparisons to assess differences between main effect means.

RC. To streamline the analysis of Experiment 3 RC behavior and make it directly comparable with Experiment 2, we subdivided the shape-pairs into three groups: pairs that differed by an expected value of 1.0 (ΔEV 1.0: 1.0+/0, 0.5+/0.5−, 0/1.0−), pairs that differed by an expected value of 0.5 (ΔEV 0.5: 1.0+/0.5+, 0.5+/0, 0/0.5−, 0.5−/1.0−), and pairs that differed by an expected value of 1.5 or more (ΔEV 1.5+: 1.0+/0.5−, 0.5+/1.0−, 1.0+/1.0−). For brevity, in the main text, we restrict our analysis to the ΔEV 1.0, as these pairs include the same pairs as Experiment 2, with the important addition of the 0.5+/0.5− pair. However, for the complete analysis and discussion of the ΔEV 0.5 and ΔEV 1.5+ pairs, we ask the reader to see the online supplemental materials and Supplemental Figure S4. Of note, the results for these other EV groups are fully consistent with the results reported here.

For the ΔEV 1.0 pairs, we first analyzed the probability of an optimal choice using a $3 \times 2 \times 4$ GG-corrected RM-ANOVA to investigate the effects of Pair (1.0+/0, 0.5+/0.5−, 0/1.0−), Timing (immediate or delay), and Experimental Repetition (4). We used Bonferroni corrected post hoc pairwise comparisons to assess differences between main effect means. We then analyzed the reach trajectories using the identical reach area measure described in Experiment 2. Recall that for the RC trials in Experiment 3 we only enforced that participants have at least two successful reaches in every Pair \times Timing \times Experimental Repetition condition. Here, because we only analyzed reaches toward optimal choices, this meant that some participants did not have sufficient optimal choices to include Experimental Repetition as a factor. Therefore, we first tested for effects that were consistent across experimental repetitions by using a 3×2 GG-corrected RM-ANOVA to investigate the effects of Pair (3) and Timing (2). Then, to examine how reaching changed across the experiment, and to explicitly test the

alternative hypothesis that our Experiment 2 trajectory effects were driven by poor loss learning (rather than a loss delay), for each participant, we selected the two experimental repetitions with the lowest P(optimal choice) on gain-pair trials (worst-gain) and the highest P(optimal choice) on loss-pair trials (best-loss). We then extracted the P(optimal choice) and reach trajectory measures for the win-pair trials from the worst-gain repetitions and the loss-pair trials from the best-loss repetitions. Across participants, the worst-gain repetitions came predominantly from the first (24/25 participants) and second (19/25 participants) repetitions, with a small minority coming from the third (6/25 participants) and fourth (1/25 participants) repetitions. By contrast, the best-loss repetitions came predominantly from the fourth (20/25 participants) and third (18/25 participants) repetitions, with a minority coming from the second (8/25 participants) and first (4/25 participants) repetitions. This is consistent with the general trend of improved choice performance with task exposure. Importantly, this analysis allows us to take advantage of the extended experimental session we employed by directly comparing reach-choice behavior on gain and loss trials in which the direction of learning (loss > gain) stands in opposition to predicted direction of preferential processing (gain > loss). For these specific trials, we then examined both the P(optimal choice) and the reach trajectories using a 2×2 GG-corrected RM-ANOVA for the effects of Pair (gain or loss) and Timing (immediate or delay).

Results

VL. The results of the following analysis are depicted graphically in Figure S1 of the online supplemental materials. The $5 \times 2 \times 2$ RM-ANOVA examining the effects of Shape (five possible shapes), Timing (immediate or delay trials), and Target Side (left or right) on RT showed a main effect of Shape, $F(3.01, 72.29) = 4.75, p < .005$, and Timing, $F(1, 24) = 240.97, p < .001$. No other effects or interactions were significant (all $ps > 0.05$). The effect of Shape was driven by a faster RT to the 1.0+ shape (256 ms) than to all other shapes (261 ms to 264 ms). The only significant pairwise comparisons were between the 1.0+ shape and the 0.5− and 0.5+ shapes. This finding suggests that the gain processing advantage we have seen in Experiments 1 and 2, and in the RC task here in Experiment 3, extend even to the case of responding to a single target. The effect of Timing was driven by much faster RTs on delay (225 ms) than immediate (298 ms) trials. This finding is intuitive in that, with a 500-ms delay, participants will have a larger urgency signal and preparedness to move (Cisek, Puskas, & El-Murr, 2009).

The $5 \times 4 \times 2$ RM-ANOVA examining the effects of Shape (5), Timing (2), and Experimental Repetition (4) on RT replicated the main effects of Shape and Timing while showing a main effect of Experimental Repetition, $F(2.24, 51.42) = 8.97, p < .001$, as well as a Timing \times Experimental Repetition interaction, $F(2.71, 62.32) = 13.71, p < .001$. We investigated this interaction by conducting simple main effects GG-corrected RM-ANOVAs of Experimental Repetition in each Timing condition. For the immediate condition, there was an effect of Experimental Repetition on RT, $F(2.49, 57.31) = 15.73, p < .001$, whereas for the delay condition, there was no significant effect of Experimental Repetition on RT. The significant effect of Experimental Repetition for the immediate condition was driven by having significantly faster

RTs (275 ms) on the first repetition compared with the other repetitions (296 ms to 309 ms; all Bonferroni corrected $ps < 0.01$). This gradual increase of RTs with experience reflects participant's ability to learn how to optimize their RTs given the constraints of the task. That is, in this task, for the immediate condition only, participants who could extend their RT as long as possible without exceeding the limit were afforded additional processing time.

RC. As depicted in Figure 7a and b, the $3 \times 2 \times 4$ RM-ANOVA examining the effects of Pair (1.0+/0, 0.5+/0.5-, 0/1.0-) on P(optimal choice) revealed main effects of Timing, $F(1, 24) = 22.67, p < .001$, and Experimental Repetition, $F(1.18, 28.40) = 17.44, p < .001$. There was no significant effect of Pair on P(optimal choice), nor were there any significant interactions. The main effect of Timing was driven by significantly higher P(optimal choice) on delay compared with immediate trials. The main effect of Experimental Repetition was driven by lower P(optimal choice) on Repetition 1 than all other repetitions and a general trend for improved P(optimal choice) across repetitions (Repetition 4 > Repetition 2; significant linear trend across repetitions, $F[1, 24] = 19.05, p < .001$).

As depicted in Figure 7c, the RM-ANOVA examining the effects of Pair (3) and Timing (2) on reach trajectories showed a main effect of Pair, $F(1.72, 41.21) = 14.26, p < .001$, a main effect of Timing, $F(1, 24) = 36.71, p < .001$, and a significant Pair \times Timing interaction, $F(1.70, 40.79) = 13.65, p < .001$. The significant interaction was investigated by performing simple main effects GG-corrected RM-ANOVAs of Pair for each Timing condition. For the immediate trials, there was a significant effect of Pair, $F(1.73, 41.53) = 15.14, p < .001$, whereas for the delay trials, there was no effect of Pair. For the immediate trials, we used Bonferroni corrected post hoc tests to confirm that the effect of Pair was driven by larger reach areas (straighter reaches) on the 1.0+/0 and 0.5+/0.5- trials compared with the 0/1.0- trials ($ps < 0.01$). There was no significant difference between the 1.0+/0 trials and the 0.5+/0.5- trials.

As depicted in Figure 7d, when we extracted the worst-gain trials (the two experimental repetitions having the lowest P[optimal choice] on gain trials) and the best-loss trials (the two experimental repetitions having the highest P[optimal choice] on loss trials) the RM-ANOVA examining the effects of Pair (gain or loss) and Timing (immediate or delay) on P(optimal choice) show a main effect of Pair, $F(1, 24) = 5.83, p < .05$, and a main effect of Timing, $F(1, 24) = 16.25, p < .001$, but no significant interaction. As depicted in Figure 7e, when the same RM-ANOVA is applied to the reach trajectories, we find a significant main effect of Pair, $F(1, 24) = 19.90, p < .001$, a significant main effect of Timing, $F(1, 24) = 32.91, p < .001$, and, importantly, a significant Pair \times Timing interaction, $F(1, 24) = 6.92, p < .05$. We investigated the significant interaction by comparing reach areas for gain and loss pairs in each timing condition. In the immediate condition, the reach area on gain trials is significantly larger than the reach area on loss trials ($p < .001$), whereas in the delay condition, there is no significant difference in reach area between gain and loss trials.

Discussion

In addition to providing an important replication of our results from Experiment 2, Experiment 3 specifically addressed three

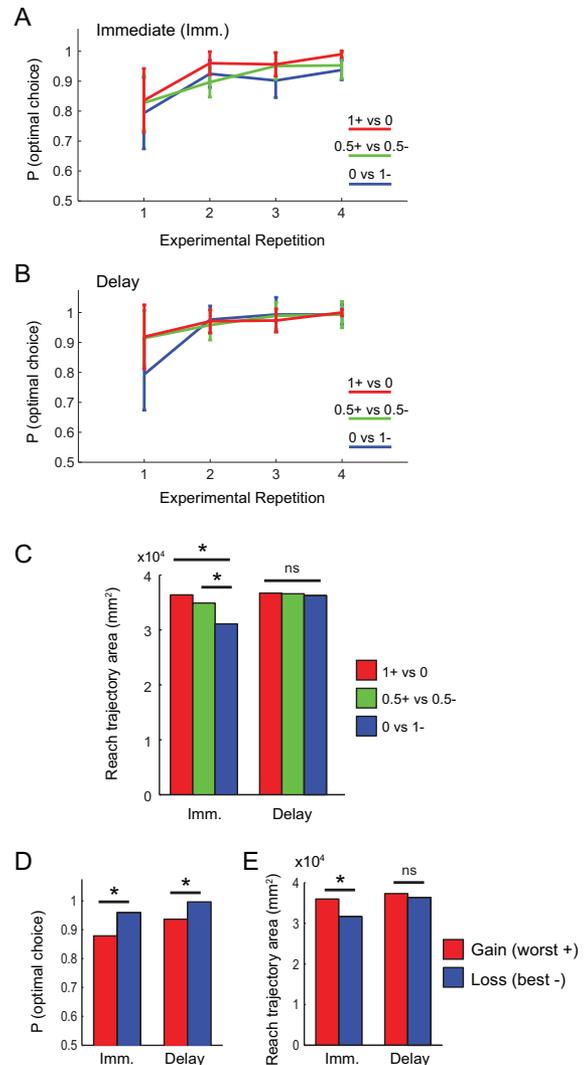


Figure 7. Results from the rapid choice task in Experiment 3 for pairs that differed by an expected value of 1 ($\Delta EV 1.0$). (a, b) Results from the choice behavior showing the P(optimal choice) for the 1+/0 (red), 0.5+/0.5- (green) and 0/1.0- (blue) pairs across experimental repetitions for the immediate (a) and delay (b) trials. Error bars in (a) and (b) represent 95% confidence intervals (CIs). (c) The total area between the curves for each pair type for both immediate (Imm; left bars) and delay (right bars) trials. Asterisk above the Imm trials indicates that the 1.0+/0 and 0.5+/0.5- pairs did not differ significantly from each other, but did differ significantly (corrected $p < .05$) from the 0/1.0- pair. The "ns" above the delay trials indicate that with the delay, there were no significant differences in reach area across the pair types (d) and (e). Results from selecting only those gain trials from the two worst-gain experimental repetitions and only those loss trials from the two best-loss experimental repetitions. (d) P(optimal choice) shows the expected learning advantage for loss trials due to our trial selection (loss > gain, $p < .05$, indicated by asterisk above bar plots). Despite this, (e) shows larger reach areas for gain relative to loss trials on immediate trials (gain > loss, $p < .05$, indicated by asterisk above bar plot) that disappears with delay (ns = no significant difference in area). See the online article for the color version of this figure.

important points. First, we were able to powerfully demonstrate the dissociation between poor loss learning and the processing delay of loss shapes. As shown in Figure 7d and e, by including a much longer experimental session, we were able to extract trials in which the loss choice performance was superior to the gain choice performance, and yet on these trials, the gain processing advantage indexed by reach area was still clearly evident. Second, we made the decision set much more complex than that employed in Experiment 2. Whereas in Experiment 2 there were only two shape-pairs and participants could potentially have used a simple go/no-go strategy, in Experiment 3, we presented 10 shape-pairs. Critically, as demonstrated in Figure 7c (immediate trials), choice performance on the 0.5+/0.5- (in which the expected value difference is the same as shapes in Experiment 2, but importantly, in which neither shape is exclusively optimal or nonoptimal) still showed a profound gain processing advantage. That is, the reach trajectories on this critical 0.5+/0.5- pair were not statistically different than the 1.0+/0 pair, but were significantly different than the 0/1.0- pair. Finally, Experiment 3 allowed us to test the novel hypothesis that by giving participants an additional 500 ms to process the shapes prior to making their choice, this would allow the loss delay to be overcome and behavior on gain and loss trials to become equivalent. As shown in Figure 7c this was exactly the case—with additional time (delay trials), the reaching on loss-pairs becomes as straight as the reaching on gain-pairs.

Modeling Stimulus Response Modulation by Reward and Probability

We developed a neurally constrained model (e.g., Purcell et al., 2010), based on known neural responses to visual stimuli and reward modulations (Kiss et al., 2009; Leathers & Olson, 2012; Pastor-Bernier & Cisek, 2011; Peck et al., 2009; Serences, 2008), to account for the different effects observed in reach trajectories due to target gain, target loss (Experiment 1, Experiment 2, and Experiment 3), and the probability of target payoff (Experiment 1 and Experiment 3). The full description of our model appears in the online supplemental materials (and see Supplemental Figures S5 and S6). Figure 8 shows the predicted results we obtained from our model for each of the three experiments we conducted. As discussed, we were particularly interested in modeling the stimulus responses for neutral, gain, and loss stimuli to determine whether a 100-ms delay in loss relative to gain modulation could produce data consistent with our primary behavioral results. Similarly, we wanted to include a term in our stimulus response model that would account for the observed differences due to the probability of a stimulus producing an outcome. As mentioned in the introduction, stimulus responses to reward-associated visual stimuli are modulated in the earliest visual processing areas, and here we sought to capture mathematically the key characteristics of those known neural responses:

1. Following stimulus onset, visual stimuli with the same physical intensity generate an identical response that rapidly rises from baseline to a peak of activity after approximately 70 ms. This effect is seen in neurophysiological recordings in primary visual cortex (V1; Stanisor, van der Togt, Pennartz, & Roelfsema, 2013) and

even in parietal cortex (Leathers & Olson, 2012; Peck et al., 2009).

2. Following this peak activity, and with continued presentation of the visual stimulus, a second steady-state level of neural activation is achieved after approximately 140 ms, which is lower than the initial response but above the original baseline. This second steady-state is particularly visible in V1 responses (Stanisor et al., 2013).
3. Reward associations affect stimulus response curves (putatively through feedback mechanisms) by changing the level of the postpeak steady-state activity at the earliest after approximately 120 ms. These reward modulations have been reported in V1 (Stanisor et al., 2013) and parietal cortex (Leathers & Olson, 2012; Peck et al., 2009), and inferred from fMRI responses in early visual areas (Lee & Shomstein, 2013).

To capture these different phases of stimulus processing, we modeled the initial response to all our visual stimuli with an identical function that rose from baseline to a peak, before declining back to an above-baseline steady state. We then modeled the reward contribution to the response as an additive (for gain) or subtractive (for loss) effect that started influencing the evolving response after a delay that depended on valence of the reward (thus, the direction of reward modulation and the time to reward modulation were the first two parameters that varied across our experimental conditions). The magnitude of the reward response was simply the probability that a given shape would payoff (thus, reward modulation strength was the third parameter we varied across conditions). By varying only these parameters (which were determined by the exact values used in, or derived from, our experiments), we were able to produce model data that were remarkably consistent with what we observed empirically.

For the RR portion of Experiment 1 (see Figure 8a), because no decision was made, we model the bias we observe in the reach trajectory as directly reflecting the stimulus response differences between the two options presented. Given that there are always exactly two alternatives, and only one effector (the hand) moving through space, we model the single observed bias as the difference between the two stimulus responses. Most notably, our model confirms that it is the pairs containing the high probability gain shape (0.8+) that show the largest bias.

Applying the model to the stimuli from the RC task in Experiment 2 and Experiment 3 produces the predicted data in Figure 8b and c. Because the RC task required participants to make a decision, in modeling these results, we use the generated stimulus responses as input into a classic DDM (Bogacz et al., 2006). Unlike traditional DDMs, which assume decisions are made once evidence for one option has reached some decision threshold, we see our timing constraints as forcing the reach movement to begin before the decision is made, while the computation of the gain and loss information continues to evolve. This is why, we believe, reach trajectories in these cases can provide such a rich source of decision information. For example, at the shortest processing time (rightmost bars in Figure 8b, which corresponds to trials in which the shape

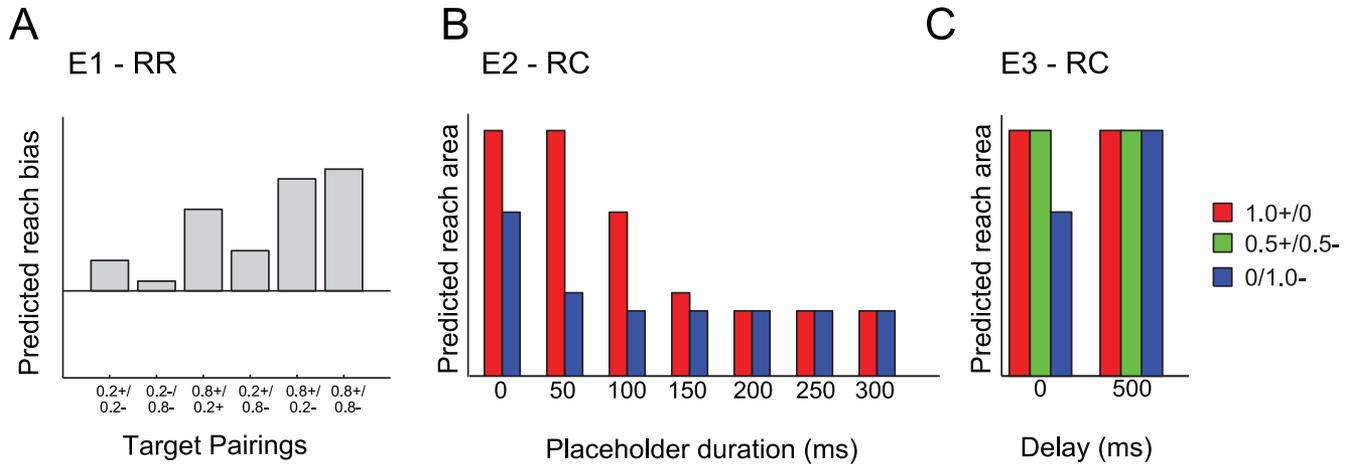


Figure 8. Model predictions for the reach area results from (a) Experiment 1, rapid reaching (E1-RR); (b) Experiment 2, rapid choice (E2-RC), and (c) Experiment 3, rapid choice (E3-RC). Gray bars in (a) indicate the predicted bias toward a reward-associated shape in the absence of a decision (as in E1-RR; compare with Figure 2d), whereas the colored bars in (b) and (c) represent the choice certainty reflected in reach areas for gain (1.0+/0, red), mixed (0.5+/0.5-, green) and loss (0/1.0-, blue) pairs during rapid choice decisions [(b) as in E2-RC; compare with Figure 4e and (c) E3-RC; compare with Figure 7c]. See the online article for the color version of this figure.

information was delayed the most, or a placeholder duration of 300 ms), there has been literally no time prior to reach onset for evidence to accumulate in favor of any particular option. Thus, there is equal competition between the two targets at the time of movement initiation and we predict small reach areas (as the hand initially moves mostly down the middle of the table). Importantly, this means that the accumulated difference for gain and loss pairs is equal for short processing times. With additional processing time prior to reach initiation (e.g., 200 ms, which corresponds to a placeholder duration of 100 ms in Figure 8b), evidence begins to accumulate in favor of one of the two choice options. With this additional time, the reach moves more directly to that location, resulting in larger reach areas. Importantly, the lag in accumulated difference for the loss-pair begins to be evident in the model data, with the bias toward gain targets becoming larger earlier than the bias away from loss targets. At the time of an immediate response (0-ms placeholder in Figure 8b; 0-ms delay in Figure 8c), we see a pronounced gain-loss asymmetry both in the model and behavioral data. But as in Experiment 3, our model predicts that with additional processing time (the extra 500 ms in the delay condition; Figure 8c), this asymmetry is attenuated, as the loss processing catches up to the gain processing and becomes equivalent.

General Discussion

In three separate experiments on rapid visuomotor decisions, high-gain targets received preferential processing relative to equally aversive loss targets. This finding stands in stark contrast to the now-classical EDM work, in which losses loom larger than gains (Kahneman & Tversky, 1979; Tversky & Kahneman, 1992), and also differs from previous RVDM work showing that losses and gains are weighted equally (Trommershäuser et al., 2008). Importantly, it brings clarity to previous PDM work that has either

not investigated the role of loss (for reviews, see Anderson et al., 2013; Chelazzi et al., 2013) or has had inconclusive results regarding the relative effects of gain and loss (O'Brien & Raymond, 2012; Painter et al., 2014; Raymond & O'Brien, 2009; Rutherford et al., 2010). Overwhelmingly, our data suggest that in the earliest stages of perceptual processing, gain targets are processed (i.e., influence decisions) about 100 ms faster than loss targets. Specifically, in Experiment 1, participants' hands were drawn toward targets previously associated with a high probability of gain, with no evidence of avoidance of targets previously associated with loss. In Experiment 2, participants' hands favored a gain target over a neutral target more rapidly than they favored a neutral target over a loss target. In Experiment 3, we replicated the results of Experiment 2 across a wider range of stimuli, and extended the finding, showing that with an additional 500 ms to process targets prior to initiating a response, the gain-loss asymmetry was eliminated—consistent with the hypothesis that loss processing is subject to a temporal lag that can be overcome with additional processing time. These behavioral data were fit well by a neurally inspired model of stimulus response that assumed a 100-ms delay in the reward modulation to a loss-item versus a gain-item, combined with individual item responses that scaled with the probability of payoff.

We situate our findings within theories that contend that (a) all objects are processed as potential action targets, and (b) that multiple potential action targets compete for action selection until an action has been completed (Baldauf & Deubel, 2010; Cisek & Kalaska, 2010). This competition results in a constantly shifting attentional landscape in which “hills” represent behaviorally relevant objects and “valleys” represent objects to be avoided. Because the hand is the physical effector that traverses these landscapes and is attracted toward hills and avoids valleys, detailed measurement of its path through space provides a sensitive readout of the current

landscape. With respect to the current work, the sizes of the hills and valleys are set by the response activity profiles associated with each choice option (see model description in the online supplemental materials for examples). Thus, in our task, in which two items appear at unique spatial locations, the spatial bias toward one of the items is a direct reflection of the amount of preferential processing for that item. For instance, when two items compete equally, the hand moves on a path between the two items. Likewise, when a bias is introduced by the valence of the target (i.e., gain or loss), it is reflected in trajectory deviations of the hand toward favored items and away from items to be avoided.

Why the Snooze of Lose?

If we assume that, on any given trial, the attentional landscape begins flat, and that the initial inputs to the landscape are the simple initial visual responses to stimuli, it is clear that all visual targets will first be represented as hills; that is, as implemented in our model (see the online supplemental materials), all stimuli with equal physical salience produce the same initial positive visual response. Only after sufficient time has passed can the landscape be sculpted by feedback signals potentially carrying reward information. By then, those targets that need to be avoided will have already developed hills, and so will require suppression before avoidance can occur. In contrast, targets that need to be selected already have this positive processing advantage. We propose that the facilitation of the already developed hill for a gain target likely relies on the same neural circuits as those that created the hill in the first place, while the circuits recruited to inhibit the hill for a loss target are distinct and take longer to engage.

Unfortunately, definitive support for separate gain and loss processing circuits in the brain is decidedly mixed. This arises predominantly because loss processing has received almost no empirical attention relative to gain processing. In fact, in a recent meta-analysis examining more than 80 fMRI articles studying the neural computation of subjective value, the authors acknowledged that they did not have sufficient power to look for separate neural responses to appetitive (gain) versus aversive (loss) stimuli (Clithero & Rangel, 2014). Moreover, the processing of loss is further complicated by what loss means in these kinds of experiments. For example, in the current study, participants are losing points that mean they will receive less potential reward. However, this is very different from incurring a cost that would subtract from their actual initial state (e.g., having to pay for poor performance), or experiencing a truly aversive outcome (e.g., a painful stimulus). In fact, both the aforementioned meta-analysis (Clithero & Rangel, 2014) and another recent fMRI article (Delgado, Jou, & Phelps, 2011) show evidence for different neural responses based on the type of reward or penalty received. Specifically, primary reinforcers, like food (positive) or shock (negative), are processed differently than secondary reinforcers, like net gains or losses of money. Despite these notable limitations, the idea of separate circuits for gain and loss processing is supported by a recent meta-analysis of human neuroimaging research (Levin et al., 2012), recent modeling work integrating real fMRI data into accumulator models of decision making (Basten, Biele, Heekeren, & Fiebach, 2010), and a review of some of the neurophysiological signals in response to aversive choice options (Hikosaka, 2010).

More generally, the idea that inhibition takes longer than facilitation resonates with previous work showing that both the eye (Van der Stigchel, Meeter, & Theeuwes, 2006) and hand (Neyedli & Welsh, 2012; Welsh & Elliott, 2004; Welsh, Neyedli, & Tremblay, 2013) are initially drawn toward distracting stimuli, and only later, after sufficient processing time, move away from them. Here, the idea is that to inhibit a task-irrelevant item, it must be first processed as an item (i.e., resulting in a hill in an attentional landscape) before it can be avoided—a conjecture that is similar to that accounting for selection versus inhibition timing asymmetries in a visual search task (Arita, Carlisle, & Woodman, 2012). If true, this stance would have significant implications for other aspects of motor behavior—most notably, it suggests that to plan to move around a potential obstacle would take longer than to plan to move toward a target, an interesting question for future work.

The idea of a spatial relevance map, shaped by potentially discrete neural inputs from multiple biasing factors (e.g., physical salience, gain, loss), is consistent with research suggesting the convergence of decision-making variables onto one specific brain area (potentially for each effector enacting a decision). Many theorists consider the parietal cortex in the primate brain to be the most likely site for the neural instantiation of a relevance map, and ultimately where the competition between potential action targets is resolved (Bisley & Goldberg, 2010; Chapman, Gallivan, Culham, & Goodale, 2011; Chelazzi et al., 2014; Cisek & Kalaska, 2010; Fecteau & Munoz, 2006; Gottlieb, 2012; Lee & Shomstein, 2014; Ptak, 2012; Summerfield et al., 2012). Importantly, however, our current findings highlight that understanding decision making in the context of relevance maps is not just to acknowledge that such maps exist but also to explore how the map is formed and shaped across time. As outlined earlier, the physical salience of a stimulus plays an important role in the initial formation of its relevance signal (Awh, Belopolsky, & Theeuwes, 2012). Recently we have shown using RR that this physical salience response occurs much earlier than an enumeration or probability of selection signal (Wood et al., 2011), which, in turn, occurs earlier than the processing of symbolic numerals that convey the same information (Chapman et al., 2014). These are specific examples of the more general point that it is not just the spatial information contained in the relevance map, but its evolving dynamics, that will shape behavior.

Extensions

The hypothesized delayed loss processing, in conjunction with the conceptualization of dynamic attentional landscapes, have explanatory power that extends to previous results from each of the decision domains we have discussed. For example, recent results exploring loss aversion in EDM tasks have found longer RTs for choices involving loss options (Xue et al., 2009; Yechiam & Telpaz, 2013). Although these authors contend this is because loss options demand more attention (another manifestation of loss aversion), we propose this RT effect is additional evidence for the delayed loss processing identified here. In a simple PDM task, our theoretical framework explains why eye movements drawn toward distractor items that had, on the previous trial, been positively rewarding (Hickey & van Zoest, 2012). Appealing to the idea of attentional landscapes and a gain-target advantage, those previously rewarded targets would receive a processing boost. Accord-

ing to our model, the sensory response to these previously rewarded targets would, on the next trial when they are now a distractor, develop more quickly than the sensory response to the actual target location. Thus, when the eye movement is launched, the relative strength of the distractor response would pull gaze toward its location. Here, the effector is the eye, but the principle of movements being biased by activity profiles competing in parallel is the same. Finally, this framework, in conjunction with the proposed loss delay, may also help explain a discrepancy that exists in the small body of RVD work. Recall that, in general, past RVD studies have found that losses and gains are treated equivalently (Trommershäuser et al., 2003a, 2003b, 2008). Recently, however, work has shown this equivalence only emerges after significant experience with the task (Neyedli & Welsh, 2015)—prior to achieving gain–loss symmetry, participants relatively undervalue loss. Following from the current results, and situating these findings in the framework of attentional landscapes, we would argue that in the initial trials, participants process the gain region faster than the loss region. As a result, the relative landscape results in a reach that nonoptimally moves toward gain, but not enough away from loss. This account does not explain why, over time, participants do learn to avoid the loss regions optimally, but one potential explanation is that an implicit strategy emerges that compensates for loss delay. It is possible that in our task, in which an explicit choice is required between specific choice options that vary unpredictably across trials, such an implicit strategy cannot be learned.

Conclusion

Decision making appears to be the deceptively simple problem of accounting for the potential gains or losses associated with each choice. However, the problem is complicated in everyday life by uncertainty over what information is relevant, by value assignments that change over time, and by the wide range of time scales over which relevance can be computed, from milliseconds to years. Here we introduce an RR task with parametric temporal delays as a means of exploring value associations over the shortest time scale. Our results suggest that the extreme gain-sensitivity we observe is a consequence of loss information being delayed relative to equivalent gain information. We interpret our results in the context of neurobehavioral frameworks arguing for a common relevance map as the arena in which choice options compete for action execution. As explored in our model, we believe that reward effects can be accounted for by modulations to the early sensory responses that serve as a primary input into these relevance maps. By virtue of being a physical stimulus, all targets are initially positively activated. Targets associated with high-probability gains build off of this initial activation due to relatively rapid feedback, which enhances their early sensory response and leads to faster actions toward them. Loss-related targets, however, must overcome their initial sensory activation and are only suppressed when a slower inhibitory feedback signal arrives, leading to delayed actions away from them. As is typified by our experimental design, results and modeling approach, we emphasize the importance of not only studying the behavioral outcome of a choice, but in tracking its progress over time.

References

- Anderson, B. A., Laurent, P. A., & Yantis, S. (2011). Value-driven attentional capture. *PNAS Proceedings of the National Academy of Sciences of the United States of America*, *108*, 10367–10371. <http://dx.doi.org/10.1073/pnas.1104047108>
- Anderson, B. A., Laurent, P. A., & Yantis, S. (2013). Neural mechanisms of value-driven attentional capture. *Journal of Vision*, *13*, 906. <http://dx.doi.org/10.1167/13.9.906>
- Arita, J. T., Carlisle, N. B., & Woodman, G. F. (2012). Templates for rejection: Configuring attention to ignore task-irrelevant features. *Journal of Experimental Psychology: Human Perception and Performance*, *38*, 580–584. <http://dx.doi.org/10.1037/a0027885>
- Awh, E., Belopolsky, A. V., & Theeuwes, J. (2012). Top-down versus bottom-up attentional control: A failed theoretical dichotomy. *Trends in Cognitive Sciences*, *16*, 437–443. <http://dx.doi.org/10.1016/j.tics.2012.06.010>
- Baldauf, D., & Deubel, H. (2010). Attentional landscapes in reaching and grasping. *Vision Research*, *50*, 999–1013. <http://dx.doi.org/10.1016/j.visres.2010.02.008>
- Basten, U., Biele, G., Heekeren, H. R., & Fiebach, C. J. (2010). How the brain integrates costs and benefits during decision making. *PNAS Proceedings of the National Academy of Sciences of the United States of America*, *107*, 21767–21772. <http://dx.doi.org/10.1073/pnas.0908104107>
- Bisley, J. W., & Goldberg, M. E. (2010). Attention, intention, and priority in the parietal lobe. *Annual Review of Neuroscience*, *33*, 1–21. <http://dx.doi.org/10.1146/annurev-neuro-060909-152823>
- Bogacz, R., Brown, E., Moehlis, J., Holmes, P., & Cohen, J. D. (2006). The physics of optimal decision making: A formal analysis of models of performance in two-alternative forced-choice tasks. *Psychological Review*, *113*, 700–765. <http://dx.doi.org/10.1037/0033-295X.113.4.700>
- Brainard, D. H. (1997). The psychophysics toolbox. *Spatial Vision*, *10*, 433–436. <http://dx.doi.org/10.1163/156856897X00357>
- Chapman, C. S., Gallivan, J. P., Culham, J. C., & Goodale, M. A. (2011). Mental blocks: fMRI reveals top-down modulation of early visual cortex when obstacles interfere with grasp planning. *Neuropsychologia*, *49*, 1703–1717. <http://dx.doi.org/10.1016/j.neuropsychologia.2011.02.048>
- Chapman, C. S., Gallivan, J. P., & Enns, J. T. (2015). Separating value from selection frequency in rapid reaching biases to visual targets. *Visual Cognition*, *23*, 249–271. <http://dx.doi.org/10.1080/13506285.2014.976604>
- 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*, 30. <http://dx.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*, 168–176. <http://dx.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*, 130–134. <http://dx.doi.org/10.1016/j.bbr.2010.05.012>
- Chelazzi, L., Estocinova, J., Calletti, R., Lo Gerfo, E., Sani, I., Della Libera, C., & Santandrea, E. (2014). Altering spatial priority maps via reward-based learning. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, *34*, 8594–8604. <http://dx.doi.org/10.1523/JNEUROSCI.0277-14.2014>

- Chelazzi, L., Perlato, A., Santandrea, E., & Della Libera, C. (2013). Rewards teach visual selective attention. *Vision Research*, *85*, 58–72. <http://dx.doi.org/10.1016/j.visres.2012.12.005>
- Cisek, P. (2012). Making decisions through a distributed consensus. *Current Opinion in Neurobiology*, *22*, 927–936. <http://dx.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. <http://dx.doi.org/10.1146/annurev.neuro.051508.135409>
- Cisek, P., Puskas, G. A., & El-Murr, S. (2009). Decisions in changing conditions: The urgency-gating model. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, *29*, 11560–11571. <http://dx.doi.org/10.1523/JNEUROSCI.1844-09.2009>
- Clithero, J. A., & Rangel, A. (2014). Informatic parcellation of the network involved in the computation of subjective value. *Social Cognitive and Affective Neuroscience*, *9*, 1289–1302. <http://dx.doi.org/10.1093/scan/nst106>
- Delgado, M. R., Jou, R. L., & Phelps, E. A. (2011). Neural systems underlying aversive conditioning in humans with primary and secondary reinforcers. *Frontiers in Neuroscience*, *5*, 71. <http://dx.doi.org/10.3389/fnins.2011.00071>
- Della Libera, C., & Chelazzi, L. (2009). Learning to attend and to ignore is a matter of gains and losses. *Psychological Science*, *20*, 778–784. <http://dx.doi.org/10.1111/j.1467-9280.2009.02360.x>
- Dshemuchadse, M., Scherbaum, S., & Goshke, T. (2013). How decisions emerge: Action dynamics in intertemporal decision making. *Journal of Experimental Psychology: General*, *142*, 93–100. <http://dx.doi.org/10.1037/a0028499>
- Fecteau, J. H., & Munoz, D. P. (2006). Saliency, relevance, and firing: A priority map for target selection. *Trends in Cognitive Sciences*, *10*, 382–390. <http://dx.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. <http://dx.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. <http://dx.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*, 803–811. <http://dx.doi.org/10.1177/0956797611408733>
- Gold, J. I., & Shadlen, M. N. (2007). The neural basis of decision making. *Annual Review of Neuroscience*, *30*, 535–574. <http://dx.doi.org/10.1146/annurev.neuro.29.051605.113038>
- Gottlieb, J. (2012). Attention, learning, and the value of information. *Neuron*, *76*, 281–295. <http://dx.doi.org/10.1016/j.neuron.2012.09.034>
- Heekeren, H. R., Marrett, S., & Ungerleider, L. G. (2008). The neural systems that mediate human perceptual decision making. *Nature Reviews Neuroscience*, *9*, 467–479. <http://dx.doi.org/10.1038/nrn2374>
- Hertwig, R., & Erev, I. (2009). The description-experience gap in risky choice. *Trends in Cognitive Sciences*, *13*, 517–523. <http://dx.doi.org/10.1016/j.tics.2009.09.004>
- Hickey, C., & van Zoest, W. (2012). Reward creates oculomotor saliency. *Current Biology*, *22*, R219–R220. <http://dx.doi.org/10.1016/j.cub.2012.02.007>
- Hikosaka, O. (2010). The habenula: From stress evasion to value-based decision-making. *Nature Reviews Neuroscience*, *11*, 503–513. <http://dx.doi.org/10.1038/nrn2866>
- Jarvstad, A., Hahn, U., Rushton, S. K., & Warren, P. A. (2013). Perceptuo-motor, cognitive, and description-based decision-making seem equally good. *PNAS Proceedings of the National Academy of Sciences of the United States of America*, *110*, 16271–16276. <http://dx.doi.org/10.1073/pnas.1300239110>
- Kahneman, D. (2003). A perspective on judgment and choice: Mapping bounded rationality. *American Psychologist*, *58*, 697–720. <http://dx.doi.org/10.1037/0003-066X.58.9.697>
- Kahneman, D. (2011). *Thinking, fast and slow*. New York, NY: Macmillan.
- Kahneman, D., & Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica*, *47*, 263–291. <http://dx.doi.org/10.2307/1914185>
- Kiss, M., Driver, J., & Eimer, M. (2009). Reward priority of visual target singletons modulates event-related potential signatures of attentional selection. *Psychological Science*, *20*, 245–251. <http://dx.doi.org/10.1111/j.1467-9280.2009.02281.x>
- Kleiner, M., Brainard, D., & Pelli, D. (2007). What's new in Psychtoolbox-3? [Abstract]. *Perception*, *36* (Suppl.), 14.
- Leathers, M. L., & Olson, C. R. (2012). In monkeys making value-based decisions, LIP neurons encode cue saliency and not action value. *Science (New York, N. Y.)*, *338*, 132–135. <http://dx.doi.org/10.1126/science.1226405>
- Lee, J., & Shomstein, S. (2013). The differential effects of reward on space- and object-based attentional allocation. *The Journal of Neuroscience*, *33*, 10625–10633. <http://dx.doi.org/10.1523/JNEUROSCI.5575-12.2013>
- Lee, J., & Shomstein, S. (2014). Reward-based transfer from bottom-up to top-down search tasks. *Psychological Science*, *25*, 466–475. <http://dx.doi.org/10.1177/0956797613509284>
- Levin, I. P., Xue, G., Weller, J. A., Reimann, M., Lauriola, M., & Bechara, A. (2012). A neuropsychological approach to understanding risk-taking for potential gains and losses. *Frontiers in Neuroscience*, *6*, 15. <http://dx.doi.org/10.3389/fnins.2012.00015>
- McKinstry, C., Dale, R., & Spivey, M. J. (2008). Action dynamics reveal parallel competition in decision making. *Psychological Science*, *19*, 22–24. <http://dx.doi.org/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*, 1456–1465. <http://dx.doi.org/10.1177/0956797612473485>
- Neyedli, H. F., & Welsh, T. N. (2012). The processes of facilitation and inhibition in a cue-target paradigm: Insight from movement trajectory deviations. *Acta Psychologica*, *139*, 159–165. <http://dx.doi.org/10.1016/j.actpsy.2011.11.001>
- 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*, 638–645. <http://dx.doi.org/10.1037/a0030518>
- Neyedli, H. F., & Welsh, T. N. (2015). Experience and net worth affects optimality in a motor decision task. *Motor Control*, *19*, 75–89. <http://dx.doi.org/10.1123/mc.2013-0024>
- O'Brien, J. L., & Raymond, J. E. (2012). Learned predictiveness speeds visual processing. *Psychological Science*, *23*, 359–363. <http://dx.doi.org/10.1177/0956797611429800>
- O'Hara, D., Dale, R., Piiroinen, P. T., & Connolly, F. (2013). Local dynamics in decision making: The evolution of preference within and across decisions. *Scientific Reports*, *3*, 2210. <http://dx.doi.org/10.1038/srep02210>
- Painter, D. R., Kritikos, A., & Raymond, J. E. (2014). Value learning modulates goal-directed actions. *Quarterly Journal of Experimental Psychology (2006)*, *67*, 1166–1175. <http://dx.doi.org/10.1080/17470218.2013.848913>
- Pastor-Bernier, A., & Cisek, P. (2011). Neural correlates of biased competition in premotor cortex. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, *31*, 7083–7088. <http://dx.doi.org/10.1523/JNEUROSCI.5681-10.2011>

- Peck, C. J., Jangraw, D. C., Suzuki, M., Efem, R., & Gottlieb, J. (2009). Reward modulates attention independently of action value in posterior parietal cortex. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, *29*, 11182–11191. <http://dx.doi.org/10.1523/JNEUROSCI.1929-09.2009>
- Pelli, D. G. (1997). The VideoToolbox software for visual psychophysics: Transforming numbers into movies. *Spatial Vision*, *10*, 437–442. <http://dx.doi.org/10.1163/156856897X00366>
- Ptak, R. (2012). The frontoparietal attention network of the human brain: Action, saliency, and a priority map of the environment. *The Neuroscientist*, *18*, 502–515. <http://dx.doi.org/10.1177/1073858411409051>
- Purcell, B. A., Heitz, R. P., Cohen, J. Y., Schall, J. D., Logan, G. D., & Palmeri, T. J. (2010). Neurally constrained modeling of perceptual decision making. *Psychological Review*, *117*, 1113–1143. <http://dx.doi.org/10.1037/a0020311>
- Rakow, T., & Newell, B. R. (2010). Degrees of uncertainty: An overview and framework for future research on experience-based choice. *Journal of Behavioral Decision Making*, *23*, 1–14. <http://dx.doi.org/10.1002/bdm.681>
- 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*, 981–988. <http://dx.doi.org/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*, 263–266. <http://dx.doi.org/10.1038/nature08275>
- Rutherford, H. J., O'Brien, J. L., & Raymond, J. E. (2010). Value associations of irrelevant stimuli modify rapid visual orienting. *Psychonomic Bulletin & Review*, *17*, 536–542. <http://dx.doi.org/10.3758/PBR.17.4.536>
- Scherbaum, S., Dshemuchadse, M., Fischer, R., & Goschke, T. (2010). How decisions evolve: The temporal dynamics of action selection. *Cognition*, *115*, 407–416. <http://dx.doi.org/10.1016/j.cognition.2010.02.004>
- Serences, J. T. (2008). Value-based modulations in human visual cortex. *Neuron*, *60*, 1169–1181. <http://dx.doi.org/10.1016/j.neuron.2008.10.051>
- Song, J. H., & Nakayama, K. (2009). Hidden cognitive states revealed in choice reaching tasks. *Trends in Cognitive Sciences*, *13*, 360–366. <http://dx.doi.org/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*, 207–211. <http://dx.doi.org/10.1111/j.1467-8721.2006.00437.x>
- Stanisor, L., van der Togt, C., Pennartz, C. M., & Roelfsema, P. R. (2013). A unified selection signal for attention and reward in primary visual cortex. *PNAS Proceedings of the National Academy of Sciences of the United States of America*, *110*, 9136–9141. <http://dx.doi.org/10.1073/pnas.1300117110>
- Summerfield, C., Tsetsos, K., Van Maanen, L., Dreher, J., & Ruff, C. (2012). Building bridges between perceptual and economic decision-making: Neural and computational mechanisms. *Frontiers in Neuroscience*, *6*, 70. <http://dx.doi.org/10.3389/fnins.2012.00070>
- Trommershäuser, J., Maloney, L. T., & Landy, M. S. (2003a). Statistical decision theory and the selection of rapid, goal-directed movements. *Journal of the Optical Society of America. A, Optics, Image Science, and Vision*, *20*, 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*, 255–275. <http://dx.doi.org/10.1163/156856803322467527>
- Trommershäuser, J., Maloney, L. T., & Landy, M. S. (2008). Decision making, movement planning and statistical decision theory. *Trends in Cognitive Sciences*, *12*, 291–297. <http://dx.doi.org/10.1016/j.tics.2008.04.010>
- Tversky, A., & Kahneman, D. (1992). Advances in prospect theory: Cumulative representation of uncertainty. *Journal of Risk and Uncertainty*, *5*, 297–323. <http://dx.doi.org/10.1007/BF00122574>
- Van der Stigchel, S., Meeter, M., & Theeuwes, J. (2006). Eye movement trajectories and what they tell us. *Neuroscience and Biobehavioral Reviews*, *30*, 666–679. <http://dx.doi.org/10.1016/j.neubiorev.2005.12.001>
- Welsh, T. N., & Elliott, D. (2004). Movement trajectories in the presence of a distracting stimulus: Evidence for a response activation model of selective reaching. *The Quarterly Journal of Experimental Psychology. A, Human Experimental Psychology*, *57*, 1031–1057. <http://dx.doi.org/10.1080/02724980343000666>
- Welsh, T. N., Neyedli, H., & Tremblay, L. (2013). Refining the time course of facilitation and inhibition in attention and action. *Neuroscience Letters*, *554*, 6–10. <http://dx.doi.org/10.1016/j.neulet.2013.08.055>
- Wolpert, D. M., & Landy, M. S. (2012). Motor control is decision-making. *Current Opinion in Neurobiology*, *22*, 996–1003. <http://dx.doi.org/10.1016/j.conb.2012.05.003>
- 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*, 16. <http://dx.doi.org/10.1167/11.10.16>
- Xue, G., Lu, Z., Levin, I. P., Weller, J. A., Li, X., & Bechara, A. (2009). Functional dissociations of risk and reward processing in the medial prefrontal cortex. *Cerebral Cortex*, *19*, 1019–1027. <http://dx.doi.org/10.1093/cercor/bhn147>
- Yechiam, E., & Hochman, G. (2013). Losses as modulators of attention: Review and analysis of the unique effects of losses over gains. *Psychological Bulletin*, *139*, 497–518. <http://dx.doi.org/10.1037/a0029383>
- Yechiam, E., & Hochman, G. (2014). Loss attention in a dual-task setting. *Psychological Science*, *25*, 494–502. <http://dx.doi.org/10.1177/0956797613510725>
- Yechiam, E., & Telpaz, A. (2013). Losses induce consistency in risk taking even without loss aversion. *Journal of Behavioral Decision Making*, *26*, 31–40.

Received March 17, 2014

Revision received April 27, 2015

Accepted April 28, 2015 ■