

Sudden onsets reflexively drive spatial attention, but those that predict reward do more

Devin R. Butler

Department of Psychology, University of Maryland,
College Park, MD, USA



Michael A. Grubb

Department of Psychology, Trinity College,
Hartford, CT, USA



The sudden appearance of an unexpected object elicits the automatic allocation of spatial attention. Even without eye movements, effortless, but transient, improvements in perception occur at the onset location. Much is known about the consequences of such exogenously elicited shifts of covert attention, but most research has used stimuli that carry very little, if any, additional information. In everyday life, attention is captured by sudden onsets that, due to past experience, alert us to more than just their appearance. An abundance of recent work has shed light on the interaction of associative learning and attention, leading to refinements in current models of attentional control; in this study, we tested two hypotheses concerning the efficacy of meaning-imbued onsets, specifically those that predict reward, to drive the reflexive allocation of covert spatial attention and to improve task performance more generally. First, spatially uninformative, abrupt-onset cues that are predictive of reward may elicit the involuntary allocation of attention more effectively than nonreward-predictive onsets; second, the presence of peripheral cues that are predictive of reward, regardless of validity, may impact global attentional processes in a spatially nonspecific manner. We paired monetary reward with one of two luminance-defined, abrupt-onset cues and measured each cue's ability to modulate performance in a visual task. Replicating research with nonmeaning-imbued stimuli, both kinds of abrupt onsets reflexively improved visual perception at attended, relative to unattended, locations. However, when features of the onset predicted the availability of monetary reward, enhancements in perception become less constrained, spreading rapidly to other task-relevant locations.

rectangle swiftly moves in from the upper-right-hand corner of your laptop screen. Within 100 ms of the notification's abrupt appearance, and without an accompanying eye movement, your visual system reflexively allocates attention to the digital notification's location, transiently biasing the processing of visual information emanating from that part of the visual field. Well-controlled, laboratory-based research has revealed a tremendous amount regarding the behavioral and neural consequences of such covert shifts of spatial attention (i.e., those occurring in the absence of observable eye movements), demonstrating that task-irrelevant, abrupt-onset stimuli, even those as simple as small colored circles, reliably elicit the involuntary allocation of visuospatial attention (for a comprehensive review, see Carrasco, 2011).

As conveyed by the opening example, however, exogenous shifts of spatial attention in everyday life are often elicited by objects in our visual world that, through learning, have become associated with specific kinds of information. For example, if you use Gmail, a pop-up in the lower-right-hand corner captures your attention. Because you have encountered this pop-up many times in the past and have learned its meaning, this sudden onset also alerts you that a chat from a friend has appeared, even before you move your eyes to identify the sender and read the message. Might such meaning-imbued onsets, particularly those that signal the impending delivery of a rewarding stimulus, be more effective at biasing the allocation of spatial attention than those that carry no such associations?

Despite an abundance of interest in and research on the interrelationships between reward, learning, and visual attention over the past decade (for reviews, see Anderson, 2018; Awh, Belopolsky, & Theeuwes, 2012; Chelazzi, Perlato, Santandrea, & Della Libera, 2013; Failing & Theeuwes, 2018; Le Pelley, Mitchell, Beesley, George, & Wills, 2016), very little is known about the impact of reward-based associative learning on an abrupt-onset cue's ability to elicit the reflexive allocation of covert spatial attention. Although some

Introduction

As you stare blankly at the cursor, pondering the opening sentence of your next manuscript, you suddenly notice a familiar intrusion: a gray

Citation: Butler, D. R., & Grubb, M. A. (2020). Sudden onsets reflexively drive spatial attention, but those that predict reward do more. *Journal of Vision*, 20(7):30, 1–10, <https://doi.org/10.1167/jov.20.7.30>.



studies have used an exogenous spatial cueing paradigm and manipulated the availability of reward from block-to-block (Bucker & Theeuwes, 2014; Engelmann & Pessoa, 2007), and another has paired exogenous cues with monetary reward while voluntary spatial attention was reliably preallocated to the target location (Munneke, Hoppenbrouwers, & Theeuwes, 2015), linking reward with specific features of an abrupt-onset cue itself and comparing the cue's efficacy with that of a nonreward-predictive onset in the absence of spatially informative endogenous cues would provide the empirical data needed to address the question asked earlier. Randomly interleaving the two cue types would also create rapid, trial-to-trial manipulation of reward availability in a way that does not rely on block-to-block changes in the incentive structure, thus allowing researchers to probe the transient impact of reward on task performance more generally (i.e., anytime the reward-predictive cue appears, regardless of its location).

In this study, we aimed to fill this gap in the literature by testing two related hypotheses. First, spatially uninformative, abrupt-onset cues that are predictive of monetary reward (as a result of associative learning), may elicit the involuntary allocation of covert spatial attention more effectively than nonreward-predictive onsets. Given the transient nature of exogenous attention (Muller & Rabbitt, 1989; Nakayama & Mackeben, 1989), we focused on the temporal dynamics of the cueing effect in addition to its overall magnitude: if reward-predictive cues capture attention more strongly, we would expect to see significantly larger attentional effects for reward-predictive cues, relative to nonreward-predictive cues; similarly, if reward-predictive onsets capture attention more quickly than do nonreward-predictive onsets, we would expect peak attentional effects for the reward-predictive cues to manifest earlier in time. Second, the presence of peripheral cues that are predictive of monetary reward, regardless of their validity, may impact global attentional processes in a spatially nonspecific manner (e.g., by transiently increasing arousal); if so, we would expect to see improvements in task performance when reward-predictive cues were present, relative to nonreward-predictive cues, irrespective of their location. Both hypotheses are inspired by a large body of research on reinforcement learning showing that the appearance of reward-predictive stimuli elicit a rapidly occurring, dopamine-dependent, reward-prediction signal (for an overview, see: Daw & Tobler, 2013). In the context of reward-predictive, abrupt onset cues, such stimuli should do double duty (i.e., they should both bias the allocation of spatial attention and engender a reward prediction signal), something that nonreward-predictive, abrupt onset cues should not do. Such a dopamine-dependent, reward-prediction signal may serve to enhance the concomitant selective

attention effect generated by these same visual cues (hypothesis 1), it may improve performance globally by, for example, modulating arousal (hypothesis 2), or it may have no impact on behavior in our visual task (the null hypothesis). Psychophysical data, collected using robust methods for manipulating exogenous spatial attention and measuring the resultant impact on behavior, are necessary to guide future research on any potential interaction and to make specific hypotheses related to the neural mechanisms underlying any observed behavioral effects.

To test these hypotheses, we paired monetary reward with one of two luminance-defined, abrupt-onset peripheral cues, each of which was presented at four distinct cue-to-target onset asynchrony durations (Experiment 1). On each trial, two grating stimuli were briefly presented, one in the left and one in the right visual periphery, and participants were asked to report the orientation of the target grating, as indicated by a response cue. Irrespective of cue type (i.e., reward-predictive or nonreward-predictive), peripheral onsets appeared near the location of the forthcoming target (valid trials) and near the location of the forthcoming distractor (invalid trials) an equal number of times, thus providing no spatially relevant information. Irrespective of cue validity, a monetary bonus was delivered after a correct judgment if the cue on that trial was reward-predictive; no bonus was delivered on nonreward-predictive cue trials. This design allowed us to characterize the impact of reward learning on the reflexive allocation of covert exogenous spatial attention and to assess the temporal dynamics of reward's impact on task performance more generally.

To preview our main results, we found evidence for our second hypothesis: reward-predictive onsets enhanced task performance, regardless of their location, but this improvement took time to manifest, becoming statistically significant only at our final stimulus-onset-asynchrony (SOA). In a follow-up study (Experiment 2), we replicated our primary finding in an independent group of naive observers and evaluated whether this boost in performance was maintained at longer SOA durations.

Methods: Experiment 1

Participants

Seventy-two participants (39 female, average age 19.59 years) completed Experiment 1. Six participants failed to perform the task above chance (see Assessing chance accuracy section), leaving 66 participants for all subsequent analyses. Experimental procedures were approved by Trinity College's institutional review board. All participants provided written, informed

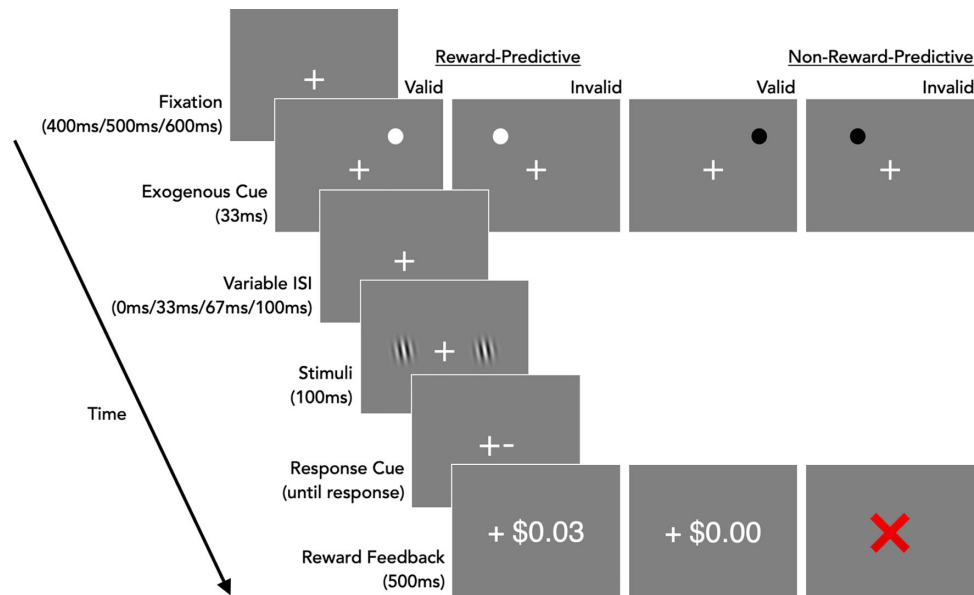


Figure 1. Trial sequence, [Experiment 1](#). Interstimulus-interval. In this example, the white cue was reward-predictive; cue color and reward predictability were counterbalanced across participants. See Methods: [Experiment 1](#) for full details.

consent, and received a base payment of \$5, plus a performance-contingent bonus for participating (see Reward manipulation section, average bonus ~\$14).

Psychophysical task

On each trial, participants reported the orientation of a target grating in a two-alternative forced choice task ([Figure 1](#)). Each trial featured two Gabor patches—one distractor, one target—positioned 8° of visual angle (d.v.a.) to the left and right of a central fixation cross (0.5×0.5 d.v.a.). The stimuli were sinusoidal gratings enveloped by a Gaussian (diameter, 4 d.v.a.; spatial frequency, 2 c/°; contrast, 100%) displayed on a gray background for 100 ms. Each grating was randomly and independently rotated 30° clockwise or counterclockwise of vertical. After the presentation of the Gabor patches, a response cue appeared either to the left or right of fixation, indicating the grating for which the orientation discrimination judgment should be made. The target grating appeared equally often on the left and right sides, and the response window was virtually unlimited, as participants had up to 20 seconds to make a response before the trial timed out. Reward feedback was displayed for 500 ms: “+ \$0.03” if the response was correct and the peripheral cue was reward-predictive (see Attentional manipulation and Reward manipulation sections), “+ \$0.00” if the response was correct but the peripheral cue was nonreward-predictive, or a red X if the response was incorrect. After a variable period of fixation (400, 500, or 600 ms), the next trial commenced. There were

eight blocks of 152 trials each, and participants saw a running total of bonus earnings at the end of each block. Participants completed 40 practice trials before starting the main experimental blocks.

Attentional manipulation

We manipulated covert exogenous spatial attention by briefly presenting a peripheral cue. The cue was a small circle (radius, 0.25 d.v.a.) that appeared for 33 ms, 2.5 d.v.a. above the forthcoming target grating’s location, on 50% of the trials (valid condition), or 2.5 d.v.a. above the location of the forthcoming distractor for the other half of the trials (invalid condition). The cue was thus spatially uninformative, as it provided no information about the location of the target grating. To investigate the temporal dynamics of the peripheral cue’s impact on task performance, we included four distinct interstimulus-interval durations (0, 33, 67, and 100 ms), yielding four SOA conditions: 33, 67, 100, and 133 ms.

When studying covert spatial attention, one must consider methodological issues carefully (see [Carrasco, 2011](#) for detailed treatment of methodological concerns). Although we did not directly monitor adherence to fixation instructions, we have three reasons to be confident that our results are not due to overt shifts of attention: (1) given the SOAs in our design, participants do not have time to make a goal-directed saccade. The 133 ms SOA coupled with a 100 ms presentation time would mean that the Gabor stimuli would have offset by the time the eyes had moved to the

distractor or the target location. (2) The exogenous cue provides no information about the target location: on half of the trials, the cue appears on the same side as the target, and on the other half, the cue appears on the side of the distractor. Participants were explicitly told this information before beginning the experiment, and they were also advised to maintain fixation to maximize their chance for successful task performance, as keeping their eyes on the center of the screen allows for their endogenous attention to be distributed across the entire screen. (3) Participants were rewarded for making correct responses (on reward-predictive cue trials, which were randomly interleaved with the nonreward-predictive cue trials), so they had a financial incentive to maintain fixation, given that the abrupt onset cue provided no information as to where the target would appear.

Reward manipulation

On half of the trials, the peripheral cue was white, and on the other half, it was black. Counterbalanced across participants, one cue color was predictive of monetary reward, irrespective of the cue's validity: after a correct response, participants received a bonus of \$0.03. Bonus earnings accumulated throughout the task and were paid in cash at the end of the session. The other cue color was nonreward-predictive, yielding no bonus earnings.

Assessing chance accuracy

Six participants performed the task below chance (proportion correct range, 0.4901, 0.5123). There were 1216 total trials, and we simulated chance performance by taking the mean of 1216 random draws from the set [0 1]. We repeated this procedure 10,000 times to build a distribution of simulated guessing behavior and excluded any participant whose overall accuracy was within the 95% confidence interval of this distribution (proportion correct ≤ 0.5280). This procedure allows us to identify participants who performed as if they were randomly pressing buttons during the task, either due to a misunderstanding or to a lack of motivation.

Bootstrapped confidence intervals

To compute 95% confidence intervals, we randomly sampled 66 observers (with replacement) and recomputed the mean exogenous cueing effect for the bootstrapped sample of observers (i.e., mean change in accuracy and response time [RT], on valid relative to invalid trials, Figure 2). We then repeated this procedure 10,000 times to build a distribution

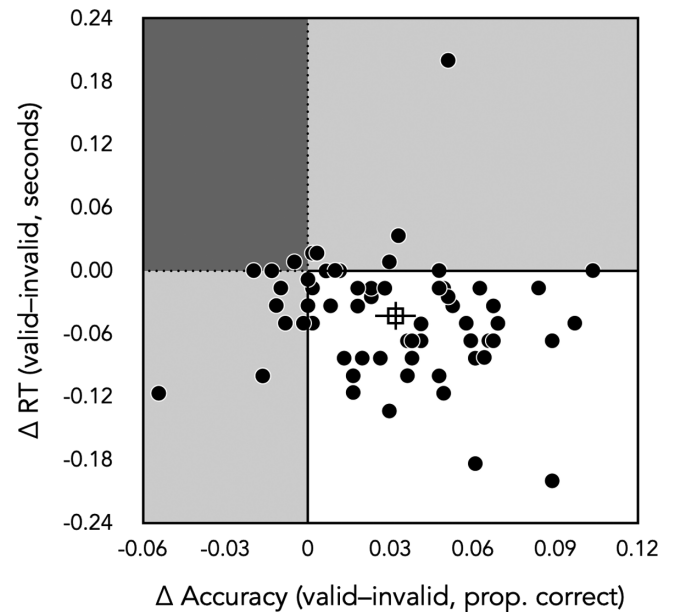


Figure 2. Overall exogenous cueing effects. Accuracy modulation is shown on the x-axis; RT modulation on the y-axis. Black dots are individual participants ($N = 66$), square icon indicates mean across participants with bootstrapped 95% confidence intervals. See Results: [Experiment 1](#) for additional details.

for each effect, drawing a new random sample of 66 observers on each iteration. The inner 95% of the bootstrapped distribution was extracted to produce the 95% confidence interval. An analogous procedure was used to create 95% confidence intervals for results presented in [Figures 3C](#) and [4B](#).

Variables

The three independent variables were cue type (reward-predictive or nonreward-predictive), cue validity (valid or invalid), and SOA (33, 67, 100, 133 ms). Each of the 16 unique trial types was presented 76 times. The dependent variables were accuracy (proportion correct), RT (median, correct trials only), and an inverse efficiency score (IES), which combines accuracy and RT into a single value (e.g., [Bucker & Theeuwes, 2017](#); [Grubb, White, Heeger, & Carrasco, 2015](#); [Kiss, Driver, & Eimer, 2009](#); [Townsend & Ashby, 1983](#)). RT is sometimes divided by proportion correct, thus penalizing speed as accuracy decreases (i.e., RT/proportion correct). One can also divide accuracy by RT, such that accuracy is penalized as RTs increase (i.e., proportion correct/RT). Although both approaches combine speed and accuracy data through a simple normalization procedure, the statistical results will differ between approaches. Because the

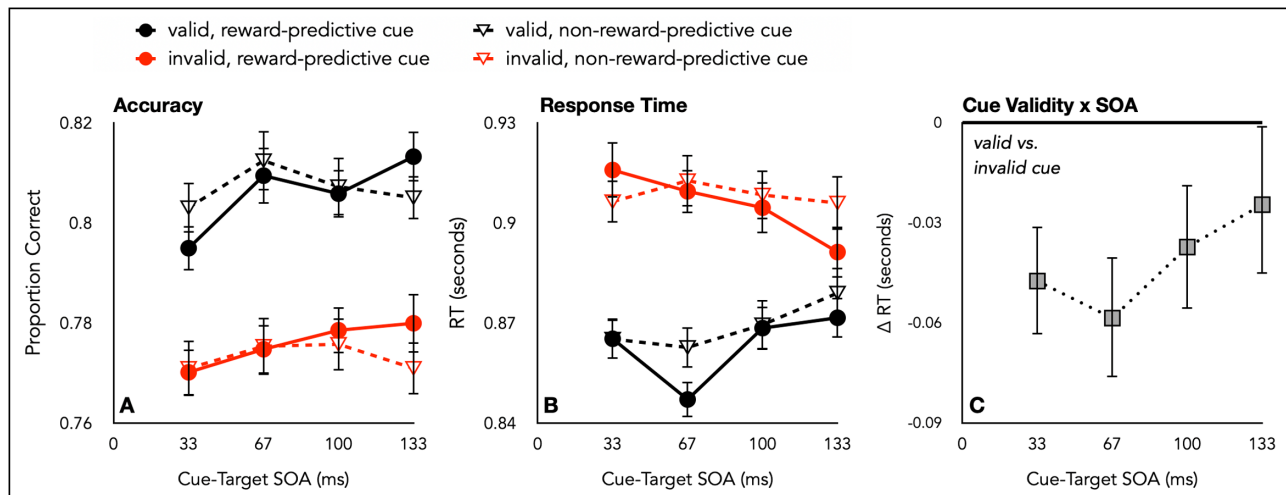


Figure 3. Comprehensive results. Accuracy (A) and RT (B) as a function of SOA for each cue condition. Error bars, within-participant error (Cousineau, 2005). N = 66. (C) Interaction follow-up. Error bars, bootstrapped 95% confidence intervals.

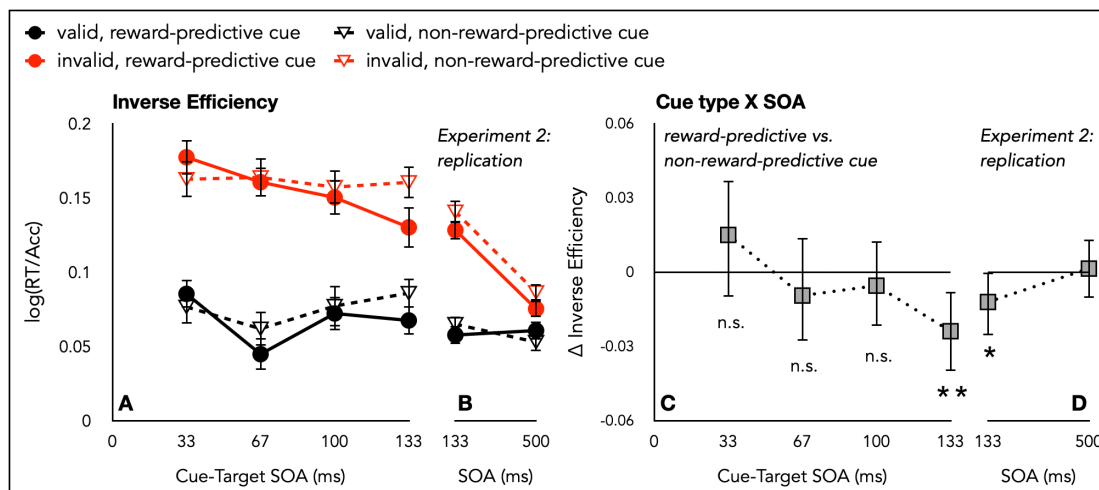


Figure 4. Inverse efficiency results. (A) Inverse efficiency as a function of SOA for each cue condition. Error bars, within-participant error. N = 66. (B) Experiment 2 replication and extension. Error bars, within-participant error. N = 82. (C) Interaction follow-up. Error bars, bootstrapped 95% confidence intervals. ** $p < 0.01$. (D) Experiment 2 replication and extension. Error bars, bootstrapped 95% confidence intervals. * $p < 0.05$.

former is the reciprocal of the latter, the two measures are mathematically related via a log transformation [$\ln(A/B) = -\ln(B/A)$]. Here we use the log-transformed measure [i.e., $\ln(\text{RT}/\text{proportion correct})$] so that it does not matter which metric is in the numerator and which is in the denominator; the values (except for a sign flip), and most importantly the statistics, are identical when IES is calculated as $\ln(\text{RT}/\text{proportion correct})$ or as $\ln(\text{proportion correct}/\text{RT})$.

Apparatus

The experiment was programmed in PsychoPy (Peirce, 2007) and run on a 3.0 GHz Dual-Core Intel Core i7 Mac Mini; stimuli were displayed on a 27-in. LED-Lit Dell Gaming Monitor (model: S2716DG). Participants were seated in a darkened experimental testing room and instructed to keep their chins on a chin rest 70 cm from the monitor.

	Accuracy			RT			Inverse efficiency		
	<i>F</i>	<i>df</i>	<i>p</i>	<i>F</i>	<i>df</i>	<i>p</i>	<i>F</i>	<i>df</i>	<i>p</i>
Cue type	0.08	(1, 65)	0.7826	2.15	(1, 65)	0.1477	2.27	(1, 65)	0.1363
Cue validity	67.54	(1, 65)	0.0000	43.36	(1, 65)	0.0000	87.07	(1, 65)	0.0000
SOA	2.00	(3, 195)	0.1155	0.39	(3, 195)	0.7624	1.26	(3, 195)	0.2877
Cue type × cue validity	0.60	(1, 65)	0.4427	0.15	(1, 65)	0.6968	0.02	(1, 65)	0.9014
Cue type × SOA	1.22	(3, 195)	0.3022	1.39	(3, 195)	0.2468	2.22	(3, 195)	0.0869
Cue validity × SOA	0.43	(3, 195)	0.7287	4.12	(3, 195)	0.0074	2.44	(3, 195)	0.0653
Cue type × cue validity × SOA	0.09	(3, 195)	0.9649	0.67	(3, 195)	0.5697	0.32	(3, 195)	0.8101

Table. *F* statistics, degrees of freedom (*df*) and *p* values from three-way ($2 \times 2 \times 4$) repeated-measures ANOVAs on accuracy, RT, and inverse efficiency data. Significant and marginal *p* values denoted in bold

Results: Experiment 1

As a preliminary sanity check, we first verified that our exogenous cues, pooled across SOA and cue type (i.e., reward-predictive and nonreward-predictive), were capable of generating a robust cueing effect. Replicating decades of previous research (Carrasco, 2011), we found that participants were both more accurate [paired *t*-test, $t(65) = 8.218$, $p < 0.0001$, mean within-participant change = 0.0318, bootstrapped 95% confidence interval (0.0243, 0.0393)] and faster [paired *t*-test, $t(65) = -6.283$, $p < 0.0001$, mean within-participant change = -43.10 ms, bootstrapped 95% confidence interval (-56.17, -29.68)] when the cue was valid compared with invalid, thus confirming the reflexive modulation of visual processing in our task (Figure 2).

To assess our two hypotheses, we completed a three-way ANOVA, with cue type (reward-predictive, nonreward-predictive), cue validity (valid, invalid), and SOA (33, 67, 100, 133 ms) as repeated-measures factors, separately for accuracy and RT data (Figures 3A, 3B). Consistent with the results discussed earlier, we found significant main effects of cue validity for both metrics ($ps < 0.0001$), but apart from an RT-based cue validity \times SOA interaction ($p = 0.0074$), which confirms well-established temporal dynamics of exogenous cueing effects (Figure 3C), there were no other significant main effects or interactions (Table). These independent accuracy- and RT-based analyses provide robust evidence for spatial cueing effects but suggest no reward-specific modulation in our task.

A close inspection of Figure 2, however, reveals two important issues for further consideration, both of which are related to the balancing of speed and accuracy at the individual level. First, although most of the participants show improvements in both speed and accuracy on valid compared with invalid trials (lower-right, white quadrant), a minority of participants do show a speed-accuracy trade-off, such that accuracy improves, whereas RTs slow (upper-right, gray quadrant) or accuracy suffers as RTs accelerate

(lower-left, gray quadrant). Second, some participants without a speed-accuracy trade-off show expected attention effects (i.e., improved discrimination accuracy or faster RTs), but only in one domain. For example, participants on the solid black line running along the horizontal show expected changes in accuracy, at no cost to RT; similarly, participants on the solid black line running along the vertical show expected changes in RT, at no cost to accuracy. Separate analyses on speed and accuracy data fail to take both of these issues into account. Combining accuracy and RT into a single metric can control for individual-level speed-accuracy trade-offs and may improve the sensitivity of the analytic approach.

To this aim, we computed an IES, which combines RT and accuracy into a single metric for each condition (Figure 4A) and recomputed our three-way ANOVA (Table). Importantly, we found a robust main effect of cue validity ($p < 0.0001$), alleviating concerns regarding individual speed-accuracy trade-offs. We also found marginal evidence for a cue validity \times SOA interaction ($p = 0.0653$), again revealing the expected temporal dynamics of the exogenous cue. Of interest to our second hypothesis, there was trend toward a cue type \times SOA interaction ($p = 0.0869$), indicating that the difference in task performance on reward-predictive versus nonreward-predictive trials was also dependent on SOA. Pooling data across cue validity conditions, we followed up this marginal interaction with paired *t*-tests at each SOA and found that inverse efficiency significantly decreased (i.e., participants became more efficient) on reward-predictive trials, but only at the final SOA of 133 ms [$t(65) = -2.88$, $p = 0.0054$, all other $ps > 0.1862$, Figure 4B]; importantly, the results of these null hypothesis tests are consistent with a bootstrapping procedure, which revealed that the 95% confidence intervals for the first three SOAs contain zero, whereas the final SOA does not.

These results suggest a spatially nonspecific impact of the rewarding cue. To evaluate this possibility further, we computed two-way ANOVAs, with cue type (reward-predictive, nonreward-predictive) and cue

validity (valid, invalid) as repeated-measures factors, separately at each SOA. Consistent with a lack of spatial specificity, none of the cue-type x cue-validity interactions were significant (all p s > 0.4627); consistent with the results described earlier, we observed a significant main effect for cue validity at each SOA (all p s < 0.001), and a significant main effect of cue type at the final SOA only (p = 0.0087).

In one final control analysis, we confirmed that our approach to counterbalancing the luminance of the reward-predictive cue was effective in controlling for low-level features of the cues themselves. The change in efficiency, for reward-predictive, relative to nonreward-predictive cues, at the final SOA of 133 ms was similar in size and in the same direction for both subsets of participants (black reward-predictive cue, mean = -0.0249 ; white reward-predictive cue, mean = -0.0226).

Interim discussion

Despite the fact that the difference in inverse efficiency at the final SOA, on trials with reward-predictive versus nonreward-predictive onsets, was statistically significant by traditional standards (p = 0.0054), the cue type x SOA interaction had a p value of 0.0869, which may give some readers cause to question the appropriateness of the subsequent follow-up tests. We are sensitive to this criticism but feel that our follow-up tests are warranted in this particular situation. With a p value less than 0.1, it would have been disingenuous to claim that there was no evidence for an impact of SOA. Intuition also suggests that, like the spatial attention effect, there should be some temporal dynamics associated with any impact of the reward-predictive cue, given its abrupt and short-lived appearance.

That said, to rule out the possibility that following up on a marginally significant interaction led to a statistically significant false positive, we ran a second experiment whose primary goal was to replicate our result in an independent group of observers. To fulfill a secondary aim of evaluating whether such a modulation in task efficiency would be observed at longer SOA durations, we also included an additional SOA condition in our replication study.

Methods: Experiment 2

Participants

Using the data from [Experiment 1](#), we conducted a nonparametric power analysis to determine our

sample size for [Experiment 2](#). From our 66 datasets, we randomly sampled (with replacement) two to 100 observers, in one-observer increments, and at each sample size we repeated our paired t -test of interest: inverse efficiency on trials with reward-predictive versus nonreward-predictive cues at the final SOA of 133 ms. We repeated this procedure 10,000 times to build a distribution at each sample size, each time drawing a new random sample of observers. We then calculated the proportion of 10,000 iterations in which the null hypothesis was rejected, separately for each sample size. This procedure indicated that a replication power of approximately 95% could be achieved with 82 participants and the use of one-tailed tests in line with the prediction from [Experiment 1](#).

Ninety-one new participants (45 female, average age 19.93 years) completed [Experiment 2](#). Nine participants failed to perform the task above chance (see Assessing chance accuracy section), leaving 82 participants for subsequent analyses.

Experimental changes

The methodological approach in [Experiment 2](#) differed from [Experiment 1](#) in two ways: (1) we included only two distinct interstimulus-interval durations (100 and 467 ms), yielding two SOA conditions (133 and 500 ms); and (2) we doubled the number of trials at each SOA condition. The first change was implemented (1) to replicate the key finding from the final SOA of [Experiment 1](#) in an independent group of naive observers; and (2) to extend our previous results by evaluating the impact of the peripheral cues at a time when the exogenous spatial attention effect will have greatly diminished (i.e., 500 ms after the onset of the cue), owing to its transient nature ([Carrasco, 2011](#)). The second change was made to match the total number of trials in [Experiment 1](#), thus equating the length of the session and the potential bonus earnings. For bootstrapped confidence intervals, we used the 82 participants from [Experiment 2](#), but all other methodological details matched exactly those reported earlier for [Experiment 1](#).

Results: Experiment 2

The primary goal of [Experiment 2](#) was to confirm the robustness of our [Experiment 1](#) finding that, regardless of their validity (i.e., location), reward-predictive onsets enhance task performance at an SOA of 133 ms. This result was replicated in an independent group of 82 participants: inverse efficiency significantly decreased (i.e., participants became more efficient) on trials in which a reward-predictive, relative to a

nonreward-predictive, cue was present [$t(81) = -1.90$, $p = 0.0308$, one-tailed in predicted direction, [Figure 4C](#)]. The result of this null hypothesis test is also consistent with a bootstrapping procedure, which revealed that the 95% confidence interval did not contain zero (95% CI = $[-0.0255, -0.003]$). At our additional SOA of 500 ms, however, the modulatory impact of the reward-predictive onset was no longer evident [$t(81) = 0.24$, $p = 0.8137$, [Figure 4C](#)], suggesting that the boost in performance observed at SOA durations of 133 ms in [Experiments 1](#) and [2](#) may be more reflexive than intentional.

Discussion

In our first experiment, we used abrupt-onset, peripheral cues to manipulate the covert allocation of exogenous spatial attention during an orientation discrimination task, employed associative learning to pair monetary reward with one of two luminance-defined cue types, and tested two related hypotheses concerning reward learning and attentional allocation. Regarding our first hypothesis, that reward-predictive, exogenous cues elicit the involuntary allocation of covert exogenous spatial attention more effectively than do nonreward-predictive cues, we found robust evidence that both cue types modulated task performance (i.e., task accuracy increased, whereas both RTs and IESs decreased, on valid relative to invalid trials). However, we found no evidence that the reward-predictive cue generated spatial attentional effects that were larger, or manifested earlier in time, than those generated by the nonreward-predictive cue. In short, there was no spatially specific impact of reward learning on the reflexive allocation of covert exogenous attention.

We did find evidence for our second hypothesis, that reward-predictive onsets, regardless of their validity, impact visual processes in a spatially nonspecific manner. Regardless of the cue's location, IESs for the orientation discrimination judgement decreased when the target was preceded by a cue that predicted reward, relative to when the target was preceded by a nonreward-predictive cue, thus confirming that task performance improved (i.e., became more efficient) when the luminance of the cue signaled that reward was on the line. Interestingly, an assessment of the temporal dynamics revealed that the impact of the reward-predictive cue grew as the trial progressed in time, reaching statistical significance only at our final SOA of 133 ms. In [Experiment 2](#), we replicated this result in an independent group of observers.

These results fit nicely within the theoretical framework proposed by [Awh et al., \(2012\)](#), wherein one's attentional history, in addition to goal-directed and stimulus-driven sources of attentional bias, is

an important component in determining the focus of attention. With sufficient exposure, the presence of the reward-predictive cue may serve to reflexively enhance attentional allocation toward the response cue (the small line at fixation that indicated the grating stimulus for which the orientation judgment should be made), which could, in turn, increase the efficiency of the response. Rapid, reflexive orienting of spatial attention would have been reinforced in the past (i.e., by the delivery of monetary reward following a correct orientation judgment), even when the reward-predictive cue was delivered near the distractor grating, creating a divergence of attentional histories between trials that contained a reward-predictive cue and those that did not, irrespective of cue validity. In addition to modulating spatial attention, these reward-predictive cues are expected to have generated a transient, dopamine-dependent reward prediction signal ([Daw & Tobler, 2013](#)). How such a reward prediction signal may modify attentional processes (e.g., by generating an increase in arousal that facilitates a reflexive orienting toward the response cue as described earlier) is not something our data can speak to, but understanding any potential interaction is an important next step and an exciting area for future research.

Although previous studies have investigated the interaction of spatial cueing and monetary incentives, we extend work in this area by associating reward directly with features of abrupt-onset cues themselves. [Engelmann and Pessoa \(2007\)](#) manipulated motivation by varying probabilistic rewards and punishments (performance-contingent monetary gains and losses) from block-to-block and, as task incentives increased, observed improvements in both validly and invalidly precued targets in a spatial localization task. [Bucker and Theeuwes \(2014\)](#) also manipulated task structure to probe the relationship between reward and spatial attention, using high and low reward blocks with spatially uninformative peripheral onsets; consistent with our results, they found significant, and comparably sized, cueing effects at their short SOA (170 ms), for both high and low reward blocks. Unlike our results, however, they did not find a spatially nonspecific effect of reward (i.e., irrespective of cue validity, task performance was not improved in high reward, relative to low reward, blocks), and this may be because of the blocking nature of the experimental design. Previous research has shown evidence of habituation to both abrupt onset cues ([Turatto, Bonetti, Pascucci, & Chelazzi, 2018](#)) and to rewarding stimuli ([Lloyd, Medina, Hawk, Fosco, & Richards, 2014](#)), so blocking rewards might lead to increased habituation, and thus, a loss in the cue's effectiveness.¹ Tying reward directly to a perceptual feature of the cue itself, rather than relying on block-by-block changes in task incentives, allowed us to probe the impact of implicitly learned stimulus-reward associations on a trial-to-trial

basis and reveal a more subtle relationship between performance-contingent reward and task performance.

Munneke et al., (2015) also used color-based features of an abrupt onset cue to signal the availability of monetary reward and found, in contrast to our results, increased capture effects for reward-predicting onsets. Directly comparing their study and ours, however, is difficult because their study combined exogenous cues with spatially informative, endogenous cues. Eighty percent of the time, participants knew in advance where the target would appear and could preallocate voluntary spatial attention to the target location. In our study, no spatially relevant information was provided until the response cue appeared. In the absence of a reward manipulation, endogenous spatial attention has been shown to modulate exogenous effects (Grubb et al., 2015), so it is possible that the discrepancy between their results and ours is because of the differences in the availability of spatially relevant information in advance of the target and distractor appearing. Future research that systematically characterizes the behavioral consequences of multiple sources of attentional biases (e.g., those arising from reward-learning, the physical salience of an abrupt-onset cue, and the top-down goals of the observer), when present in varying combinations, will fill a critical gap in this literature.

In our assessment, the most fruitful next step will be to focus exclusively on the spatially nonspecific impact of reward-predictive onsets on visual processing. Because the most sensitive behavioral metric combined both accuracy and RT (i.e., IES), another important direction for future research concerns the temporal dynamics of the processing of the visual target itself. Previous psychophysical studies using a speed-accuracy trade-off approach (Wickelgren, 1977), which conjointly measures discriminability and speed of information processing by using response tones to sample task accuracy at distinct points in time during visual processing, has shown repeatedly shown that spatial attention speeds the rate at which visual information is processed by the visual system, as well as increases asymptotic discriminability (e.g., Carrasco & McElree, 2001; Giordano, McElree, & Carrasco, 2009; Grubb et al., 2015). Future work using an speed-accuracy trade-off (SAT) procedure will help to pinpoint whether the results reported here are due to increases in the rate of visual processing, improvements in asymptotic discriminability, or both.

Conclusions

An abrupt onset elicits the rapid allocation of spatial attention, reflexively improving visual perception at the attended location, no matter what one may have learned about the visual features of the onset itself. However,

when experience has shown that specific features of the onset predict the availability of monetary reward, enhancements in visual perception become less spatially constrained, spreading rapidly to other task-relevant locations.

Keywords: covert attention, exogenous attention, reward

Acknowledgments

Open practices statement: The data on which these results are based are available on reasonable request from the corresponding author (MG). Study supported by Trinity College start-up funding to MG.

Commercial relationships: none.

Corresponding author: Michael A. Grubb.

Email: michael.grubb@trincoll.edu.

Address: Department of Psychology, Trinity College, Hartford, CT, USA.

Footnote

¹We thank an anonymous reviewer for raising this possibility.

References

- Anderson, B. A. (2018). Neurobiology of value-driven attention. *Current Opinion in Psychology*, 29, 27–33, doi:10.1016/j.copsyc.2018.11.004.
- Awh, E., Belopolsky, A. V., & Theeuwes, J. (2012). Top-down versus bottom-up attentional control: A failed theoretical dichotomy. *Trends in Cognitive Sciences*, 16(8), 437–443, doi:10.1016/j.tics.2012.06.010.
- Bucker, B., & Theeuwes, J. (2014). The effect of reward on orienting and reorienting in exogenous cuing. *Cognitive Affective & Behavioral Neuroscience*, 14(2), 635–646, doi:10.3758/s13415-014-0278-7.
- Bucker, B., & Theeuwes, J. (2017). Pavlovian reward learning underlies value driven attentional capture. *Attention, Perception & Psychophysics*, 79(2), 415–428, doi:10.3758/s13414-016-1241-1.
- Carrasco, M. (2011). Visual attention: The past 25 years. *Vision Research*, 51(13), 1484–1525, doi:10.1016/j.visres.2011.04.012.
- Carrasco, M., & McElree, B. (2001). Covert attention accelerates the rate of visual information processing. *Proceedings of the National Academy of Sciences of the United States of America*, 98(9), 5363–5367, doi:10.1073/pnas.081074098.

- Chelazzi, L., Perlato, A., Santandrea, E., & Della Libera, C. (2013). Rewards teach visual selective attention. *Vision Research*, *85*, 58–72, doi:10.1016/j.visres.2012.12.005.
- Daw, N., & Tobler, P. (2013). Value learning through reinforcement: The basics of dopamine and reinforcement learning. In P. W. Glimcher, & E. Fehr (Eds.), *Neuroeconomics: Decision making and the brain* (pp. 283–298). Cambridge, MA: Academic Press.
- Engelmann, J. B., & Pessoa, L. (2007). Motivation sharpens exogenous spatial attention. *Emotion*, *7*(3), 668–674, doi:10.1037/1528-3542.7.3.668.
- Failing, M., & Theeuwes, J. (2018). Selection history: How reward modulates selectivity of visual attention. *Psychonomic Bulletin & Review*, *25*(2), 514–538, doi:10.3758/s13423-017-1380-y.
- Giordano, A. M., McElree, B., & Carrasco, M. (2009). On the automaticity and flexibility of covert attention: a speed-accuracy trade-off analysis. *Journal of Vision*, *9*(3):30, 31–10, doi:10.1167/9.3.30.
- Grubb, M. A., White, A. L., Heeger, D. J., & Carrasco, M. (2015). Interactions between voluntary and involuntary attention modulate the quality and temporal dynamics of visual processing. *Psychonomic Bulletin & Review*, *22*(2), 437–444, doi:10.3758/s13423-014-0698-y.
- Kiss, M., Driver, J., & Eimer, M. (2009). Reward priority of visual target singletons modulates event-related potential signatures of attentional selection. *Psychological Science*, *20*(2), 245–251, doi:10.1111/j.1467-9280.2009.02281.x.
- Le Pelley, M. E., Mitchell, C. J., Beesley, T., George, D. N., & Wills, A. J. (2016). Attention and associative learning in humans: An integrative review. *Psychological Bulletin*, *142*(10), 1111–1140, doi:10.1037/bul0000064.
- Lloyd, D. R., Medina, D. J., Hawk, L. W., Fosco, W. D., & Richards, J. B. (2014). Habituation of reinforcer effectiveness. *Frontiers in Integrative Neuroscience*, *7*, 107, doi:10.3389/fnint.2013.00107.
- Muller, H. J., & Rabbitt, P. M. A. (1989). Reflexive and voluntary orienting of visual-attention—Time course of activation and resistance to interruption. *Journal of Experimental Psychology—Human Perception and Performance*, *15*(2), 315–330, doi:10.1037/0096-1523.15.2.315.
- Munneke, J., Hoppenbrouwers, S. S., & Theeuwes, J. (2015). Reward can modulate attentional capture, independent of top-down set. *Attention Perception & Psychophysics*, *77*(8), 2540–2548, doi:10.3758/s13414-015-0958-6.
- Nakayama, K., & Mackeben, M. (1989). Sustained and transient components of focal visual-attention. *Vision Research*, *29*(11), 1631–1647, doi:10.1016/0042-6989(89)90144-2.
- Peirce, J. W. (2007). PsychoPy—Psychophysics software in Python. *Journal of Neuroscience Methods*, *162*(1-2), 8–13, doi:10.1016/j.jneumeth.2006.11.017.
- Townsend, J. T., & Ashby, F. G. (1983). *The stochastic modeling of elementary psychological processes*. Cambridge: Cambridge University Press.
- Turatto, M., Bonetti, F., Pascucci, D., & Chelazzi, L. (2018). Desensitizing the attention system to distraction while idling: A new latent learning phenomenon in the visual attention domain. *Journal of Experimental Psychology. General*, *147*(12), 1827–1850, doi:10.1037/xge0000503.
- Wickelgren, W. A. (1977). Speed-accuracy tradeoff and information-processing dynamics. *Acta Psychologica*, *41*(1), 67–85, doi:10.1016/0001-6918(77)90012-9.