Experience-driven recalibration of learning from surprising events

Leah Bakst* and Joseph T. McGuire

Department of Psychological & Brain Sciences, Boston University, Boston, MA

Center for Systems Neuroscience, Boston University, Boston, MA

Mailing address: 677 Beacon Street, Boston, MA 02215

Email addresses: lbakst@bu.ed, jtmcg@bu.edu

Author note: Some of the data in this article were presented at the Society for Neuroeconomics annual meeting in 2019. Experiment 2 was preregistered at <https://aspredicted.org/cd54j.pdf>. Task code and data are available at <https://osf.io/5pmhg/> and analysis code is available upon request.

ABSTRACT

Different contexts favor different patterns of adaptive learning. A surprising event that in one context would accelerate belief updating might, in another context, be downweighted as a meaningless outlier. Here, across two experiments, we examined whether participants performing a perceptual judgment task under spatial uncertainty ($n=29$, $n=63$) would spontaneously adapt their patterns of predictive gaze according to the informativeness or uninformativeness of surprising events in their current environment. Uninstructed predictive eye movements exhibited a form of *metalearning* in which surprise came to modulate event-by-event learning rates in opposite directions across contexts. Participants later appropriately readjusted their patterns of adaptive learning when the statistics of the environment underwent an unsignaled reversal. Although significant adjustments occurred in both directions, performance was consistently superior in contexts in which surprising events reflected meaningful change, potentially reflecting a bias toward interpreting surprise as informative. Overall, our results demonstrate remarkable flexibility in contextually adaptive *metalearning*.

INTRODUCTION

Different environments involve different sources of uncertainty, with profound implications for how decision makers should respond to surprising events. An event that violates one's expectations might be interpreted either as a meaningless outlier or a sign of fundamental change. For example, suppose you supervise a reliable employee who one day misses work without notice. Whether you dismiss the event as a one-off or modify your predictions about the employee's future reliability will depend on your background knowledge and assumptions about the sources of uncertainty in this context.

Previous research has approached the problem of inference and prediction in dynamic environments by making a distinction between *expected uncertainty* and *unexpected uncertainty* (Payzan-LeNestour et al., 2013; Payzan-LeNestour & Bossaerts, 2011; Soltani & Izquierdo, 2019; A. J. Yu & Dayan, 2005). *Expected* uncertainty relates to noisy event-to-event fluctuations that are compatible with one's current understanding of how events are generated. A less-than-fully-predictable event, like rolling a six-sided die and getting a six, need not drive updates to one's understanding of the outcome-generating mechanism. *Unexpected* uncertainty relates to events that do not conform to the known unreliability of predictive cues (such as rolling a six-sided die and getting a seven) and signal a need to update one's beliefs about the underlying context. Previous empirical work has shown that people adaptively modulate the rate at which they learn from new observations by balancing context-appropriate estimates of expected uncertainty, the likelihood of true change in the outcome-generating process, and other factors (Bakst & McGuire, 2021; Behrens et al., 2007; Cheadle et al., 2014; Diederer et al., 2016; Farashahi et al., 2019; Lee et al., 2020; Massi et al., 2018; McGuire et al., 2014; Nassar et al., 2010, 2019; Ossmy et al., 2013). These different factors have been linked to dissociable patterns

of activity in arousal-related brain systems, frontoparietal cortex, and ventral striatum (Jepma et al., 2016, 2018; Kao et al., 2020; McGuire et al., 2014; Nassar et al., 2012; O'Reilly et al., 2013). Reduced flexibility of adaptive learning has been associated with dimensions of psychopathology including trait anxiety (Browning et al., 2015; Kraus et al., 2021; Pulcu & Browning, 2019).

Previous research on adaptive learning has largely focused on scenarios in which events that signify fundamental change tend to be more extreme and surprising than events that reflect expected uncertainty. For example, in spatial prediction paradigms with “change point” structure, the generative statistics that govern target positions occasionally undergo discrete shifts. Targets subsequent to a change point tend to be associated with high *prediction error* between the anticipated and actual position. In this type of environment, an approximately Bayesian theoretical account (Nassar et al., 2016) holds that a larger prediction error should lead to a higher inferred probability that a change point has occurred, resulting in a higher trial-specific learning rate (that is, a large belief update expressed as a proportion of the prediction error). A related theoretical proposal holds that learning is gated by the associability of a stimulus, which scales with surprise (Li et al., 2011; Pearce & Hall, 1980).

However, larger-magnitude prediction errors need not always lead to higher learning rates (O'Reilly et al., 2013; L. Q. Yu et al., 2021). In some situations, good decisions require maintaining stable beliefs in the face of salient but non-predictive events, such as when an investor avoids overreacting to a transient market downturn. People can adapt to environments in which the ideal learning rate is equal or lower for surprising events relative to more moderate events (Cheadle et al., 2014; d'Acremont & Bossaerts, 2016; Lee et al., 2020; Nassar et al., 2019; O'Reilly et al., 2013; Summerfield & Tsetsos, 2015). Downregulating learning in response

to non-predictive “oddball” events has been suggested to require vigilance and control (d’Acremont & Bossaerts, 2016), although other evidence suggests that actively updating beliefs requires additional cognitive resources compared with holding them stable (O’Reilly et al., 2013).

Given that surprising events can either upregulate or downregulate learning in different contexts, an outstanding question is to what degree adaptive learning undergoes spontaneous, experience-driven recalibration to track the current context’s higher-order statistics. Calibration of adaptive learning, a form of *metalearning* (Griffiths et al., 2019; Wang, 2021), has previously been studied in human behavioral paradigms in which experimental participants were explicitly tasked with optimizing and reporting their predictions about upcoming events and received explicit description, training, and/or event-specific cues to ensure they clearly understood the relevant sources of uncertainty (d’Acremont & Bossaerts, 2016; Nassar et al., 2019; O’Reilly et al., 2013). In this previous work, participants successfully adapted their behavior to their environment, showing dynamic learning rate modulation appropriate for their context. These studies provide persuasive evidence for flexible learning-rate modulation and raise the question of how successfully people would spontaneously infer a context-appropriate pattern of learning rate modulation solely from experience in a new environment.

We recently developed an eye-tracking paradigm for measuring spontaneous belief updating amid spatial uncertainty via predictive gaze (Bakst & McGuire, 2021). The paradigm is premised on the idea that prediction is a natural aspect of moment-to-moment cognition (Clark, 2013; Hayhoe et al., 2012; Henderson, 2017; Summerfield & Lange, 2014). While making perceptual judgments about briefly presented targets, participants’ eye movements tended to anticipate target onsets, which were temporally predictable but spatially uncertain. Despite not having been

informed about the behaviors of interest or the optimal strategy, participants exhibited well-calibrated predictive inference and adapted to a manipulation of expected uncertainty in a change-point environment (Bakst & McGuire, 2021).

Here, across two experiments, we extended this eye-tracking paradigm to test the hypothesis that people would spontaneously adapt to the informativeness of surprising events. In Experiment 1, we used a between-participant design to contrast environments in which surprise reflected either meaningful change or uninformative outliers. In Experiment 2, using a within-participant design, we tested people's ability to adjust their behavioral strategy when the environment's structure reversed. This allowed us to investigate the possibility of an asymmetry in the difficulty of the two environments, such that participants would adapt more easily to environments in which surprising outcomes warranted updating or to environments in which surprising outcomes were best ignored.

RESULTS

General Task Performance

We used an implicit spatial prediction paradigm (Bakst & McGuire, 2021) across two experiments to evaluate how successfully people could learn through experience to interpret surprising events either as meaningful changes or random outliers (Figure 1A). Participants were given a nominal task of reporting whether briefly presented numerical digits were even or odd. The task implicitly led participants to make anticipatory saccades to the predicted locations of upcoming digits on the screen and to update their spatial predictive beliefs continually in response to new observations.

The task had two conditions: In the *Change Point* (CP) condition, peripheral digit locations were drawn from a one-dimensional Gaussian distribution on the perimeter of a large circle ($\sigma =$

11.25° of angular distance; see Figure 1A) with the mean of the Gaussian distribution occasionally resampled from a uniform distribution spanning the entire circumference of the circle. Peripheral trials alternated with central trials (see Figure 1A), which served to recenter the participant's gaze prior to each prediction. Small changes in digit location from one peripheral trial to the next therefore tended to represent noise around a stable mean, whereas large changes in digit location tended to represent meaningful shifts in the underlying position of the Gaussian (Figure 1B). To maximize accuracy, participants in the CP condition ideally should converge on an increasingly precise estimate of the current generative mean and downregulate learning from small errors to avoid chasing noise, but should transiently raise their learning rate in response to large errors reflective of change points (Nassar et al., 2010).

In the *Random Walk* (RW) condition, the contingencies were reversed: the Gaussian-sampled digit location on one trial became the generative mean for the next trial. This random walk was punctuated by occasional outliers drawn from a uniform distribution spanning the circumference of the circle. Small changes in digit location therefore tended to indicate meaningful shifts in the underlying Gaussian distribution, whereas large changes tended to represent meaningless outliers. To maximize performance, participants in the RW condition ideally should fully update their beliefs in response to small errors but should downregulate learning from extreme events.

Participants performed the nominal task with high accuracy for both central and peripheral trial types in both the CP and RW conditions. Mean accuracy for odd/even judgments in Experiment 1 was 82% (SD=5.6%) for peripheral trials and 98% (SD=1.4%) for central trials (Figure 1D), with no significant differences between conditions (Wilcoxon rank sum, both $p \geq 0.616$). Participants were also successful at the uninstructed task of predicting digit locations. Because task events were temporally predictable, we used gaze position at the time of digit

appearance as an indication of the participant's prediction (Bakst & McGuire, 2021). Data from a representative participant are shown in Figure 1C. Gaze was directed near the appropriate eccentricity for both central and peripheral trials (Figure 1C, top panel).

Figure 1. Task summary. (A) A digit appeared between flanking Xs for 180 ms before being masked by a third X. The participant then had unlimited time to indicate via a keypress whether the digit was even or odd. Accuracy feedback was displayed at the same location for 500 ms, followed by a 750-ms inter-trial interval (ITI). The total reward earned in each run was represented by the width of a bar at the bottom of the screen. Each run began with a central trial, and alternated between central and peripheral trials. (B) Task conditions. In the Change Point (CP) condition, small changes in digit location from one peripheral trial to the next tended to represent noise around a stable mean, whereas large changes in digit location tended to represent meaningful shifts in the underlying position of the Gaussian. In the Random Walk (RW) condition, small changes in digit location tended to indicate meaningful shifts in the underlying Gaussian distribution, whereas large changes tended to represent meaningless outliers. (C) Data

from an example participant. Top: Eccentricity. Dashed lines show the true eccentricity of central and peripheral digits in degrees of visual angle (deg). Gaze-based predictions were defined as the gaze position at the time of digit onset, shown in triangles (light grey: peripheral trials, dark grey: central trials). Middle: CP condition, peripheral trials only. Digit location (angular position on the peripheral circle) shown in black points, gaze-based predictions shown in blue triangles. Bottom: RW condition. Same conventions as above, with dark red triangles indicating gaze-based predictions. (D) Odd/even task accuracy in Experiment 1 by condition and trial type. Individual participant performance shown in open circles, group means and SEM shown atop in black. (E) Prediction error (PE; mean absolute value across trials), conventions as in D. Left axis shows prediction error in degrees around the circle (for peripheral trials), right axis shows corresponding degrees of visual angle (for central trials).

Examining peripheral trials only, the angle of the prediction was near to the subsequent digit location for both the CP and RW conditions (Figure 1C, middle and bottom panels, respectively). The mean prediction error was 28.2° of angular distance ($SD=4.7^\circ$) for peripheral trials, and 0.84° of visual angle ($SD=0.53^\circ$) for central trials (Figure 1E), with no significant differences between conditions (Wilcoxon rank sum, both $p \geq 0.230$). As expected, both odd/even task accuracy and predictive gaze accuracy decreased transiently around extreme events (Supplemental Figure 1).

Adaptive learning: Experiment 1

Effects of prediction error magnitude

We tested whether participants modulated their learning rate across trials as a function of prediction error magnitude. Per-trial learning rate was empirically estimated as the gaze-based prediction update (from trial t to $t+1$) expressed as a proportion of the gaze-based prediction error (on trial t). We hypothesized that in the CP condition, participants would use higher learning rates for larger prediction errors, which tended to reflect meaningful change points. In the RW condition, in contrast, larger prediction errors tended to reflect noisy outliers and should be associated with lower learning rates.

To evaluate whether these patterns were present in the data, we first calculated the median observed trialwise learning rate for each participant for different bins of prediction error magnitude (bin edges were 0°, 5°, 10°, 15°, 20°, 25°, 35°, 45°, 75°, 105°, 135°, and 180°). Group mean and SEM (of within-subject medians) are shown in Figure 2A.

Figure 2. Learning rate as a function of absolute prediction error. (A) Learning rate as a function of prediction error magnitude (PE) in degrees of angular position on the peripheral circle (deg). Within-subject median learning rate was calculated for PE bins (bin edges = 0°, 5°, 10°, 15°, 20°, 25°, 35°, 45°, 75°, 105°, 135°, and 180°) for each participant. Group mean and SEM are shown for each group: RW (red), CP (blue). Left panel is using all task data, middle is from the Early epoch (first four minutes of task experience), right is Late epoch (last four minutes). (B) Linear models with a probit link were fit to the trial-by-trial data for each participant separately. Individual fits are shown in thin dashed lines; group median slope and intercept are shown in the thicker line. Left is RW (red), right is CP (blue).

Because initial task instructions were identical in the CP and RW conditions, any systematic differences in behavior presumably emerged through metalearning over time, as participants used direct experience with the task's statistics to recalibrate the rate at which they learned from individual events. To examine calibration of learning over time, we repeated the procedure for Early and Late task epochs (Figure 2A, middle and right panels), defined as the first and last four minutes of the task.

To facilitate a quantitative test of the effects depicted descriptively in Figure 2A, we fit probit functions to each participant's per-trial learning rates as a function of prediction error magnitude (Figure 2B, left). Individual participant fits are shown in dotted lines and group medians are shown in thicker lines. Slopes differed significantly between the CP and RW groups over the full task ($p = 0.014$, Wilcoxon rank sum test), with the CP group showing a median slope of 0.009 (interquartile range [IQR] = 0.002, 0.037) and RW a median of -0.001 (IQR = -0.006, 0.002).

Early behavior showed minimal differences between the two groups (Figure 2B, middle), with no significant difference in slopes (Wilcoxon rank-sum $p = 0.348$; CP median [IQR] = -0.003 [-0.010, 0.006], RW = 0.004 [-0.007, 0.043]). However, slopes in the Late epoch differed significantly between conditions (Wilcoxon rank-sum $p < 0.001$; CP = 0.023 [0.007, 0.050]; RW = -0.008 [-0.011, 0.001]; Figure 2D right). The change in slope between the Early and Late epochs (Late minus Early) likewise differed between conditions (Wilcoxon rank sum $p = 0.011$; CP median [IQR] = 0.020 [0.002, 0.059], RW = -0.006 [-0.051, 0.007]). However, the magnitude of the change in slope from Early to Late did not differ between groups (Wilcoxon rank sum $p = 0.326$), implying there was no evidence of faster metalearning in one context than the other.

Effect of extreme events

We focused our next analyses on non-Gaussian extreme events, which represented change points in the CP condition and outliers in the RW condition. Analyses were restricted to digits that were sampled from the uniform distribution and appeared more than two standard deviations from the previous Gaussian generative mean ($> 22.5^\circ$ of angular distance) to ensure we focused only on events that were distinguishable from Gaussian samples. In addition to being of theoretical interest, extreme events allowed for more accurate measurement of empirical learning rates compared with less-extreme events. Because per-trial learning rate estimates were

calculated as a ratio with prediction error in the denominator, the estimates were less susceptible to measurement error on large-prediction-error trials.

To evaluate performance over time, we looked at extreme-event learning rates over the course of the task, with time expressed in terms of the cumulative number of extreme events encountered (Figure 3A). The CP and RW groups exhibited different trajectories of learning rate as a function of experience. The overall difference in learning rate between the two conditions can be summarized in terms of the median extreme-event learning rate per participant (Figure 3B). The median learning rate was significantly greater in the CP condition than in the RW condition (Wilcoxon rank sum $p < 0.001$; CP median [IQR] = 0.99 [0.97, 1.01]; RW = 0.18 [0.04, 0.74]). Extreme-event learning rates also displayed a bimodal within-subject distribution in both groups, with a peak centered at the optimal behavior for each condition (Supplemental Figure 2).

Next, we evaluated whether the two groups differed in their final learning rate (median learning rate for extreme events during the Late epoch; Figure 3C). In the CP condition, the median final learning rate was 1.00 (IQR = 0.98, 1.02), which was not significantly different from the optimal learning rate of 1 (Wilcoxon signed rank $p = 0.952$). In comparison, in the RW condition, the median final learning rate was 0.08 [0.01, 0.65], which was significantly greater than the optimal learning rate of 0 (Wilcoxon signed rank $p = 0.002$,) and significantly lower than the final learning rate for the CP group (Wilcoxon rank sum $p < 0.00$).

Figure 3. Rates of learning from extreme events differed by condition and evolved with task experience. (A) Per-trial learning rate (LR) over time. Here, time is expressed in terms of the cumulative number of extreme events experienced. Group median and SEM are shown for each condition. Traces were smoothed with a 5-trial window. Data are plotted until the latest point at which at least 50% of participants per group contributed data. (B) Overall median extreme-event learning rate for each condition. Circles show individual participants and triangles show the group mean and SEM. (C) Final learning rate was calculated as the median extreme-event learning rate per participant during the Late epoch. Conventions as in B. (D) Change in median extreme-event learning rate (Δ LR) between Early and Late epochs of the task, conventions as in B-C.

Finally, we quantified the amount of metalearning as the change in learning rate for extreme events (Δ LR) over the course of the task (Figure 3D). We estimated Δ LR by calculating each participant's median learning rate for extreme events in Early and Late epochs and taking the difference (Late minus Early). For each group individually, the difference was not significantly different from zero (RW: Wilcoxon signed rank $p = 0.151$; median [IQR] = -0.28 [-0.81, 0.11]; CP: $p = 0.091$; median [IQR] = 0.05 [-0.03, 0.75]), likely because of consistently well-calibrated behavior over time in the CP condition and high inter-individual variability in the amount of

metalearning in the RW condition. The signed value of ΔLR differed significantly between the two groups (Wilcoxon rank sum $p = 0.025$) and its magnitude did not (Wilcoxon rank sum $p = 0.898$).

Adaptive learning: Experiment 2

Replication

To further investigate experience-driven metalearning, we conducted a preregistered second experiment in which the conditions reversed midway through each participant's experimental session. Each participant experienced three 8-minute runs of one condition (CP or RW) followed by three runs of the other, with the order counterbalanced across participants. As before, participants were not informed about the structure of the task or the variables of interest.

Pre-reversal data from the first three runs provided an opportunity to replicate the analyses reported above for Experiment 1. In examining learning rate as a function of prediction error magnitude, we obtained similar results (see Figure 4A for per-participant fits analogous to Figure 2B and see Supplemental Figure 3 for binned mean-of-median learning rates analogous to Figure 2A). Probit slope coefficients summarizing the overall effect of prediction error magnitude on learning rate differed significantly between the two groups (Wilcoxon rank sum $p < 0.001$; Figure 4A, left). The CP group showed a median slope of 0.009 (IQR = 0.003, 0.029; $p < 0.001$, Wilcoxon signed rank test against zero) and the RW group showed a median of -0.004 (IQR = -0.008, -0.0002; $p = 0.010$). The two groups did not differ during the Early epoch (Wilcoxon rank sum $p = 0.660$; CP median [IQR] = 0.005 [-0.004, 0.022]; RW median [IQR] = 0.012 [-0.009, 0.050]; Figure 4A middle) but differences were found in the Late epoch (Wilcoxon rank sum $p < 0.001$; CP median [IQR] = 0.024 [0.004, 0.050]; RW median [IQR] = -0.006 [-0.014, 0.001]; Figure 4A right). The change in slope between the Early and Late epochs differed between

groups (Wilcoxon rank sum $p = 0.002$) but the magnitude of the change did not (Wilcoxon rank sum $p = 0.093$).

A Pre-reversal

B Post-reversal

Figure 4. Learning rate as a function of absolute prediction error in Experiment 2. (A) Pre-reversal data: Linear models with a probit link were fit separately to each participant's trial-by-trial learning rate estimates. Overall fits are shown at left, Early epoch at middle, and Late epoch at right. Conventions as in Figure 2B. (B) Equivalent results for the post-reversal phase of Experiment 2. Across both panels, dashed lines represent the participant group that experienced the CP condition in the pre-reversal phase and the RW condition in the post-reversal phase; solid lines represent the group that experienced the opposite order.

Patterns of learning from extreme events also matched those seen in Experiment 1. Both groups demonstrated marked metalearning (Figure 5A), with a significant difference between groups in overall median learning rate (Wilcoxon rank sum $p < 0.001$; Figure 5B). A preregistered replication analysis showed that final learning rates in the pre-reversal phase differed significantly between groups (Wilcoxon rank sum $p < 0.001$; CP median [IQR] = 0.98 [0.93, 1.04]; RW median [IQR] = 0.07 [0.00, 0.41]; see Figure 5C). The final learning rate for the CP group was not significantly different from the optimal rate of 1 (Wilcoxon signed rank p

= 0.121), whereas the final learning rate in the RW group was significantly greater than the optimal rate of 0 (Wilcoxon signed rank $p < 0.001$). Signed ΔLR significantly differed between the two conditions (Wilcoxon rank sum $p < 0.001$; Figure 5D) but the magnitude of ΔLR did not (Wilcoxon rank sum $p = 0.092$). We again observed a bimodal pattern of individual trial learning rates in both conditions similar to that seen in Experiment 1 (Supplemental Figure 2).

Post-reversal behavior

When the CP and RW conditions reversed halfway through the experimental session, participants were given no explicit indication that the environment had changed. We investigated whether they would revise their understanding of the task's statistical structure through experience alone.

In examining learning rate as a function of prediction error magnitude, slopes in the post-reversal block as a whole did not significantly differ between groups (Wilcoxon rank sum $p = 0.532$), with both groups showing positive median slopes (CP median [IQR] = 0.003 [0.001, 0.023]; RW = 0.003 [-0.001, 0.007]; Figure 4B, left). The two groups also did not differ during the Early epoch (Wilcoxon rank sum $p = 0.132$; CP = 0.003 [-0.004, 0.014]; RW = 0.010 [-0.002, 0.049]; Figure 4B, middle) and although behavior trended in the expected direction, there was still no significant difference between groups in the Late epoch (Wilcoxon rank sum $p = 0.204$; CP = 0.011 [-0.002, 0.046]; RW = -0.001 [-0.005, 0.040]; Figure 4B, right), driven primarily by persistent suboptimal behavior in the group that experienced the RW condition after the reversal point. Post-reversal Late-epoch slopes did not significantly differ from zero in the RW condition (Wilcoxon signed rank $p = 0.421$), but were significantly positive in the CP condition (Wilcoxon signed rank $p < 0.001$).

In assessing learning from extreme events, participants who switched from the RW condition to the CP condition (RW $\rightarrow$ CP) showed a substantial increase in their learning rate shortly after the reversal point, indicative of rapid metalearning (Figure 5A). In comparison, participants who experienced CP and then RW (CP $\rightarrow$ RW) showed more gradual metalearning post-reversal. They appeared to maintain their previous beliefs about the environment for an extended period before gradually adapting their behavior.

Figure 5. Effects of the mid-session reversal of the task's statistical structure. (A) Learning rate (LR) for extreme events over time, pre- and post-reversal. RW (red), CP (blue). Time is expressed in terms of the cumulative number of extreme events experienced. Group median and SEM is shown for each condition. Data are plotted until the latest point at which at least 50% of participants per group contributed data in each half of the session. Traces were smoothed with a 5-trial window. (B) Overall median learning rate was calculated for each participant both pre- and post-reversal. Circles show individual participants and triangles show the group mean and SEM. (C) Final learning rate, conventions as in B. (D) Change in learning rate (Δ LR) between

Early and Late epochs, shown separately for the pre- and post-reversal phases. Conventions as in B.

To assess whether the order of the conditions affected behavior, we compared the overall median extreme-event learning rate for each condition between groups (comparing one group's pre-reversal behavior to the other group's post-reversal behavior; Figure 5B). Learning rates in the CP condition did not differ between groups (Wilcoxon rank sum $p = 0.514$), while those in the RW condition did (Wilcoxon rank sum $p < 0.001$; pre-reversal median [IQR] = 0.13 [0.04, 0.45], post-reversal = 0.94 [0.60, 0.98]), with pre-reversal learning rates closer to the optimal value of zero.

Analyses of final learning rates (extreme-event learning rates in the final 4 minutes of each phase; Figure 5C) provided further evidence that participants' ability to calibrate to the RW environment was impacted by previous experience in the CP environment. In a pre-registered analysis, there was a significant difference in final learning rates between participants who experienced the RW condition in the pre- versus post-reversal phases. The pre-reversal group showed better-calibrated behavior (Wilcoxon rank sum $p = 0.009$; pre-reversal median [IQR] = 0.07 [0.00, 0.41]; post-reversal = 0.41 [0.06, 0.99]). In contrast, the CP condition showed no significant difference between pre- and post-reversal phases (Wilcoxon rank sum $p = 0.685$; pre-reversal median [IQR] = 0.98 [0.93, 1.04]; post-reversal = 0.99 [0.91, 1.02]). In another pre-registered analysis, the change in learning rate (ΔLR , the difference between Early and Late epochs of each phase) did not significantly differ between the pre-reversal and post-reversal phases within either condition (Wilcoxon rank sum, both $p \geq 0.115$; Figure 5D), though both were significantly different than zero in the post-reversal phase (Wilcoxon signed rank, both $p \leq 0.050$). Additionally, the raw ΔLR differed between groups in the post-reversal phase (Wilcoxon rank sum $p < 0.001$), while the magnitude did not (Wilcoxon rank sum $p = 0.905$).

DISCUSSION

Across two experiments, we investigated the extent to which participants spontaneously used experience with the statistics of their environment to regulate the influence of surprise on learning. Learning was assessed via predictive eye movements while participants performed a perceptual judgment task under spatial uncertainty. Gaze-based predictions exhibited successful metalearning, tending to approach context-appropriate patterns of learning rate modulation over time. At the same time, learning rates for extreme events were generally better calibrated in a context in which large prediction errors were associated with meaningful change (CP) compared to one in which they were not (RW). We found asymmetric order effects: participants who experienced the RW condition second showed worse performance compared to those who saw RW first, whereas order had no effect on performance in the CP condition.

Learning initial task structure

Our findings agree with previous results from explicit prediction tasks, which demonstrated better performance in conditions in which surprise was meaningful than in conditions in which surprising events needed to be ignored (d'Acremont & Bossaerts, 2016; Nassar et al., 2019). However, the amount of metalearning (measured in terms of within-session change) was generally similar between conditions. This could suggest that performance differences are driven by initial assumptions about the environment. Participants may have held an initial bias or default towards assuming large predictions errors indicated meaningful change, consistent with the proposal that suppressing the influence of salient outliers is cognitively taxing (d'Acremont & Bossaerts, 2016). A related possibility is that maintaining stable beliefs across an intervening outlier entails elevated working memory demands. Such a default could be adaptive insofar as, in

real-world environments, there might be more dire consequences for erroneously ignoring events that differ from expectations than for erroneously overweighting them.

Questions remain as to whether other task contexts or framings could evoke different patterns of initial assumptions and inductive biases. For instance, O'Reilly and colleagues (2013) found additional behavioral costs associated with belief updating compared with merely reacting to isolated surprising events. Future work should examine the flexibility of default strategies and the cues relevant to assessing the relative costs of different patterns of learning. Future work should also examine how metalearning might be altered by task manipulations of arousal, motivation, or incentives (Jepma et al., 2018; Nassar et al., 2012; Urai et al., 2017).

We also noted substantial differences in behavioral policies across individuals, especially early on in the task. Behavior tended to converge over time, showing decreasing inter-individual variability in later epochs of the session. Why participants have such a variety of initial strategies is an open question. Whether it relates to differing interpretations of the experimental context, broader influence of their recent state and/or experiences, or trait-level factors (Browning et al., 2015; Kraus et al., 2021) has yet to be identified. A related goal for future work is to identify computational parameters that enable some individuals to adjust their behavioral policy rapidly and successfully while others maintain suboptimal strategies for extended periods. Previous modeling work has demonstrated the varying influence of factors such as surprise and environmental volatility on participant behavior (Behrens et al., 2007; Lee et al., 2020; McGuire et al., 2014; Nassar et al., 2010, 2019) and further efforts to model metalearning could help understand the sources of the behavioral patterns we identified (d'Acremont & Bossaerts, 2016; Nassar et al., 2019).

Learning after environmental change

In our second experiment, the conditions underwent an unsignaled reversal halfway through, and participants revised their behavioral strategy in line with the new statistical structure they encountered. The trajectory of metalearning appeared to exhibit condition-specific order effects. Participants who experienced the RW condition after having previously experienced the CP condition displayed CP-like strategies for longer than those who saw RW first. In comparison, behavior in the CP condition did not appear to depend on condition order. If participants tended to hold initial biases towards a CP-like strategy, then experiencing the CP condition could have served to reinforce those biases, making them yet more difficult to overcome when the participant was suddenly thrust into the RW condition.

A possibility that merits further investigation is that the asymmetry in metalearning between the two conditions might be rooted in an asymmetric performance cost of miscalibration. A CP-optimized agent could be reasonably successful in the RW environment by treating an uninformative outlier as if it were two sequential change points. In contrast, an RW-optimized agent might suffer more prolonged performance costs in the CP environment.

Limitations

There are limitations on our ability to generalize our findings to other situations and participant populations. An intriguing possibility is that different task contexts might evoke different patterns of bias and flexibility in metalearning. For example, our study focused solely on rapid, momentary decisions. Whether the type of spontaneous metalearning observed here would generalize to slower, deliberative processes is currently unknown. In addition, given that our participant samples were drawn from the Boston University community, the extent to which the results would generalize across populations, cultures, or clinically defined groups is an open question. Though we would expect the visual system's propensity for predictive gaze to hold

across populations, the relevant prior beliefs and the dynamics of feedback-driven cognitive flexibility could very well differ.

Conclusions

Two experiments demonstrated that participants calibrated patterns of learning rate modulation to the structure of their environment through experience alone. Participants displayed metalearning, adapting to the informativeness of surprising events in an initial context and readjusting after the environment's statistics reversed. The findings motivate new questions about sources of bias and individual variation in the cognitive processes that guide learning, prediction, and decision making in complex environments.

METHODS

Participants

All human participant procedures were approved by the Boston University Institutional Review Board. Informed consent was obtained from all participants. Participants were recruited as a convenience sample from the Boston University community (Experiment 1: N=29, 18 female and 11 male, Age: mean=20.5, range 18-24; Experiment 2: N=63, 48 female and 15 male, Age mean=20.3, range 18-31). All participants had normal or corrected-to-normal vision. Nine additional participants were excluded (5 in Experiment 1 and 4 in Experiment 2): three due to technical problems, three due to task accuracy below 60%, and three due to more than half of their trials not meeting eccentricity thresholds as explained below. The sample size for Experiment 1 was determined based on effect sizes derived from pilot data, while the sample size for Experiment 2 was determined based on effect sizes observed in Experiment 1, both using a power of 80% and a significance threshold of 5%. For Experiment 2, the sample size, exclusion criteria, and key analyses were preregistered with AsPredicted (aspredicted.org/cd54j.pdf).

Task

In two experiments, participants performed an implicit spatial prediction task based on Bakst & McGuire (2021; Figure 1A) programmed in Python using PsychoPy (Peirce et al., 2019) while their eye movements were tracked (EyeLink 1000+ desk-mounted eye tracker, SR Research Ltd, Osgoode, Canada). Gaze position was collected monocularly at 1000 Hz. Head position was stabilized using a chinrest positioned 57 cm from the display (BENQ XL2430 with a resolution of 1920x1080), and the eye tracker was calibrated at the beginning of each run. An additional post-hoc calibration was performed following data collection (see below). The nominal task was to report whether briefly presented numerical digits were even or odd in order to earn reward. The digits were presented in varying locations, implicitly requiring participants to anticipate the location of the next digit so that they could use central vision to make the odd/even judgment.

Each digit appeared between two flanking Xs for 180 ms before being backward-masked by another X. The participant then had unlimited time to respond by pressing “1” with their left hand for “odd” or “0” with their right hand for “even.” Accuracy feedback (a filled or empty circle signifying a correct or incorrect response) was then displayed in the same location as the digit for 500 ms, followed by a 750-ms inter-trial interval. Thus, there was a fixed 1250 ms interval between the keypress and the appearance of the next digit, making targets temporally predictable.

Two types of trials, *central* and *peripheral*, occurred in alternating order. Central trials displayed the digit at the center of the screen, which was marked by a small white point throughout the task. Peripheral trials displayed the digit somewhere on the perimeter of a circle centered on the screen with a radius of 6.9° of visual angle, which was marked throughout the

task with a white outline. The purpose of including central trials was to recenter gaze at a point equidistant from all possible digit locations prior to each peripheral trial.

Two conditions – *change point* (CP) and *random walk* (RW) – governed the peripheral digit locations. Average spatial predictability was similar between the two conditions but the sequential contingencies differed. In the CP condition, each digit's location was drawn from a Gaussian distribution with a fixed width ($\sigma = 11.25^\circ$ in angular distance around the circle) and a mean that usually remained fixed from one peripheral trial to the next but was resampled from a uniform distribution spanning the entire circle at occasional unsignaled change points. The generative mean was not resampled during a two-trial refractory period after each change point, and was resampled with a probability of 0.167 thereafter, leading to an overall change point probability of approximately 0.125. In the RW condition, digit locations were usually drawn from a Gaussian distribution ($\sigma = 11.25^\circ$) with a mean equal to the digit's location on the previous peripheral trial. This random walk around the circle was punctuated by occasional uniformly distributed outliers, which occurred at the same frequency as change points in the CP condition. In contrast to the CP condition, these noisy outliers did not reset the mean of the Gaussian; the subsequent peripheral digit location was drawn from a Gaussian distribution centered on the location of the peripheral digit prior to the outlier.

Therefore, in both conditions, peripheral digits had a high probability of appearing at a Gaussian-distributed location near the previous generative mean and a lower probability of appearing at an arbitrary new location anywhere on the circle. The two conditions differed in the predictive significance of large versus small movements. Small trial-to-trial differences in the CP condition tended to reflect noise around a stable mean, whereas large differences indicated a meaningful shift. In contrast, small trial-to-trial differences in the RW condition reflected

meaningful shifts in the mean, whereas large differences represented non-predictive outliers (Figure 1B).

Experiment 1 used a between-subject design in which each participant experienced only one condition, performing four 8-minute runs of either the CP or RW condition. Experiment 2 used a within-subject design in which each participant experienced three 8-minute runs of one condition followed by three 8-minute runs of the other, with the order counterbalanced across participants. Participants were not informed that the conditions changed, or that there were, in fact, multiple conditions at all. They were only told that they would complete four 8-minute runs in Experiment 1, and six 8-minute runs in Experiment 2.

Because run duration was based on time rather than number of trials (and the duration of individual trials depended on the participant's response latency), participants varied in the total number of trials they completed. In Experiment 1, participants completed an average of 456 peripheral trials in total across all four runs (range 382-515). They completed 462 peripheral trials on average in the CP condition (range 432-515), and 450 trials on average in the RW condition (range 382-484). In Experiment 2, participants completed an average of 687 trials in total across all six runs (range 556-780). They completed 345 trials on average in the CP condition (range 283-393) and 342 trials on average in the RW condition (range 258-393).

Monetary compensation for both studies consisted of a show-up payment of \$12 in addition to a bonus payment that scaled with the mean proportion of correct odd/even responses (95% accuracy resulted in a bonus of an additional \$12 and payments were rounded to the nearest \$0.25). A horizontal bar centered at the bottom of the screen increased in length proportional to the total number of correct responses in each run.

Following the experiment, participants were asked to respond to an open-ended question that asked, “What do you think the study was about?” Answers including any reference to digit locations/patterns or prediction were considered evidence of explicit awareness of task structure. Two independent raters assessed the responses and any discrepancies were resolved through discussion. Eleven individuals (38%) in the first experiment and 30 individuals (47%) in the second experiment were coded as having evidence for explicit awareness. Explicit awareness was not associated with odd/even judgment accuracy or predictive gaze accuracy in either experiment (all $p \geq 0.544$, Wilcoxon rank sum test).

Post-hoc calibration

Calibration of the eye tracker was performed at the start of each run. In addition, we performed post-hoc calibration of gaze position (Supplemental Figure 4). Using gaze position at the time of peripheral digit appearance (which served as our measure of predictive beliefs), we calculated the residuals according to the following:

$$Residuals = D - \sqrt{(B_x + X_{adj})^2 + (B_y + Y_{adj})^2}$$

where $D=6.9^\circ$ of visual angle, the eccentricity of peripheral digits. Measured gaze position was decomposed into its x and y components B_x and B_y relative to screen center, and X_{adj} and Y_{adj} represent additive post-hoc calibration adjustments. We then found the adjustments that minimized the following expression using `fminsearch` (Matlab):

$$\frac{\sum |Residuals|}{n} + 0.1 * (|X_{adj}| + |Y_{adj}|)$$

where n is the number of residuals. We only included trials on which gaze position was closer to the peripheral circle than the center (eccentricity $> 3.45^\circ$ of visual angle) so as to avoid basing the calibration on beliefs that lingered near the center of the screen. This procedure was

completed for each run separately. Runs were only recalibrated if they included more than 40 trials that met the eccentricity-based inclusion criterion. Two additional runs were excluded from recalibration after manual inspection revealed large, unexpected shifts in gaze position.

For X_{adj} , the mean across participants from both experiments was -0.086° , with a range from -4.908° to 6.862° . Y_{adj} had an average of 0.018° with a range from -2.757° to 3.621° . On average, the residuals per run decreased by 0.23° following post-hoc calibration.

Learning rate analyses

A participant's predictive belief on each peripheral trial, B_t , was operationalized as the gaze position at the time of digit onset, quantified in terms of angular position on the peripheral circle. The learning rate on trial t (LR_t) was estimated as the belief update from trial t to $t+1$, scaled by the prediction error on trial t :

$$LR_t = \frac{B_{t+1} - B_t}{X_t - B_t}$$

where B represents gaze-derived belief, and X is the angular position of the digit. The numerator and denominator were circular differences with a possible range of $\pm 180^\circ$. Analyses only included trials in which gaze was nearer to the peripheral circle than to the center (eccentricity $> 3.45^\circ$ of visual angle), to exclude beliefs inaccurately represented by the angle on the peripheral circle. Three participants in Experiment 1 and one participant in Experiment 2 were excluded entirely because fewer than half their trials met this eccentricity threshold.

Analyses assessed the relationship between per-trial learning rates and the absolute value of prediction errors. Median learning rate for each participant was calculated for different prediction error bins (bin edges were $0^\circ, 5^\circ, 10^\circ, 15^\circ, 20^\circ, 25^\circ, 35^\circ, 45^\circ, 75^\circ, 105^\circ, 135^\circ$, and 180°). This procedure was repeated separately for trials in “Early” and “Late” epochs, defined as

the first half (four minutes) of the first run and second half (four minutes) of the last run, respectively. In Experiment 1, this analysis compared the first half of Run 1 to the second half of Run 4. In Experiment 2, it compared the first half of Run 1 to the second half of Run 3 for the pre-reversal phase, and the first half of Run 4 to the second half of Run 6 for the post-reversal phase.

Additionally, a linear model with a probit link was fit to learning rate as a function of prediction error magnitude across individual trials, separately for each participant, with learning rate estimates constrained to the range -0.5 to 1.5:

$$\widehat{LR}_t = \Phi(\beta_0 + \beta_1 \cdot |PE_t|)$$

The `fmincon` function in Matlab (Mathworks, Natick, MA) was used to determine the least-squares best-fitting probit coefficients, using bounds of ± 3 for the intercept term (β_0) and ± 0.05 for the slope term (β_1). This was repeated separately for Early and Late epochs. Note that the probit function can take on values from 0 to 1, matching the theoretical range of learning rates. We allowed the input data (the empirical per-trial learning rate estimates) to retain a larger range to avoid truncating the measurement error distribution.

Subsequent analyses focused solely on change point or outlier trials, on which digit locations were sampled from a uniform distribution. These analyses were restricted to digits that appeared more than two standard deviations from the previous generative mean ($>22.5^\circ$ in angular distance) to focus only on extreme events. Outlier learning rate estimates were removed beforehand, defined as being outside the group interquartile range (IQR) by more than 1.5 times the IQR (which generally removed trials with learning rate estimates $>>1$ or $<<0$).

We calculated each participant's overall median learning rate for extreme events separately for the entire task experience, as well as the Early and Late epochs. The median Late-epoch

learning rate per participant was considered their “final LR.” The change in learning rate (ΔLR) was computed as the Late-minus-Early median learning rate for extreme events per participant. Participants with no valid learning-rate estimates for extreme events during the Early or Late epoch were excluded from this analysis ($n=3$ in Experiment 1, $n=15$ in the pre-reversal phase of Experiment 2, and $n=1$ in the post-reversal phase of Experiment 2). The retained participants had a median of three extreme events included in the Early epoch and five in the Late epoch.

The time course of metalearning was estimated by taking the group median and standard error of the learning rate for each successive extreme event experienced. Data were plotted until the latest point at which at least 50% of participants per group contributed data. A larger number of extreme events were available for Experiment 1 than for the pre-reversal phase of Experiment 2 because participants completed an additional run of the task. Additionally, a larger number of extreme events were available post-reversal than pre-reversal in Experiment 2 because response times tended to decrease with experience in the task, and participants therefore encountered more trials in the second half of the experiment. We also fit probability distributions to each participant’s set of extreme-event learning rates using a Gaussian kernel (bandwidth = 0.15).

The following comparisons were preregistered as primary confirmatory analyses in Experiment 2: (1) comparison of pre-reversal “final LR” between groups (to serve as a replication of Experiment 1); (2) comparison of post-reversal ΔLR between groups and against zero in each group separately (to test for recalibration after the reversal point); and (3) comparisons of both “final LR” and ΔLR between participant groups within each condition (to test for differences as a function of whether the condition was encountered first or second).

REFERENCES

- Bakst, L., & McGuire, J. T. (2021). Eye Movements Reflect Adaptive Predictions and Predictive Precision. *Journal of Experimental Psychology: General*, 150(5), 915–929.
<https://doi.org/10.1037/xge0000977>
- Behrens, T. E. J., Woolrich, M. W., Walton, M. E., & Rushworth, M. F. S. (2007). Learning the value of information in an uncertain world. *Nature Neuroscience*, 10(9), 1214–1221.
<https://doi.org/10.1038/nn1954>
- Browning, M., Behrens, T. E., Jocham, G., O'Reilly, J. X., & Bishop, S. J. (2015). Anxious individuals have difficulty learning the causal statistics of aversive environments. *Nature Neuroscience*, 18(4), 590–596. <https://doi.org/10.1038/nn.3961>
- Cheadle, S., Wyart, V., Tsetsos, K., Myers, N., de Gardelle, V., Herce Castañón, S., & Summerfield, C. (2014). Adaptive Gain Control during Human Perceptual Choice. *Neuron*, 81(6), 1429–1441. <https://doi.org/10.1016/j.neuron.2014.01.020>
- Clark, A. (2013). Whatever next? Predictive brains, situated agents, and the future of cognitive science. *Behavioral and Brain Sciences*, 36(3), 181–204.
<https://doi.org/10.1017/s0140525x12000477>
- d'Acremont, M., & Bossaerts, P. (2016). Neural Mechanisms Behind Identification of Leptokurtic Noise and Adaptive Behavioral Response. *Cerebral Cortex*, 26(4), 1818–1830.
<https://doi.org/10.1093/cercor/bhw013>
- Diederen, K. M. J., Spencer, T., Vestergaard, M. D., Fletcher, P. C., & Schultz, W. (2016). Adaptive Prediction Error Coding in the Human Midbrain and Striatum Facilitates Behavioral Adaptation and Learning Efficiency. *Neuron*, 90(5), 1127–1138.
<https://doi.org/10.1016/j.neuron.2016.04.019>
- Farashahi, S., Donahue, C. H., Hayden, B. Y., Lee, D., & Soltani, A. (2019). Flexible combination of reward information across primates. *Nature Human Behaviour*, 3(11), 1215–1224. <https://doi.org/10.1038/s41562-019-0714-3>
- Griffiths, T. L., Callaway, F., Chang, M. B., Grant, E., Krueger, P. M., & Lieder, F. (2019). Doing more with less: meta-reasoning and meta-learning in humans and machines. *Current Opinion in Behavioral Sciences*, 29, 24–30. <https://doi.org/10.1016/j.cobeha.2019.01.005>
- Hayhoe, M. M., McKinney, T., Chajka, K., & Pelz, J. B. (2012). Predictive eye movements in natural vision. *Experimental Brain Research*, 217(1), 125–136.
<https://doi.org/10.1007/s00221-011-2979-2>

641 Henderson, J. M. (2017). Gaze Control as Prediction. *Trends in Cognitive Sciences*, 21(1), 15–
642 23. <https://doi.org/10.1016/j.tics.2016.11.003>

643 Jepma, M., Brown, S. B. R. E., Murphy, P. R., Koelewijn, S. C., Vries, B. de, Maagdenberg, A.
644 M. van den, & Nieuwenhuis, S. (2018). Noradrenergic and Cholinergic Modulation of Belief
645 Updating. *Journal of Cognitive Neuroscience*, 30(12), 1803–1820.
646 https://doi.org/10.1162/jocn_a_01317

647 Jepma, M., Murphy, P. R., Nassar, M. R., Rangel-Gomez, M., Meeter, M., & Nieuwenhuis, S.
648 (2016). Catecholaminergic Regulation of Learning Rate in a Dynamic Environment. *PLOS*
649 *Computational Biology*, 12(10), e1005171. <https://doi.org/10.1371/journal.pcbi.1005171>

650 Kao, C.-H., Lee, S., Gold, J. I., & Kable, J. W. (2020). Neural encoding of task-dependent errors
651 during adaptive learning. *ELife*, 9, e58809. <https://doi.org/10.7554/elife.58809>

652 Kraus, N., Niedeggen, M., & Hesselmann, G. (2021). Trait anxiety is linked to increased usage
653 of priors in a perceptual decision making task. *Cognition*, 206, 104474.
654 <https://doi.org/10.1016/j.cognition.2020.104474>

655 Lee, S., Gold, J. I., & Kable, J. W. (2020). The Human as Delta-Rule Learner. *Decision*, 7(1),
656 55–66. <https://doi.org/10.1037/dec0000112>

657 Li, J., Schiller, D., Schoenbaum, G., Phelps, E. A., & Daw, N. D. (2011). Differential roles of
658 human striatum and amygdala in associative learning. *Nature Neuroscience*, 14(10), 1250–
659 1252. <https://doi.org/10.1038/nn.2904>

660 Massi, B., Donahue, C. H., & Lee, D. (2018). Volatility Facilitates Value Updating in the
661 Prefrontal Cortex. *Neuron*, 99(3), 598–608.e4. <https://doi.org/10.1016/j.neuron.2018.06.033>

662 McGuire, J. T., Nassar, M. R., Gold, J. I., & Kable, J. W. (2014). Functionally Dissociable
663 Influences on Learning Rate in a Dynamic Environment. *Neuron*, 84(4), 870–881.
664 <https://doi.org/10.1016/j.neuron.2014.10.013>

665 Nassar, M. R., Bruckner, R., & Frank, M. J. (2019). Statistical context dictates the relationship
666 between feedback-related EEG signals and learning. *ELife*, 8, e46975.
667 <https://doi.org/10.7554/elife.46975>

668 Nassar, M. R., Bruckner, R., Gold, J. I., Li, S.-C., Heekeren, H. R., & Eppinger, B. (2016). Age
669 differences in learning emerge from an insufficient representation of uncertainty in older
670 adults. *Nature Communications*, 7(1), 11609. <https://doi.org/10.1038/ncomms11609>

671 Nassar, M. R., Rumsey, K. M., Wilson, R. C., Parikh, K., Heasly, B., & Gold, J. I. (2012).
672 Rational regulation of learning dynamics by pupil-linked arousal systems. *Nature*
673 *Neuroscience*, 15(7), 1040–1046. <https://doi.org/10.1038/nn.3130>

- 674 Nassar, M. R., Wilson, R. C., Heasly, B., & Gold, J. I. (2010). An Approximately Bayesian
675 Delta-Rule Model Explains the Dynamics of Belief Updating in a Changing Environment.
676 *The Journal of Neuroscience*, 30(37), 12366–12378. [https://doi.org/10.1523/jneurosci.0822-](https://doi.org/10.1523/jneurosci.0822-10.2010)
677 10.2010
- 678 O'Reilly, J. X., Schüffelgen, U., Cuell, S. F., Behrens, T. E. J., Mars, R. B., & Rushworth, M. F.
679 S. (2013). Dissociable effects of surprise and model update in parietal and anterior cingulate
680 cortex. *Proceedings of the National Academy of Sciences of the United States of America*,
681 110(38), E3660-9. <https://doi.org/10.1073/pnas.1305373110>
- 682 Ossmy, O., Moran, R., Pfeiffer, T., Tsetsos, K., Usher, M., & Donner, T. H. (2013). The
683 Timescale of Perceptual Evidence Integration Can Be Adapted to the Environment. *Current*
684 *Biology*, 23(11), 981–986. <https://doi.org/10.1016/j.cub.2013.04.039>
- 685 Payzan-LeNestour, E., & Bossaerts, P. (2011). Risk, Unexpected Uncertainty, and Estimation
686 Uncertainty: Bayesian Learning in Unstable Settings. *PLoS Computational Biology*, 7(1),
687 e1001048. <https://doi.org/10.1371/journal.pcbi.1001048>
- 688 Payzan-LeNestour, E., Dunne, S., Bossaerts, P., & O'Doherty, J. P. (2013). The Neural
689 Representation of Unexpected Uncertainty during Value-Based Decision Making. *Neuron*,
690 79(1), 191–201. <https://doi.org/10.1016/j.neuron.2013.04.037>
- 691 Pearce, J. M., & Hall, G. (1980). A model for Pavlovian learning: Variations in the effectiveness
692 of conditioned but not of unconditioned stimuli. *Psychological Review*, 87(6), 532–552.
693 <https://doi.org/10.1037/0033-295x.87.6.532>
- 694 Peirce, J., Gray, J. R., Simpson, S., MacAskill, M., Höchenberger, R., Sogo, H., Kastman, E., &
695 Lindeløv, J. K. (2019). PsychoPy2: Experiments in behavior made easy. *Behavior Research*
696 *Methods*, 51(1), 195–203. <https://doi.org/10.3758/s13428-018-01193-y>
- 697 Pulcu, E., & Browning, M. (2019). The Misestimation of Uncertainty in Affective Disorders.
698 *Trends in Cognitive Sciences*, 23(10), 865–875. <https://doi.org/10.1016/j.tics.2019.07.007>
- 699 Soltani, A., & Izquierdo, A. (2019). Adaptive learning under expected and unexpected
700 uncertainty. *Nature Reviews Neuroscience*, 20(10), 635–644. [https://doi.org/10.1038/s41583-](https://doi.org/10.1038/s41583-019-0180-y)
701 019-0180-y
- 702 Summerfield, C., & Lange, F. P. de. (2014). Expectation in perceptual decision making: neural
703 and computational mechanisms. *Nature Reviews Neuroscience*, 15(11), 745–756.
704 <https://doi.org/10.1038/nrn3838>
- 705 Summerfield, C., & Tsetsos, K. (2015). Do humans make good decisions? *Trends in Cognitive*
706 *Sciences*, 19(1), 27–34. <https://doi.org/10.1016/j.tics.2014.11.005>

707 Urai, A. E., Braun, A., & Donner, T. H. (2017). Pupil-linked arousal is driven by decision
708 uncertainty and alters serial choice bias. *Nature Communications*, 8(1), 14637.
709 <https://doi.org/10.1038/ncomms14637>

710 Wang, J. X. (2021). Meta-learning in natural and artificial intelligence. *Current Opinion in*
711 *Behavioral Sciences*, 38, 90–95. <https://doi.org/10.1016/j.cobeha.2021.01.002>

712 Yu, A. J., & Dayan, P. (2005). Uncertainty, Neuromodulation, and Attention. *Neuron*, 46(4),
713 681–692. <https://doi.org/10.1016/j.neuron.2005.04.026>

714 Yu, L. Q., Wilson, R. C., & Nassar, M. R. (2021). Adaptive learning is structure learning in time.
715 *Neuroscience & Biobehavioral Reviews*, 128, 270–281.
716 <https://doi.org/10.1016/j.neubiorev.2021.06.024>

717

718 **ACKNOWLEDGEMENTS**

719 This work was supported by National Science Foundation grants 1755757 and 1809071,
720 National Eye Institute F32 Award EY029134, Office of Naval Research MURI Award N00014-
721 16-1-2832 and DURIP Award N00014-17-1-2304, and the Center for Systems Neuroscience
722 Postdoctoral Fellowship at Boston University. We thank Matthew Nassar, Sam Ling, and David
723 Somers for helpful discussion, and Lila Wright for assistance with data collection in Experiment
724 1.

725

726 Author contributions: LB and JTM contributed to experimental design. LB collected the data and
727 performed the analysis. LB prepared the manuscript with assistance from JTM.

729 **Supplemental Figure 1.** Effect of extreme events on task accuracy. (A) Experiment 1. Left,
730 Proportion correct odd/even judgments. Group mean and SEM shown for both central and
731 peripheral trials in each condition. Right, Gaze-based prediction error, same conventions as in A.
732 Peripheral prediction error shown in degrees of angular position (left axis), central prediction
733 errors shown in degrees of visual angle (right axis). (B) Experiment 2, pre-reversal. Same
734 conventions as in A.

Supplemental Figure 2. Within-participant distributions of extreme-event learning rates. (A) Experiment 1: Probability distributions were fit to individual participants' per-trial learning rates following extreme events (thin dotted lines). RW is shown in red, CP in blue. Group means are shown atop in thicker lines. Vertical dashed lines mark the optimal learning rate for CP (1.0) and RW (0.0). (B) Data from the pre-reversal phase of Experiment 2. Same conventions as in A. (C) Data from the post-reversal phase of Experiment 2. Same conventions as in A.

Supplemental Figure 3. Learning rate as a function of absolute prediction error in Experiment 2. (A) Pre-reversal: Learning rate as a function of prediction error (PE) magnitude in degrees of angular position on the peripheral circle (deg). Within-subject median learning rate was calculated for PE bins (bin edges = 0°, 5°, 10°, 15°, 20°, 25°, 35°, 45°, 75°, 105°, 135°, and 180°). Group mean-of-medians and SEM are shown for each condition: RW (red), CP (blue). Left panel is using all task data, middle is from the Early epoch (first four minutes of task experience), right is Late epoch (last four minutes). (B) Post-reversal data. Across both panels, dashed lines represent the participant group that experienced the CP condition in the pre-reversal phase and the RW condition in the post-reversal phase; solid lines represent the group that experienced the opposite order.

Supplemental Figure 4. Post-hoc calibration. (A) Example raw data. Left: digit locations around peripheral circle shown in black. Raw gaze position shown in blue. Red triangles indicate gaze position at the time of digit appearance on peripheral trials (gaze-based predictions). Right: Error between gaze-based prediction and digit location (in degrees of visual angle), as a function of angular position. (B) Same conventions as in A, but after the post-hoc calibration procedure was applied. (C) Summary of calibration over all participants in both experiments. Left: histogram of amount of adjustment in degrees of visual angle. Right: Average error between eye and digit per participant, pre- compared to post-calibration. Line represents equal pre- and post-calibration error.