

1 **Long title:** Differentiating Bayesian model updating and model revision based on their
2 prediction error dynamics

3 **Short title:** Model revision and model updating as two distinct learning mechanisms

4

5

6 Danaja Rutar^{12*}, Olympia Colizoli¹, Luc Selen¹, Lukas Spieß³, Johan Kwisthout^{1&} & Sabine

7 Hunnius^{1&}

8

9

10

11

12 ¹ Donders Institute for Brain, Cognition and Behaviour, Radboud University, Nijmegen, The

13 Netherlands

14 ² Leverhulme Centre for the Future of Intelligence, University of Cambridge, Cambridge, UK

15 ³ Qlip B.V., P.O. Box 119, 7200 AC Zutphen, The Netherlands

16

17

18

19

20

21

22 *Corresponding author

23 E-mail: danaja.rutar@donders.ru.nl

24

25 &JK and SK are joint senior authors.

26 **Abstract**

27 Within predictive processing learning is construed as Bayesian model updating with the degree
28 of certainty for different *existing* hypotheses changing in light of new evidence. Bayesian model
29 updating, however, cannot explain how *new* hypotheses are added to a model. Model revision,
30 unlike model updating, makes structural changes to a generative model by altering its causal
31 connections or adding or removing hypotheses. Whilst model updating and model revision have
32 recently been formally differentiated, they have not been empirically distinguished. The aim of
33 this research was to empirically differentiate between model updating and revision on the basis
34 of how they affect prediction errors and predictions over time. To study this, participants took
35 part in a within-subject computer-based learning experiment with two phases: updating and
36 revision. In the updating phase, participants had to predict the relationship between cues and
37 target stimuli and in the revision phase, they had to correctly predict a change in the said
38 relationship. Based on previous research, phasic pupil dilation was taken as a proxy for
39 prediction error. During model updating, we expected that the prediction errors over trials
40 would be gradually decreasing as a reflection of the continuous integration of new evidence.
41 During model revision, in contrast, prediction errors over trials were expected to show an abrupt
42 decrease following the successful integration of a new hypothesis within the existing model.
43 The opposite results were expected for predictions. Our results show that the learning dynamics
44 as reflected in pupil and accuracy data are indeed qualitatively different between the revision
45 and the updating phase, however in the opposite direction as expected. Participants were
46 learning more gradually in the revision phase compared to the updating phase. This could imply
47 that participants first built multiple models from scratch in the updating phase and updated them
48 in the revision phase.

49 **Introduction**

50 Imagine you recently moved from one desert city in Australia to another. In your old
51 neighbourhood all front yards were brown due to severe drought. However, in your new
52 neighbourhood, to your surprise, all the front yards appear to be green. You wonder how this is
53 possible since both cities have the same climate. You come up with potential reasons for this
54 and decide that probably the new neighbourhood has just experienced a period of heavy rain. It
55 is not until a week later that you see your neighbour changing their lawn for a new one. Now,
56 you realise that all the front yards in this town have artificial green lawns. At that moment you
57 add a new hypothesis to your model that you did not possess before.

58 What challenge does the example above pose? We make sense of the world through models
59 and use them for generating hypotheses about what the causes of sensory evidence are (Clark,
60 2016; Kwisthout et al., 2017). When a model fails to adequately predict the sensory evidence,
61 it needs to be adjusted. In the case above, a new hypothesis needs to be added to account for
62 the new observation. How does this happen? What mechanism allows for the integration of new
63 hypotheses into an existing model? This question has been one of the central questions in
64 cognitive science, but it still remains unanswered.

65 Here, we address this long-standing question within a predictive-processing framework, a
66 popular and influential theoretical framework in computational cognitive neuroscience (e.g.,
67 Friston et al., 2009; Friston et al., 2016; Friston et al., 2017; Clark, 2013; 2016; Hohwy, 2013).
68 The predictive-processing framework aims to become a unifying framework for understanding
69 the entirety of human cognition and behaviour from visual processing (Rao & Ballard, 1999;
70 Edwards et al., 2017; Petro & Muckli, 2016) to mentalizing (Kilner et al., 2007; Koster-Hale &
71 Saxe, 2013). According to the theory, the brain embodies a hierarchical generative model that
72 “aim[s] to capture the statistical structure of some set of observed inputs by tracking (one might

73 say, by schematically recapitulating) the causal matrix responsible for that very structure”
74 (Clark, 2013, p. 2). Based on this hierarchical model, the brain generates top-down predictions,
75 which are compared to the incoming sensory signals. The difference between the predicted and
76 the actual sensory input, i.e., the prediction error, is computed. From a predictive-processing
77 perspective, minimising reducible prediction error is the primary goal of computations in the
78 brain and occurs mainly as a result of learning (Friston et al., 2016; Smith et al., 2020;
79 Fitzgerald, 2015; Clark, 2016; Zenon, 2019).

80 Existing theories of learning in predictive-processing make use of Bayesian models of
81 cognition. As a natural consequence of using the Bayesian framework, most research on
82 learning has focused on one specific form of learning, Bayesian model updating, which explains
83 how model hypotheses are *updated* in the light of new evidence using Bayes’ rule (da Costa et
84 al., 2020; Friston et al., 2016; Smith et al., 2020). In other words, every time new sensory
85 evidence comes in, hypotheses in the learner’s model gain and lose probability (Kruschke,
86 2014; Perfors et al., 2011; Gopnik & Tenenbaum, 2007; Tenenbaum et al., 2006). Learning
87 accounts within the Bayesian framework are well suited for explaining updating of hypotheses
88 and selection among these existing hypotheses. However, they assume that the entire set of
89 relevant hypotheses and the mechanism for comparing hypotheses are predefined. Learning
90 under this view is reduced to *choosing* among existing hypotheses, whereas the process of
91 *generating* novel hypotheses is neglected (Perfors, 2012). Unless one assumes that learners are
92 equipped from the start with the complete set of hypotheses that they will ever consider, one
93 needs to account for how novel hypotheses are *added* to an existing model (Gentner & Hoyos,
94 2017; Schulz, 2012; Christie & Gentner, 2010). Bayesian updating, however, does not account
95 for this creative aspect of learning (Perfors, 2012; Gentner & Hoyos, 2017; Schulz, 2012; Xu,
96 2019; Christie & Gentner, 2010; Kwisthout et al., 2017). Kwisthout and colleagues (2017)
97 proposed that learning of novel hypotheses should be computationally differentiated from

98 Bayesian model updating and coined the term “model revision” for this alternative learning
99 mechanism. Model revision, unlike model updating, changes the structure of a generative model
100 by altering its causal connections or by adding and removing hypotheses (Kwisthout et al.,
101 2017). Similarly, Heald et al. (2022) have recently presented a theory for sensorimotor learning,
102 called contextual inference, that differentiates between the adaptation of behaviour based on
103 updating of existing and creation of new motor memories and adaptation due to changes in the
104 relative weighting of these motor memories.

105 Building on the theoretical argument that model updating and model revision are
106 computationally distinct, the aim of this study was to investigate whether these two learning
107 mechanisms can be empirically distinguished. To investigate this, we created an experiment
108 with two phases: an updating and a revision phase. In the updating phase, participants had to
109 learn the relationship between cues and a target stimulus. We induced the need for model
110 updating through a change in this relationship and examined whether the two experimental
111 phases would lead to two distinct patterns of learning. More specifically, we hypothesised that
112 the change in prediction error over time would follow distinct trajectories in the model updating
113 and model revision phases. In the model updating phase, the prediction error should *gradually*
114 decrease due to gradual updating of the probability distribution of existing model hypotheses.
115 In the model revision phase, prediction error should *abruptly* decrease once a new model
116 hypothesis has been added to the model (see section Hypotheses for a more detailed
117 description).

118 Whilst prediction error cannot be measured directly, it has been assessed indirectly, for instance
119 using EEG (Sambrook et al., 2018; Brydevall et al., 2018; Maier et al, 2019) or pupillometry
120 (Preuschhoff et al., 2011; Nassar et al., 2013; Browning et al., 2015; O’Reilly et al., 2013;
121 Colizoli et al., 2018; Lawson et al., 2017). In this study, we examined pupil dilation as a proxy
122 for so-called cognitive prediction error (den Ouden et al., 2012).

123 We hypothesised that prediction-error dynamics resulting from model updating and model
124 revision should also be reflected in participants' predictions of the upcoming target stimuli.
125 Prediction error minimisation is the result of learning more adequate models of the world
126 (Clark, 2013; 2016; Hohwy, 2013), reflected in the fact that predictions gradually become more
127 correct over time. We thus expected that in the model updating phase, predictions would only
128 gradually become better, whereas in the model revision phase, predictions should show a more
129 instantaneous change from incorrect to correct. Our initial idea was to study model updating
130 and model revision based on their *prediction error* minimisation dynamics. However, the
131 results we obtained were unexpected and, therefore, in addition to investigating prediction error
132 minimisation dynamics, we also performed an exploratory behavioural analysis focusing on
133 *predictions* based on button presses.

134 **Method**

135 **Participants**

136 Participants were recruited using Radboud University's online recruitment system. The only
137 restriction for participation was a minimum age of 16 years. Thirty-three healthy adults with
138 normal or corrected-to-normal vision participated in our study. One participant was excluded
139 for not following the instructions properly. Two participants were excluded due to equipment
140 malfunction. The final sample consisted of 30 participants (9 males, 21 females, mean age =
141 23.3, range: 19 – 42 years). The Ethical Committee of the Social Sciences Faculty at Radboud
142 University approved the study, and all participants gave written informed consent. Participants
143 received 15 euros for participating in the experiment.

144 **Task and procedure**

145 Participants carried out a computer-based two-alternative forced-choice (2AFC) task on the
146 *expected* orientation (left vs. right) of Gabor patches (Fig 1A). All participants took part in two
147 consecutive phases of the task, first the updating and then the revision phase. The entire
148 experiment took 1.5 hours and had three breaks (two short breaks halfway each phase and one
149 longer break between the two phases). After each break, recalibration of the eye-tracker took
150 place (see below). At the beginning of the experiment, participants were told that they would
151 be presented with auditory and visual cues. The instructions were to use the cues to predict the
152 orientation of the target stimulus in each trial. Before the start of the experiment, participants
153 were presented with an example trial. They indicated their prediction by pressing either the
154 right or left button on a button box for left or right orientation, respectively. Participants were
155 instructed to press the button as soon as they thought they knew which target orientation would
156 appear. At the end of the updating phase, participants were told that the cue-target rule would
157 change (i.e., the model revision phase) but not what the change would be. In the revision phase,
158 participants were similarly instructed to continue to use the cues to predict the orientation of
159 the target stimulus. The ‘model revision’ was a reversal of the cue-target contingencies in the
160 revision phase only for the trials which contained an auditory tone. Each phase (updating and
161 revision) of the experiment consisted of 200 trials (400 trials in total).

162 All participants were seated approximately 50 cm from the screen with their chin in a chin rest.
163 Stimuli were isoluminant and the environmental illumination was the same for all participants.
164 Stimuli were presented on a computer screen with the spatial resolution of 1920 × 1080 pixels.
165 One trial lasted approximately nine seconds. Each trial consisted of a fixation period, a cue
166 period, a response window followed by a delay period, and a target period. For all periods,
167 except for the target period, a vertically oriented Gabor patch was presented. The target stimulus

168 was a Gabor patch oriented to the left or to the right, with a spatial frequency of 0.033 and
169 opacity 0.5. Each trial started with a fixation cross on the vertical Gabor patch, which was
170 shown on the screen for 1500 ms. Afterward, the visual cue, which was either a square or a
171 diamond, was presented in the middle of the screen for 1000 ms. In 50% of the trials the visual
172 cue was paired with an auditory cue (i.e., tone “C”), which was presented for 300 ms. In the
173 other 50% of the trials, the auditory cue was absent. During the subsequent interval, the fixation
174 cross and the vertically oriented Gabor patch were presented, and participants were asked to
175 indicate their prediction about the upcoming orientation of the target by pressing a button. There
176 was no maximum response window. After a button had been pressed, a delay period started
177 with the vertical Gabor patch on the screen for additional 3000 ms. After the delay period, the
178 target stimulus (the Gabor patch tilted either left or right) was shown for 3000 ms. The durations
179 of the response period and target window were chosen in order to avoid contamination of the
180 pupil dilation response to a previous event. As our dependent measure was the pupil dilation
181 response to the target stimulus, the response window was sufficiently long to ensure that the
182 pupil response would not be contaminated by the motion response of the button press (Colizoli
183 et al., 2018). It is important to note that the target Gabor patch served as trial-by-trial feedback
184 on the accuracy of participants’ cue-target predictions. The prediction error in this context is
185 therefore the difference between the expected orientation of the target and the target’s actual
186 orientation.

187 The design contained probabilistic cue-target mappings to introduce uncertainty in the
188 predictions, thus simulating uncertainty that is inherent to perception in the real world. The
189 visual and auditory cues predicted whether the target Gabor patch was tilted to the right or to
190 the left with either an 80% or 20% probability (Fig 1B). In the updating phase, only the visual
191 cue had predictive value for learning the cue-target probabilities (i.e., the auditory tone did not
192 predict the target). In the revision phase, the visual cue-target mappings were reversed only for

193 trials in which an auditory cue was presented. Cue-target mappings were counterbalanced
194 between participants. Note that we define mapping 1 (M1) to correspond to the 80% visual cue-
195 target pairs in the updating phase and mapping 2 (M2) to correspond to the 20% visual cue-
196 target pairs in the updating phase.

197

198

199 **Figure 1. Experimental design and hypotheses** (A) Behavioural task. Participants performed a 2AFC task on the expected
200 orientation (left/right) of upcoming Gabor patches while pupil dilation was recorded. Each trial consisted of a fixation period,
201 a cue period, a response window followed by a delay period, and finally a target period. The decision interval ranged from
202 onset of the cue to the participant's response. The target interval ranged from target onset into the subsequent inter-trial interval
203 (3 s). The target served as feedback on the accuracy of participants' predictions in the decision interval. (B) The participants
204 had to learn cue-target contingencies in order to accurately predict the orientation of the upcoming Gabor patch (target).
205 Mapping 1 (M1) was defined as the visual cue-target pairs that occurred in 80% of trials in the updating phase; Mapping 2
206 (M2) was defined as the visual cue-target pairs that occurred in 20% of trials in the updating phase. Mappings were
207 counterbalanced between participants. At the start of the revision phase, the frequencies (80% vs. 20%) of the cue-target
208 mappings were reversed for trials containing the auditory tone cue only. (C) Main hypotheses for the dynamics of prediction
209 errors and accuracy over the course of the 2AFC task. The first 200 trials of the task represent the updating phase in which a
210 gradual increase in accuracy and a gradual decrease in the absolute value of prediction errors were expected (represented by
211 exponential curves). Within the revision phase (the last 200 trials), an abrupt increase in accuracy and an abrupt decrease in
212 prediction errors were expected (represented by sigmoidal curves).

213 **Hypotheses**

214 *Main hypotheses*

215 We hypothesised that in the updating phase, participants would be gradually learning the
216 probabilistic relationship between the cues and the Gabor patch orientation (target). New
217 sensory evidence, in this case the target-stimulus presentation, should deviate less from
218 expectations as evidence is accumulated over trials and thus generate smaller prediction errors
219 for the high frequency trials (i.e., their absolute values). We thus expected that in the updating

220 phase, the prediction errors would show a *gradual* decrease over time, illustrated by an
221 exponential curve. At the beginning of the revision phase, participants were instructed that
222 something would change during this revision phase. We expected that participants would
223 discover the new rule by integrating the now meaningful tone into their predictive models. The
224 new hypothesis should account for observations that were previously generating high prediction
225 errors. After an initial increase (compared to the ending point in the updating phase) in
226 prediction errors in the revision phase, an integration of a new hypothesis should lead to an
227 *abrupt* decrease in prediction errors (i.e., an “aha” moment), illustrated by sigmoidal curves.
228 Fig 1C (bottom) illustrates the expected prediction error during the 2AFC task. We expected a
229 similar dynamic for the tone and no tone trials in the updating and revision phase, however, it
230 could be that the learning curves for these trial types differed during revision due to the change
231 in contingency. Therefore, we investigated the tone and no-tone trials separately.

232 *Exploratory hypotheses*

233 While prediction errors reflect the deviation between expectations and feedback presentation,
234 the behavioural responses (i.e., button presses) are supposed to reflect the *predictions*
235 themselves. The experimental manipulation should affect accuracy data in a similar way to
236 prediction errors but in the opposite direction. We hypothesised that in the updating phase, the
237 accuracy of predictions should increase gradually and then plateau, while in the revision phase,
238 accuracy should initially decrease then abruptly plateau again as soon as the change in the cue-
239 target contingency is learned.

240 *Estimating prediction errors from pupil dilation*

241 Under constant luminance conditions, pupils dilate more when an error is made as compared
242 with a correct response (e.g., Braem et al., 2015). Recently, it was shown that pupil dilation
243 furthermore reflected an interaction between response accuracy and task difficulty in a 2AFC

244 motion-discrimination task with reward-coupled feedback, indicating its potential for
245 investigating reward prediction errors (Colizoli et al., 2018). In this previous work, the effect
246 of response accuracy (errors vs. correct trials) in the feedback-locked pupil dilation was larger
247 for hard trials as compared with easy trials on average. Following these results, we investigated
248 whether target-locked pupil dilation could be used as a measure of prediction errors in the
249 current study.

250 **Data acquisition and analyses**

251 *Data acquisition and pre-processing*

252 Changes in pupil dilation were recorded using an SMI RED500 eye-tracker (SensoMotoric
253 Instruments, Teltow/Berlin, Germany) with a sampling rate of 500 Hz. We analysed the pupil
254 dilation data of the right eye for each participant. The timing of blinks and saccade events was
255 not saved in the output of the eye-tracker; therefore, we did not attempt to categorize separate
256 blink and saccade events for pre-processing purposes. Pre-processing was applied to the entire
257 pupil dilation time series of each participant and consisted of: i) interpolation around missing
258 samples (0.15 s before and after each missing period), ii) interpolation around blinks or saccade
259 events based on spikes in the temporal derivative of the pupil time series (0.15 s before and
260 after each blink or saccade period), iii) band-pass filtering (third-order Butterworth, passband:
261 0.01–6 Hz), iv) removing responses to nuisance events using multiple linear regression (missing
262 periods and blink or saccade events were all categorized together as a single ‘nuisance’ event
263 type; responses were estimated by deconvolution, [Knapen et al., 2016]), and v) the residuals
264 of the nuisance regression were transformed to percent signal change with respect to the
265 temporal mean of the time series.

266 For each trial, intervals corresponding to the onset of the cue were extracted from each
267 participant’s pupil dilation time series (cue-locked and target-locked, respectively). The cue-

268 locked and target-locked intervals were baseline-corrected separately for each trial. The
269 baseline pupil was defined as the mean pupil in the time window -0.5 to 0 seconds with respect
270 to the cue or target for the cue-locked and target-locked intervals, respectively. The cue-locked
271 pupil response was analysed for data quality purposes, while the target-locked pupil response
272 was the main dependent variable of interest.

273

274 The temporal window of interest was independently defined as 1 to 2 seconds after the target
275 onset based on the pupil's canonical impulse response function (Knapen et al., 2016). For each
276 trial, a single value for the target-locked pupil response was computed as the mean pupil dilation
277 within this temporal window of interest. The event-locked and baseline-corrected scalar value
278 representing pupil dilation on each trial is referred to as the phasic pupil response in the
279 literature. Note that the pre-trial baseline pupil dilation is expected to reflect tonic fluctuations
280 in arousal in contrast to the task-evoked (or phasic) pupil responses investigated in the current
281 study (Murphy, O'Connell et al., 2014; Murphy, Vandekerckhove et al., 2014).

282 Trials were excluded if the reaction time (RT) was three standard deviations above the
283 participant's mean RT or lower than 100 ms (minimal time needed for the preparation of a
284 motor response; Dahan et al., 2019; Siegelman et al., 2019).

285 *Confirmatory analyses*

286 **Behaviour:** We expected that, on average, participants would learn the cue-target contingencies
287 in both phases of the task, which would be reflected in higher accuracy and faster responses for
288 high frequency trials as compared with low frequency trials. The effect of the visual cue-target
289 mapping condition was expected to interact with the auditory cue condition and task phase,
290 reflecting the reversal of the visual cue-target contingencies in the revision phase during the
291 tone trials only. These hypotheses were tested in two 3-way repeated measures ANOVAs,

292 separately for accuracy (as percentage of correct trials) and RT with factors: cue-target mapping
293 (M1 vs. M2), auditory cue (tone vs. no tone), and phase (updating vs. revision).

294

295 **Evoked pupil response:** The analysis of the evoked pupil response due to the cue presentation
296 served as a quality-control check on the goodness of the data acquisition. The visual and/or
297 auditory cues indicated to the participants that they should make a button press based on their
298 choice. A button press in the response phase was expected to evoke a canonical pupil impulse
299 response which should be reflected in the mean cue-locked pupil response. In addition, we
300 expected to see larger pupil dilation on average during tone trials as compared with no tone
301 trials in the cue-locked pupil response, as auditory cues are known to be arousing (e.g., Liao et
302 al., 2016). The cue-locked effect of tone was expected to return to baseline before the target
303 was presented on screen. We expected to see larger pupil dilation on average during errors as
304 compared with correct trials in the evoked target-locked pupil response.

305

306 **Phasic pupil response:** Following from the hypothesised pattern of behavioural results, we
307 expected to see an interaction between cue-target mapping, auditory cue, and task phase in the
308 target-locked phasic pupil response in the average phasic pupil response (this was tested with a
309 3-way repeated measures ANOVA as described above for behaviour). In the updating phase,
310 the average pupil dilation for the M2 mapping was expected to be larger as compared with the
311 M1 mapping, because low frequency trials (M2 in updating) should contain more errors overall.
312 The direction of the mapping effect was expected to reverse in the revision phase for tone trials
313 only.

314 **Curve fits on target-locked pupil data:** To test our main hypothesis concerning prediction
315 error dynamics across the updating and revision phases, we assessed whether the time course
316 of the target-locked pupil dilation showed the difference in the range parameter, σ , of the

317 sigmoid curve fits. The target-locked pupil dilation is taken as a proxy of prediction errors
318 within two phases of the experiment. The logic is that sensory evidence for the cue-target
319 contingencies is accumulated as the task progresses over time, and learning should be evident in
320 a reduction in prediction error magnitude as a function of trial order. The sigmoid function used
321 for the curve fits on target-locked pupil responses is given in Equation 1.

$$322 \quad f(x) = a_0 + G \int_{-\infty}^x \frac{1}{\sigma\sqrt{2\pi}} e^{-\frac{(x-\mu)^2}{2\sigma^2}} \quad (\text{Equation 1})$$

323 We fit four parameters, μ , σ , a_0 and G , of the above sigmoid function to the target-locked pupil
324 dilation across the high-frequency trials (80%) for the tone condition only, comparing the
325 updating and revision phases of the experiment. We differentiated between the high-frequency
326 as compared with low-frequency trials, in order to fit only a single direction of the (expected)
327 signed prediction errors (Colizoli et al., 2018; see den Ouden et al., 2012). The starting point of
328 the curve is reflected in the a_0 parameter. The inflection point of the curve is reflected in the μ
329 parameter. The gain parameter, G , allowed for negative scaling of the curves given the nature
330 of the pupil signal as dependent variable (i.e., percent signal change). The range parameter, σ ,
331 reflects the range over where the curve rises. A larger σ parameter is associated with a larger
332 range over which the transition (i.e., from $f(x) = a_0$ to $f(x) = 1$) takes place. We placed linear
333 constraints on the curve fits so that σ could not exceed three times the value of μ . We bound
334 both μ and σ so they could not exceed the number of trials in each phase of the experiment
335 (range between 1 and 200 trials). The parameters were furthermore not constrained or bounded.
336 The parameters were determined by minimising the ordinary least squares cost function.

337 Our main hypothesis was that learning would extend across more trials in the updating phase
338 as compared with in the revision phase, reflecting the difference between model updating and
339 model revision processes. Therefore, we expected that the σ would be higher for the high

340 frequency tone trials in the updating phase as compared with the high frequency tone trials in
341 the revision phase (note that the cue-target mappings were flipped for tone trials in the revision
342 phase; see Fig 1C). We furthermore expected the sign of the G to be negative in both phases,
343 indicating a decreasing trend. A larger value of σ together with a negative gain parameter, G,
344 thus reflects a more gradual reduction of prediction errors across trials. This analysis enables
345 us to examine whether pupil dilation depended on the experimental phase and if it scaled with
346 our hypotheses in Figure 1C. Descriptive statistics for all free parameters are presented in
347 Supplementary Table 1.

348 Finally, in the Supplementary Materials we describe the prediction errors as the 4-way
349 interaction between accuracy, cue-target mapping, auditory cue, and phase in the target-locked
350 pupil dilation.

351 *Exploratory analyses*

352 **Psychometric curve fits on accuracy data:** To test the hypothesis concerning the dynamics
353 of the target-orientation predictions, we assessed whether the range parameter, σ , of the
354 psychometric curve fits differed between the updating and revision phases. In a psychometric
355 curve, accuracy is plotted against signal intensity (May & Solomon, 2013). We explored the
356 resulting psychometric curves when the number of trials completed over time was taken as a
357 proxy for signal intensity. The sigmoid function used for the psychometric curve fits is given
358 in Equation 2.

$$359 \quad f(x) = a_0 + (1 - a_0) \int_{-\infty}^x \frac{1}{\sigma\sqrt{2\pi}} e^{-\frac{(x-\mu)^2}{2\sigma^2}} \quad (\text{Equation 2})$$

360 We fit three parameters, μ , σ , and a_0 , to the individual participant's response accuracy across
361 all trials for the tone condition only (i.e., the frequency conditions were not differentiated),

362 separately for the updating and revision phases of the experiment. We placed linear constraints
363 on the curve fits so that σ could not exceed three times the value of μ . We bound both μ and σ
364 so they could not exceed the number of trials in each phase of the experiment (range between
365 1 and 200 trials). The starting point, a_0 , was bound between 0 and 1. The parameters were
366 determined by minimising the negative log-likelihood cost function.

367 If our hypotheses about the difference between model revision and model updating are correct,
368 then σ should be higher for the tone trials in the updating phase as compared with the tone trials
369 in the revision phase. Descriptive statistics for all free parameters are presented in
370 Supplementary Table 1.

371 *Software*

372 The prediction task was administered with Psychopy (version 1.81; Psychopy, 2018). The
373 behavioural and pupil data were processed with custom software using Python (version 3.6;
374 Python Software Foundation, 2016). The evoked pupil responses were statistically assessed
375 with a cluster-level permutation test as part of the MNE-Python package (Gramfort et al., 2013).
376 Repeated measures ANOVAs and Bayesian tests were carried out in JASP (version 0.13.1;
377 JASP Team, 2020). All data and code are publicly available
378 (https://github.com/colizoli/learning_environmental_contingencies).

379 **Results**

380 The current study aimed to compare prediction-error dynamics, quantified in terms of pupil
381 dilation, during model updating and model revision. As predictions become more accurate over
382 time, due to learning more adequate models of the world, prediction errors are minimised
383 (Clark, 2013; 2016; Hohwy, 2013). Model revision, unlike model updating, changes the
384 structure of a generative model by altering its causal connections or by adding and removing

385 hypotheses (Kwisthout et al., 2017). To test these hypotheses, participants performed a 2AFC
386 task on the *expected* orientation (left vs. right) of an upcoming Gabor patch (target) task while
387 pupil dilation was recorded (Fig 1A). The participants had to learn cue-target contingencies in
388 order to accurately predict the orientation of the target. Cue-target contingencies changed at the
389 start of the revision phase in the following way: the cue-target mapping was reversed but only
390 for trials containing the auditory tone cue. The target served as feedback on the accuracy of
391 participants' predictions in the decision interval.

392 **Confirmatory analyses**

393 We first evaluated whether participants performed the 2AFC task as expected. Main effects and
394 interactions between the visual cue-target mapping condition (M1 vs. M2), task phase (updating
395 vs. revision), and the presence of the auditory cue (tone vs. no tone) were assessed in three
396 independent 3-way repeated measures ANOVAs on the dependent variables of mean accuracy
397 (Fig 2A), RT (Fig 2B), and target-locked pupil dilation (Fig 2C) (see Fig 2 for the analysis of
398 evoked pupil responses). The ANOVA results are presented in Table 1, and relevant post hoc
399 comparisons are presented in Fig 2. At the mean group level, a significant 3-way interaction
400 was obtained between visual cue-target mapping condition, auditory cue condition, and task
401 phase for accuracy and target-locked pupil response, but not for RT. Participants accurately
402 predicted the cue-target contingencies in both phases of the experiment, illustrated by the main
403 effect of visual cue-target mapping condition in the updating phase and the mapping reversal in
404 the revision phase for tone trials only (Fig 2A). As expected, the target-locked pupil response
405 was larger for the M2 trials as compared with the M1 trials during the updating phase, and the
406 presence of the auditory cue reversed the direction of the target-locked pupil response in the
407 revision phase for the tone trials (Fig 2C). Supplementary Fig 1 illustrates how the behaviour
408 and target-locked pupil dilation changed as a function of time (25 trials per bin).

409

410

411 **Figure 2. Cue-target prediction task results.** (A) Prediction accuracy, (B) mean RT, and (C) target-locked pupil dilation as
 412 a function of visual cue-target mapping condition (M1 vs. M2), the presence of the auditory cue (tone vs. no tone), and task
 413 phase (updating vs. revision). Results of the 3-way repeated measures ANOVAs are given in Table 1. Significance refers to
 414 post hoc t-tests: $**p < .01$, $***p < .001$. Error bars, s.e.m. ($N = 30$). Note that the frequencies of the visual cue-target mappings
 415 change in the revision phase for the tone trials.

416

417 **Table 1. Results of the 3-way repeated measures ANOVAs on accuracy, RT, and target-locked pupil response.** Factors
 418 of interest were cue-target mapping condition (levels: M1 vs. M2), auditory cue (levels: tone vs. no tone), and task phase
 419 (levels: updating vs. revision). Accuracy data were percentage of correct trials; RT was analysed in seconds. $*p < .05$, $**p <$
 420 $.01$, $***p < .001$

Effect	Accuracy			RT			Pupil response		
	<i>F</i> (1,29)	<i>p</i>	η^2_G	<i>F</i> (1,29)	<i>p</i>	η^2_G	<i>F</i> (1,29)	<i>p</i>	η^2_G
Mapping	301.22	< .001***	0.68	< .01	0.964	< .01	21.17	< .001***	0.05
Auditory cue	6.69	0.015*	< .01	4.91	0.035*	0.01	3.67	0.065	0.01
Phase	0.01	0.939	< .01	0.14	0.711	< .01	0.69	0.414	< .01
Mapping * Auditory cue	226.34	< .001***	0.65	0.16	0.694	< .01	6.79	0.014*	0.01
Mapping * Phase	201.03	< .001***	0.61	0.39	0.535	< .01	20.10	< .001***	0.04
Auditory cue * Phase	0.40	0.533	< .01	3.24	0.082	< .01	0.57	0.458	< .01
Mapping * Auditory cue * Phase	205.49	< .001***	0.64	3.52	0.071	< .01	9.71	0.004**	0.02

421

422 The data quality of the pupil dilation measures was assessed with several checks based on
 423 previous literature before testing our hypotheses about the dynamics of prediction errors across
 424 the experimental phases. First, evoked pupil dilation was present in response to the (visual and
 425 auditory) cue onsets as expected in the decision phase, here reflecting both decision preparation
 426 as well as the upcoming motor output in the form of a button press (Colizoli et al., 2018; Urai
 427 et al., 2017) (Fig 3A). Furthermore, the temporal window (1 to 2 s) chosen for the phasic pupil
 428 analysis contained the peak of the group-level cue-locked evoked response as expected from
 429 the pupil's canonical impulse response function (Fig 3A, grey box). Second, as expected, errors
 430 resulted in larger pupil dilation as compared with correct trials following the target presentation,

431 and this accuracy effect was significant within the temporal window of interest (Fig 3B, grey
432 box). Third, the presence of the auditory cue during the decision interval (cue-locked) was
433 associated with larger pupil dilation as compared with the absence of the tone (Fig 3C), likely
434 reflecting a difference in phasic arousal state during tone trials (Zekveld et al., 2018).
435 Importantly, the (unwanted) arousal effect related to the auditory tone was no longer present by
436 the time the target was presented for the participants (Fig 3D). We note that all further pupil
437 analyses used the (phasic) target-locked pupil dilation, averaged within the temporal window
438 of interest (Fig 3B, D, grey boxes), as the dependent variable. In sum, the pupil data fit all the
439 basic confirmatory hypotheses based on previous literature.

440

441

442 **Figure 3. Evoked pupil responses locked to the cue and target in the prediction task across both phases.** All trials within
443 the updating and revision phases of the prediction task were included in the evoked pupil response analysis. (A) Mean cue-
444 locked pupil responses in the prediction interval. Black bar indicates main effect of cue, $p < 0.05$ (cluster-based permutation
445 test). (B) Evoked pupil responses for correct and error trials in the feedback interval (target-locked). Black bar indicates correct
446 vs. error effect, $p < 0.05$ (cluster-based permutation test). (C) Evoked pupil responses for tone and no tone trials in
447 the prediction interval (cue-locked) and (D) in the feedback interval (target-locked). Black bar indicates tone vs. no tone
448 effect, $p < 0.05$ (cluster-based permutation test). In all panels: variability around the mean responses is illustrated as the 68%
449 confidence interval (bootstrapped; $N = 30$); the grey box indicates the temporal window of interest (1–2 s) with respect to event
450 onset for the phasic pupil responses. Note that the temporal window of interest was independently defined based on the pupil's
451 canonical impulse response function.

452 **Curve fits on target-locked pupil data:** Finally, we tested the main hypothesis concerning
453 prediction error dynamics reflected in pupil dilation across the updating and revision phases.
454 Sigmoid curves were fit for each participant's target-locked pupil response for the tone trials in
455 the high-frequency condition only, separately per phase (updating vs. revision). Individual
456 curve fits are shown in Supplementary Fig 2. Our main hypothesis was that σ would be larger
457 for the high-frequency tone trials in the updating phase as compared with the high-frequency

458 tone trials in the revision phase (i.e., the cue-target mappings were flipped for tone trials only
459 in the revision phase). For the trials without an auditory cue, we did not have expectations about
460 the difference in σ between the updating and revision phases of the task. Therefore, we tested
461 only tone trials for phase-dependent differences (updating vs. revision) of the mean σ parameter
462 (Fig 4B). To examine whether the σ parameters differed, we used Bayesian Wilcoxon Signed-
463 Rank Test (Fig 4B, right column). Whilst differences between σ parameters were obtained, they
464 were not in accordance with our expectations, since σ in the revision phase was larger as
465 compared with the updating phase.

466

467 **Figure 4. Comparisons of σ parameter of the curve fits on tone trials.** For each participant, sigmoid curves were fit to the
468 tone trials (i.e., trials with an auditory cue) and compared between the updating and revision phases of the experiment. (A) For
469 accuracy data, all tone trials (i.e., the high- and low-frequency conditions) were used to fit the curves. (B) For the target-locked
470 pupil response data, curves were fit to the high-frequency (80%) tone trials only. Note that the cue-target mappings (M1 and
471 M2) which correspond to the high-frequency trials differ per phase depending on the presence of the auditory cue. Results of
472 the Bayesian Wilcoxon Signed-Rank Test are shown for the accuracy and pupil data (right column). Error bars, s.e.m. ($N =$
473 30).

474 We also expected that the size of the prediction errors as measured with pupil dilation would
475 decrease as a result of learning, reflected in a negative gain parameter (G) in our curve fits for
476 both phases. In the updating phase, G was negative on average as expected ($M = -4.1$, $SD =$
477 11.94), however, G was positive on average in the revision phase ($M = 2.4$, $SD = 9.3$), against
478 our hypothesis of reduced prediction errors in both phases of the experiment. Furthermore, it
479 was apparent that the sign of G was not consistent at the individual level (see Supplementary
480 Fig 2).

481 **Exploratory analyses**

482 Given that the sizes of the σ parameters obtained from the curve fits of pupil data were the
483 opposite of what we had hypothesised, we decided to run an exploratory analysis on the
484 behavioural data. We expected that these additional analyses would provide us with insight into
485 what was going on during the task and would allow us to draw stronger conclusions.

486 **Psychometric curve fits on accuracy data:** For accuracy data, psychometric curves were fit
487 for each participant's response accuracy for the tone trials only, separately per phase (updating
488 vs. revision). Individual curve fits are shown in Supplementary Fig 3. As with pupil, our
489 exploratory hypothesis was that σ would be larger for the tone trials in the updating phase as
490 compared with the tone trials in the revision phase (i.e., the cue-target mappings were flipped
491 for tone trials only in the revision phase). For the trials without an auditory cue, we did not have
492 expectations about the difference in σ between the updating and revision phases of the task.
493 Therefore, we tested only tone trials for phase-dependent differences (updating vs. revision) of
494 the mean σ parameter (Fig 4A). To examine whether the σ parameters differed, we used a
495 Bayesian Wilcoxon Signed-Rank Test (Fig 4A, right column). Whilst differences between the
496 σ parameters were obtained, they were not in accordance with our expectations, since σ in the
497 revision phase was larger than in the updating phase.

498 **Discussion**

499 Our research was motivated by the observation that learning in the predictive-processing
500 framework is often reduced to Bayesian model updating (Perfors, 2012; Schulz, 2012; Gentner
501 & Hoyos, 2017). We argued that existing accounts of learning in the predictive-processing
502 framework currently lack a crucial component: a constructive learning mechanism that accounts
503 for changing models structurally when new hypotheses need to be learnt. Kwisthout and

504 colleagues (2017) proposed that model revision is a learning mechanism that is distinct from
505 Bayesian model updating and accounts for such a structural change in generative models.
506 Similarly, Heald and colleagues (2022) have recently presented a theory for sensorimotor
507 learning, called contextual inference, that differentiates between the adaptation of behaviour
508 based on updating of existing and creation of new motor memories and adaptation due to
509 changes in the relative weighting of these motor memories. Building on this, we investigated
510 whether we could empirically distinguish between Bayesian model updating and model
511 revision. We expected to be able to differentiate between the two mechanisms based on their
512 prediction error dynamics, which we assumed to be reflected in pupil dilation (Preuschoff et
513 al., 2011; Browning et al., 2015; O'Reilly et al., 2013; Colizoli et al., 2018; Lawson et al.,
514 2017). After obtaining unexpected results based on pupil data, we formed additional,
515 exploratory hypotheses based on the prediction dynamics reflected in accuracy. To investigate
516 our question, we created a within-subject computer-based experiment with two phases (the
517 updating and the revision phase). In the updating phase, participants had to learn and predict
518 the probabilistic relationship between the cues and a target stimulus, and in the revision phase,
519 a change occurred to what had been learnt previously.

520

521 To test whether participants performed the task as expected, we performed some basic quality
522 controls on behavioural and pupil data. Behavioural data showed that participants on average
523 correctly predicted the cue-target contingencies in both experimental phases. This was reflected
524 in the cue-target mapping effect for tone and no tone trials in the updating phase, and a reversed
525 effect of cue-target mapping effect was observed in the revision phase for tone trials (when the
526 mapping switched) only. As expected, a similar pattern was also observed in pupil data.

527 Before turning to the main research question, the quality of the pupil measures was checked
528 and compared to effects observed in previous literature. Decision preparation and the

529 preparation of the motor response were reflected in the evoked pupil response to the visual and
530 auditory cue onsets replicating previous work (Colizoli et al., 2018; Urai et al., 2017). We also
531 observed the peak of the group-level cue-locked evoked response around 1.5 seconds, which is
532 in line with previous findings (e.g., Mathôt, 2018). Furthermore, erroneous responses resulted
533 in larger pupil dilation as compared to correct responses, an effect that has also been previously
534 reported (Braem et al., 2015). Finally, we found a well-known auditory effect on pupil
535 responses with pupil responses being bigger on trials with a tone compared to trials without a
536 tone (e.g., Liao et al., 2016).

537 Given that participants understood and correctly performed the task, we turn to the results
538 related to our main hypotheses. Multiple studies (Friston et al., 2016; Smith et al., 2020;
539 Fitzgerald, 2015; Clark, 2016; Zenon, 2019) have shown that prediction error decreases as a
540 result of learning. Building on this we hypothesised that in the updating and the revision phase,
541 different temporal dynamics in prediction error decrease would be observed. We hypothesised
542 that in the updating phase participants would be gradually learning the probabilistic relationship
543 between the cues and the target orientation, leading to model updating. As a result of model
544 updating, participants should become better at predicting future sensory input. New sensory
545 observations should thus yield prediction errors that are smaller than the preceding ones, and
546 the degree to which the model needs to be updated should decrease over time. In other words,
547 we expected prediction error to decrease *gradually* over time. In the revision phase the rules
548 switched for the tone trials, which should initially lead to an increase in prediction error. When
549 the change in rules is learned, a new hypothesis is added to a learner's model, resulting in model
550 revision. An integration of a new hypothesis should lead to an *abrupt* decrease in prediction
551 errors (i.e., an “aha” moment). Whilst prediction errors reflect the deviation between
552 expectations and target-feedback presentation, the behavioural responses (i.e., button presses)

553 are supposed to reflect the *predictions* themselves. Thus, we expected to find a similar pattern
554 of responses in accuracy data as in prediction errors but in the opposite direction.

555 Curve fitting revealed significant differences between the phases of the task in pupil and
556 accuracy data as reflected in the σ parameter. However, the difference was in the opposite
557 direction as expected; the σ was significantly smaller in the updating phase compared to the
558 revision phase, in accuracy and pupil data. These results suggest that the participants were
559 learning more gradually in the revision phase compared to the updating phase of the experiment,
560 contrary to our expectations.

561 One possible interpretation of these results is that our experimental manipulation induced model
562 revision in the updating phase and model updating in the revision phase. It might have been that
563 in the updating phase participants built multiple internal models, that they thought could capture
564 the structure of the task, from scratch. As they were learning the task they were alternating
565 between these models and upon the realisation of the rule participants settled for the correct
566 model, resulting in a rapid decrease in prediction error and accuracy as the data in the updating
567 phase shows. We assumed we prevented participants from learning models from scratch in the
568 updating phase by providing them with detailed instructions and pictorial representation of the
569 stimuli resented in the task, before the task started. By that, we thought, we equipped
570 participants with a crude model that would contain hypotheses about all the relevant variables
571 of the task. However, in light of the current results we believe that our instructions did not result
572 in the construction of a simple model that participants could use as a baseline upon entering the
573 task.

574 At the beginning of the revision phase, participants expected for the rules of the task to change.
575 All the experimental variables (e.g., target, visual cues) in the revision phase were the same as
576 in the updating phase, possibly signalling to the participants that one of the models that they

577 had already constructed in the updating phase could be suitable for explaining the change in the
578 revision phase. If this was the case, then participants in the revision phase merely reused a
579 correct, existing model, and started gradually updating that model rather than adding a new
580 hypothesis (based on a new experimental variable) to a model they had built in the updating
581 phase.

582 We could possibly have avoided participants building models from scratch in the updating
583 phase had we ran a computerised familiarisation phase with all the relevant experimental
584 variables with the participants prior to the experiment. This would make sure that participants
585 have constructed crude models of the task before the experiment started and that they could
586 then use and update in the update phase. Additionally, to make sure that participants in the
587 revision phase were not just reusing one of the models they had constructed in the updating
588 phase but instead build on an existing model, we could have introduced a new experimental
589 variable in the revision phase that was present in the familiarisation but not in the updating
590 phase. In that case, participants would have to add a new hypothesis to a model (instantiating
591 model revision) constructed in the updating phase, if they were to successfully learn the new
592 rule in the revision phase.

593 Our results also revealed that the gain parameter, G , which indicates the direction of the σ
594 parameter, was negative in the updating phase as expected and positive in the revision phase
595 contrary to our expectations. This suggests that pupil dilation and hence prediction error was
596 on average decreasing in the updating phase and increasing in the revision phase. These results
597 are unexpected, because, according to existing literature, prediction error is decreasing as a
598 result of learning (Friston et al., 2016; Smith et al., 2020; Fitzgerald, 2015; Clark, 2016; Zenon,
599 2019) and so should pupil dilation (Kayhan, 2019). However, a look at the individual level pupil
600 data (see supplementary Fig 2) reveals large variability in the sign of the σ parameter for both

601 phases, potentially suggesting individual differences in the size of pupil dilation over time. This
602 suggestion is in line with recent findings of substantial inter- and intra-individual variation in
603 the size of pupil dilation over trials (Sybley et al., 2020). More specifically, the study shows
604 that in a simple digit-span memory task pupil dilation was consistently increasing over trials
605 for some participants and for others it was decreasing. There were also participants for whom
606 the trend was changing throughout the task. Another factor that could explain why pupil dilation
607 was increasing in the revision phase is that participants in this phase were more fatigued than
608 in the updating phase at the beginning of the experiment. As a consequence, they had to exert
609 more effort to maintain concentrated on the task and process the change in the task rules.
610 Increased cognitive effort would result in increasing pupil dilation as has been shown many
611 times before (for review see van der Wel & van Steenbergen, 2018; Hyönä et al., 1995; Porter
612 et al., 2007).

613 All in all, our data shows that there exists a qualitative difference between model updating and
614 model revision, following the initial theoretical proposal from Kwisthout and colleagues
615 (2017). However, model updating seemed to have occurred in the model revision phase and
616 model revision in the model updating phase of the experiment for reasons described above.
617 Future studies should therefore make sure to induce experimental manipulations that have the
618 initially intended effect. Alternatively, future studies might investigate empirically the
619 interpretation of the current results: that participants construct multiple models at the beginning
620 and later choose among them and update them. Lastly, our study is one of the few that studied
621 how the change in prediction error, as captured by pupil dilation, over time (for a few exceptions
622 see Kayhan, 2019; Koenig et al., 2017). Therefore, little is known about how pupil dynamics
623 and hence prediction errors change over extended periods of time and whether individual
624 differences exist in this process. More studies should thus examine pupil dilation and hence
625 prediction error in such a manner in the future.

626

627

628 **References**

629 Braem S, Coenen E, Bombeke K, van Bochove ME, Notebaert W. Open your eyes for
630 prediction errors. *Cognitive, Affective, & Behavioral Neuroscience*. 2015; 15(2): 374-380.

631 Browning M, Behrens TE, Jocham G, O'Reilly JX, Bishop SJ. Anxious individuals have
632 difficulty learning the causal statistics of aversive environments. *Nature Neuroscience*. 2015;
633 18(4): 590-596.

634 Brydevall M, Bennett D, Murawski C, Bode S. The neural encoding of information prediction
635 errors during non-instrumental information seeking. *Scientific Reports*. 2018; 8(6134): 1-11.

636 Christie S, Gentner D. Where hypotheses come from: Learning new relations by structural
637 alignment. *Journal of Cognition and Development*. 2010; 11(3): 356-373.

638 Clark A. *Surfing Uncertainty: Prediction, Action, and the Embodied Mind*. New York, NY:
639 Oxford University Press; 2016.

640 Clark A. Whatever next? Predictive brains, situated agents, and the future of cognitive science.
641 *Behavioral and Brain Sciences*. 2013; 36(3): 181-253.

642 Colizoli O, de Gee JW, Urai AE, Donner TH. Task-evoked pupil responses reflect internal
643 belief states. *Scientific Reports*. 2018; 8(1): 1-13.

644 da Costa L, Parr T, Sajid N, Veselic S, Neascu V, Friston KJ. Active inference on discrete state-
645 spaces: A synthesis. *Journal of Mathematical Psychology*. 2020; 99(102447): 102447.

- 646 Dahan A, Bennet R, Reiner M. How Long Is Too Long: An individual time-window for motor
647 planning. *Frontiers in Human Neuroscience*. 2019; 13(238): 1-9.
- 648 de Gee JW, Colizoli O, Kloosterman NA, Knapen T, Nieuwenhuis S, Donner TH. Dynamic
649 modulation of decision biases by brainstem arousal systems. *Elife*. 2007; 6: e23232.
- 650 den Ouden HEM, Kok P, de Lange, F. How prediction errors shape perception, attention, and
651 motivation. *Frontiers in Psychology*. 2012; 3, 548.
- 652 Edwards G, Vetter P, McGruerm F, Petro LS, Muckli L. Predictive feedback to V1 dynamically
653 updates with sensory input. *Scientific Reports*. 2017; 7(1): 1-12.
- 654 FitzGerald THB, Dolan RJ, Friston KJ. Dopamine, reward learning, and active inference.
655 *Frontiers in Computational Neuroscience*. 2015; 9(136).
- 656 Friston KJ, FitzGerald T, Rigoli F, Schwartenbeck P, O'Doherty J, Pezzulo G. Active inference
657 and learning. *Neuroscience & Biobehavioral Reviews*. 2016; 68: 862-879.
- 658 Friston KJ, Kiebel S. Predictive coding under the free-energy principle. *Philosophical
659 Transaction of the Royal Society B*. 2009; 364(1521): 1211-1221.
- 660 Friston KJ, Lin M, Frith CD, Pezzulo G, Hobson JA, Ondobaka S. Active Inference, Curiosity
661 and Insight. *Neural Computation*. 2017; 29(10): 2633-2683.
- 662 Gentner D, Hoyos C. Analogy and abstraction. *Topics in Cognitive Science*. 2017; 9: 672-639.
- 663 Gopnik A, Tenenbaum JB. Bayesian networks, Bayesian learning and cognitive development.
664 *Developmental Science*. 2007; 10(3): 281-287.
- 665 Gramfort A, Luessi M, Larson E, Engemann DA, Strohmeier D, Brodbeck C, et al. MEG and
666 EEG data analysis with MNE-Python. *Frontiers in Neuroscience*. 2013; 7, 267.

- 667 Heald JB, Lengyel M, Wolpert, DM. Contextual inference underlies the learning of
668 sensorimotor repertoires. *Nature*. 2021; 600(7889): 489-493.
- 669 Hohwy J. *The Predictive Mind*. Oxford: Oxford University Press; 2013.
- 670 Hyönä J, Tommola J, Alaja AM. Pupil dilation as a measure of processing load in simultaneous
671 interpretation and other language tasks. *The Quarterly Journal of Experimental Psychology*.
672 1995; 48(3), 598-612.
- 673 Kayhan E, Heil L, Kwisthout J, van Rooij I, Hunnius S, Bekkering H. (2019). Young children
674 integrate current observations, priors and agent information to predict others' actions. *Plos One*.
675 2019; 14(5), e0200976.
- 676 Kilner JM, Friston KJ, Frith CC. Predictive coding: an account of the mirror neuron system.
677 *Cognitive Processing*. 2007; 8(3): 159-166.
- 678 Knapen T, de Gee JW, Brascamp J, Nuiten S, Hoppenbrouwers S, Theeuwes J. Cognitive and
679 ocular factors jointly determine pupil responses under equiluminance. *Plos One*. 2016; 11(5):
680 e0155574.
- 681 Koenig S, Uengoer M, Lachnit H. Pupil dilation indicates the coding of past prediction errors:
682 Evidence for attentional learning theory. *Psychophysiology*. 2018; 55(4), e13020.
- 683 Koster-Hale J, Saxe R. Theory of mind: a neural prediction problem. *Neuron*. 2013; 79(5): 836-
684 848.
- 685 Kruschke JK. *Doing Bayesian Data Analysis: A Tutorial with R, JAGS, and Stan*. 2nd ed.
686 Amsterdam: Academic Press; 2015.

- 687 Kwisthout J, Bekkering H, van Rooij I. To be precise, the details don't matter: On predictive
688 processing, precision, and level of detail of predictions. *Brain and Cognition*. 2017; 112: 84-
689 91.
- 690 Larsen RS, Waters J. Neuromodulatory correlates of pupil dilation. *Frontiers in Neural*
691 *Circuits*.. 2018; 12, 21.
- 692 Lawson RP, Mathys C, Rees G. Adults with autism overestimate the volatility of the sensory
693 environment. *Nature Neuroscience*. 2017; 20(9): 1293-1299.
- 694 Liao HI, Yoneya M, Kidani S, Kashino M, Furukawa S. Human pupillary dilation response to
695 deviant auditory stimuli: Effects of stimulus properties and voluntary attention. *Frontiers in*
696 *Neuroscience*. 2016; 10, 43.
- 697 Maier EM, Ernst B, Steinhauser M. Error-related pupil dilation is sensitive to the evaluation of
698 different error types. *Biological Psychology*. 2019; 141: 25-34.
- 699 Mathôt S. Pupillometry: psychology, physiology, and function. *Journal of Cognition*.
700 2018; 1(1).
- 701 May KA, Solomon JA. Four theorems on the psychometric function. *Plos One*. 2013; 8(10):
702 e74815.
- 703 Murphy PR, O'Connell RG, O'Sullivan M, Robertson IH, Balsters JH. Pupil diameter covaries
704 with BOLD activity in human locus coeruleus. *Human Brain Mapping*. 2014; 35(8), 4140-
705 4154.
- 706 Murphy PR, Vandekerckhove J, Nieuwenhuis S. Pupil-linked arousal determines variability in
707 perceptual decision making. *Plos Computational Biology*. 2014; 10(9), e1003854.

- 708 Nassar MR, Rumsey KM, Wilson RW, Parikh K, Heasley B, Gold JJ. Rational regulation of
709 learning dynamics by pupil-linked arousal systems. *Nature Neuroscience*. 2013; 15(7): 1040-
710 1046.
- 711 O'Reilly JX, Schüffelgen U, Cuell SF, Behrens TE, Mars RB, Rushworth MF. Dissociable
712 effects of surprise and model update in parietal and anterior cingulate cortex. *PNAS*. 2013;
713 110(38): E3660-E3669.
- 714 Perfors A, Tenenbaum JB, Griffiths RL, Xu F. A tutorial introduction to Bayesian models of
715 cognitive development. *Cognition*. 2011; 120(3): 302-21.
- 716 Perfors A. Bayesian Models of Cognition: What's Built in After All? *Philosophy Compass*.
717 2012; 7(2): 127-138.
- 718 Petro LS, Muckli L. The brain's predictive prowess revealed in primary visual cortex. *PNAS*.
719 2016; 113(5): 1124-1125.
- 720 Porter G, Troscianko T, Gilchrist ID. Effort during visual search and counting: Insights from
721 pupillometry. *Quarterly Journal of Experimental Psychology*. 2007; 60(2), 211-229.
- 722 Preuschoff K, 't Hart BM, Einhäuser W. Pupil dilation signals surprise: evidence for
723 noradrenaline's role in decision making. *Frontiers in Neuroscience*. 2011; 5: 115.
- 724 Psychopy. (Version 1.81). [Computer software]. 2018. University of Nottingham; Available
725 from: <https://psychopy.org/index.html>
- 726 Python Software Foundation. Python (Version 3.6.0) [Computer software]. 2016. Available
727 from: <https://www.python.org/downloads/release/python-360/>

- 728 Rao RP, Ballard DH. Predictive coding in the visual cortex: a functional interpretation of some
729 extra-classical receptive-field effects. *Nature Neuroscience*. 1999; 2(1): 79-87.
- 730 Sambrook TD, Hardwick B, Wills AJ, Goslin J. Model-free and model-based reward prediction
731 errors in EEG. *Neuroimage*. 2018; 178: 162-171.
- 732 Schulz L. Finding new facts; Thinking new thoughts. *Rational Constructivism in Cognitive*
733 *Development*. 2012; 43: 269-294.
- 734 Sibley C, Foroughi CK, Brown NL, Phillips H, Drollinger S, Eagle M, Coyne JT. More than
735 means: characterizing individual differences in pupillary dilations. Proceedings of the Human
736 Factors and Ergonomics Society Annual Meeting; 2020 Dec; Los Angeles. Los Angeles, CA:
737 SAGE Publications; 2020. p. 57-61.
- 738 Siegelman N, Bogaerts L, Armstrong BC, Frost R. What exactly is learned in visual statistical
739 learning? Insights from Bayesian modeling. *Cognition*. 2019; 192: 104002.
- 740 Smith R, Schwartenbeck P, Parr T, Friston KJ. An Active Inference Approach to Modeling
741 Structure Learning: Concept Learning as an Example Case. *Frontiers in Computational*
742 *Neuroscience*. 2020; 14(41).
- 743 Spiess AN, Neumeier N. An evaluation of R2 as an inadequate measure for nonlinear models
744 in pharmacological and biochemical research: a Monte Carlo approach. *BMC Pharmacology*.
745 2010; 10(1): 1-11.
- 746 Tenenbaum JB, Griffiths TL, Kemp C. Theory-based Bayesian models of inductive learning
747 and reasoning. *TRENDS in Cognitive Science*. 2006; 10(7): 309-318.
- 748 The JASP Team. JASP (Version 0.13.1) [Computer software]. 2020. Amsterdam: Univeristy
749 of Amsterdam; Available from: <https://jasp-stats.org/2020/07/21/jasp-0-13-1-small-fixes/>

- 750 Urai AE, Braun A, Donner TH. Pupil-linked arousal is driven by decision uncertainty and
751 alters serial choice bias. *Nature Communications*. 2017; 8(1), 1-11.
- 752 van der Wel P, van Steenbergen H. Pupil dilation as an index of effort in cognitive control tasks:
753 A review. *Psychonomic Bulletin & Review*. 2018; 25(6), 2005-2015.
- 754 van der Wel P, van Steenbergen H. Pupil dilation as an index of effort in cognitive control
755 tasks: A review. *Psychonomic Bulletin & Review*. 2018; 25(6), 2005-2015.
- 756 van Slooten JC, Jahfai, S, Knapen, T, Theeuwes, J. How pupil responses track value-based
757 decision-making during and after reinforcement learning. *Plos Computational Biology*. 2018;
758 14(11): e1006632.
- 759 Wichmann FA, Hill NJ. The psychometric function: I. Fitting, sampling, and goodness of
760 fit. *Perception & Psychophysics*. 2001; 63(8): 1293-1313.
- 761 Xu F. Towards a Rational Constructivist Theory of Cognitive Development.
762 *Psychological Review*. 2019; 126(6): 841-864.
- 763 Zekveld AA, Koelewijn T, Kramer SE. The pupil dilation response to auditory stimuli:
764 current state of knowledge. *Trends in Hearing*. 2018; 22: 2331216518777174.
- 765 Zénon A. Eye pupil signals information gain. *Proceedings of the Royal Society B*. 2019;
766 286(1911): 20191593.

767 **Supplementary material**

768 **Supplementary Figure 1. Cue-target prediction task results as a function of trial bin.** Prediction accuracy (top row), mean
769 RT (middle row), and target-locked pupil dilation (bottom row) as a function of visual cue-target mapping condition (M1 vs.
770 M2), the presence of the auditory cue (tone vs. no tone), and consecutive trial bin. Trial bins 1-8 correspond to the updating
771 phase, and trial bins 9-16 correspond to the revision phase (25 trials per bin).

772

773 **Supplementary Figure 2. Curve fits for the target-locked pupil response (single-trial).** Individual curves were fit on the
 774 target-locked pupil response data for the high-frequency tone trials (i.e., the auditory cue and 80% frequency condition)
 775 separately for each phase (updating vs. revision).

776

777 **Supplementary Figure 3. Psychometric curve fits for accuracy (single trial).** Individual psychometric curves were fit on
 778 the response accuracy data for the tone trials (auditory cue condition) separately for each phase (updating vs. revision).

779

780 **Supplementary Table 1. Descriptive statistics on free parameters from curve fits for the accuracy and pupil data.**

Accuracy

	μ		σ		$a0$	
	Updating	Revision	Updating	Revision	Updating	Revision
<i>Valid</i>	30	30	30	30	30	30
<i>Missing</i>	0	0	0	0	0	0
<i>Median</i>	191.29	400	1	132.437	0.711	0.649
<i>Mean</i>	187.996	389.102	17.737	108.104	0.69	0.617
<i>Std. Deviation</i>	13.053	22.626	26.1	36.207	0.08	0.108
<i>Minimum</i>	156.964	323.083	1	41.918	0.492	0.351
<i>Maximum</i>	200	400	66.667	133.333	0.779	0.77

Pupil response

	μ		σ		$a0$		G	
	Updating	Revision	Updating	Revision	Updating	Revision	Updating	Revision
<i>Valid</i>	30	30	30	30	30	30	30	30
<i>Missing</i>	0	0	0	0	0	0	0	0
<i>Median</i>	130.827	246.795	0.099	1.766	-1.269	-0.365	-2.522	1.036
<i>Mean</i>	122.354	262.328	0.612	10.772	0.767	-1.99	-4.136	2.397
<i>Std. Deviation</i>	63.868	60.348	1.656	19.487	9.846	8.267	11.938	9.307
<i>Minimum</i>	0.747	186.918	1.369e -5	8.487e -4	-5.082	-43.248	-51.324	-7.209
<i>Maximum</i>	198.993	457.277	7.008	69.092	51.521	6.503	8.48	45.274


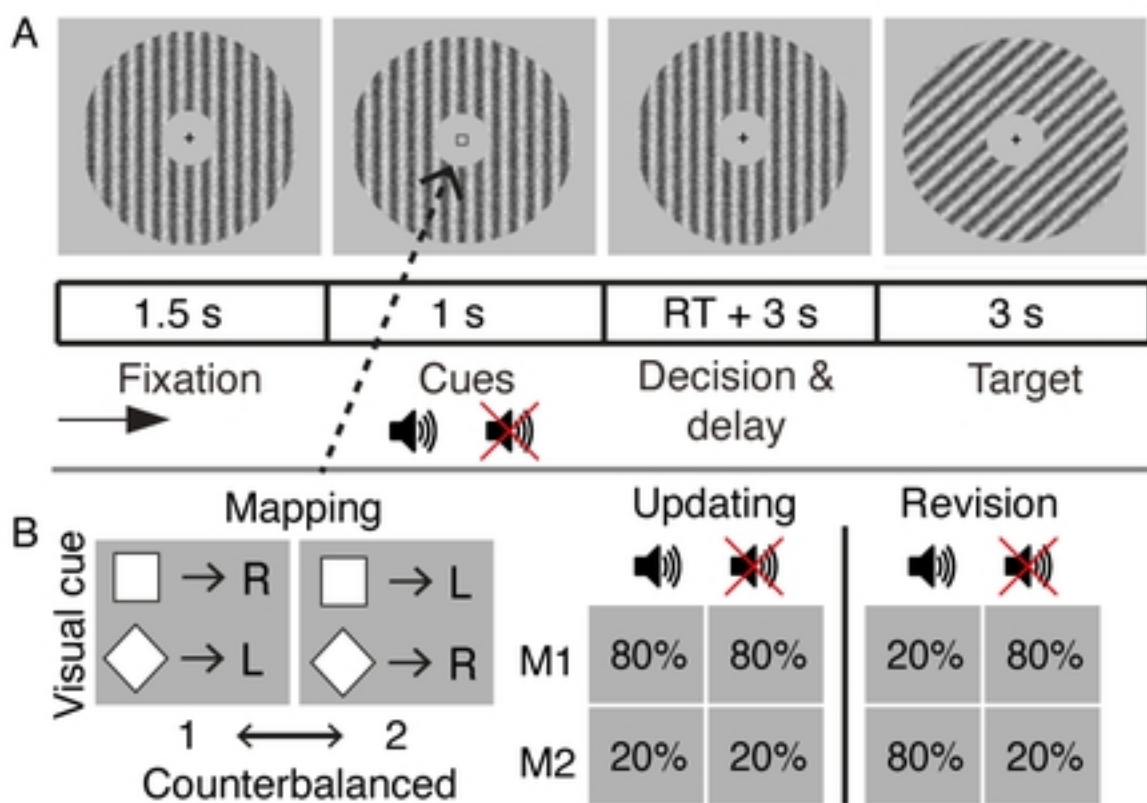
781

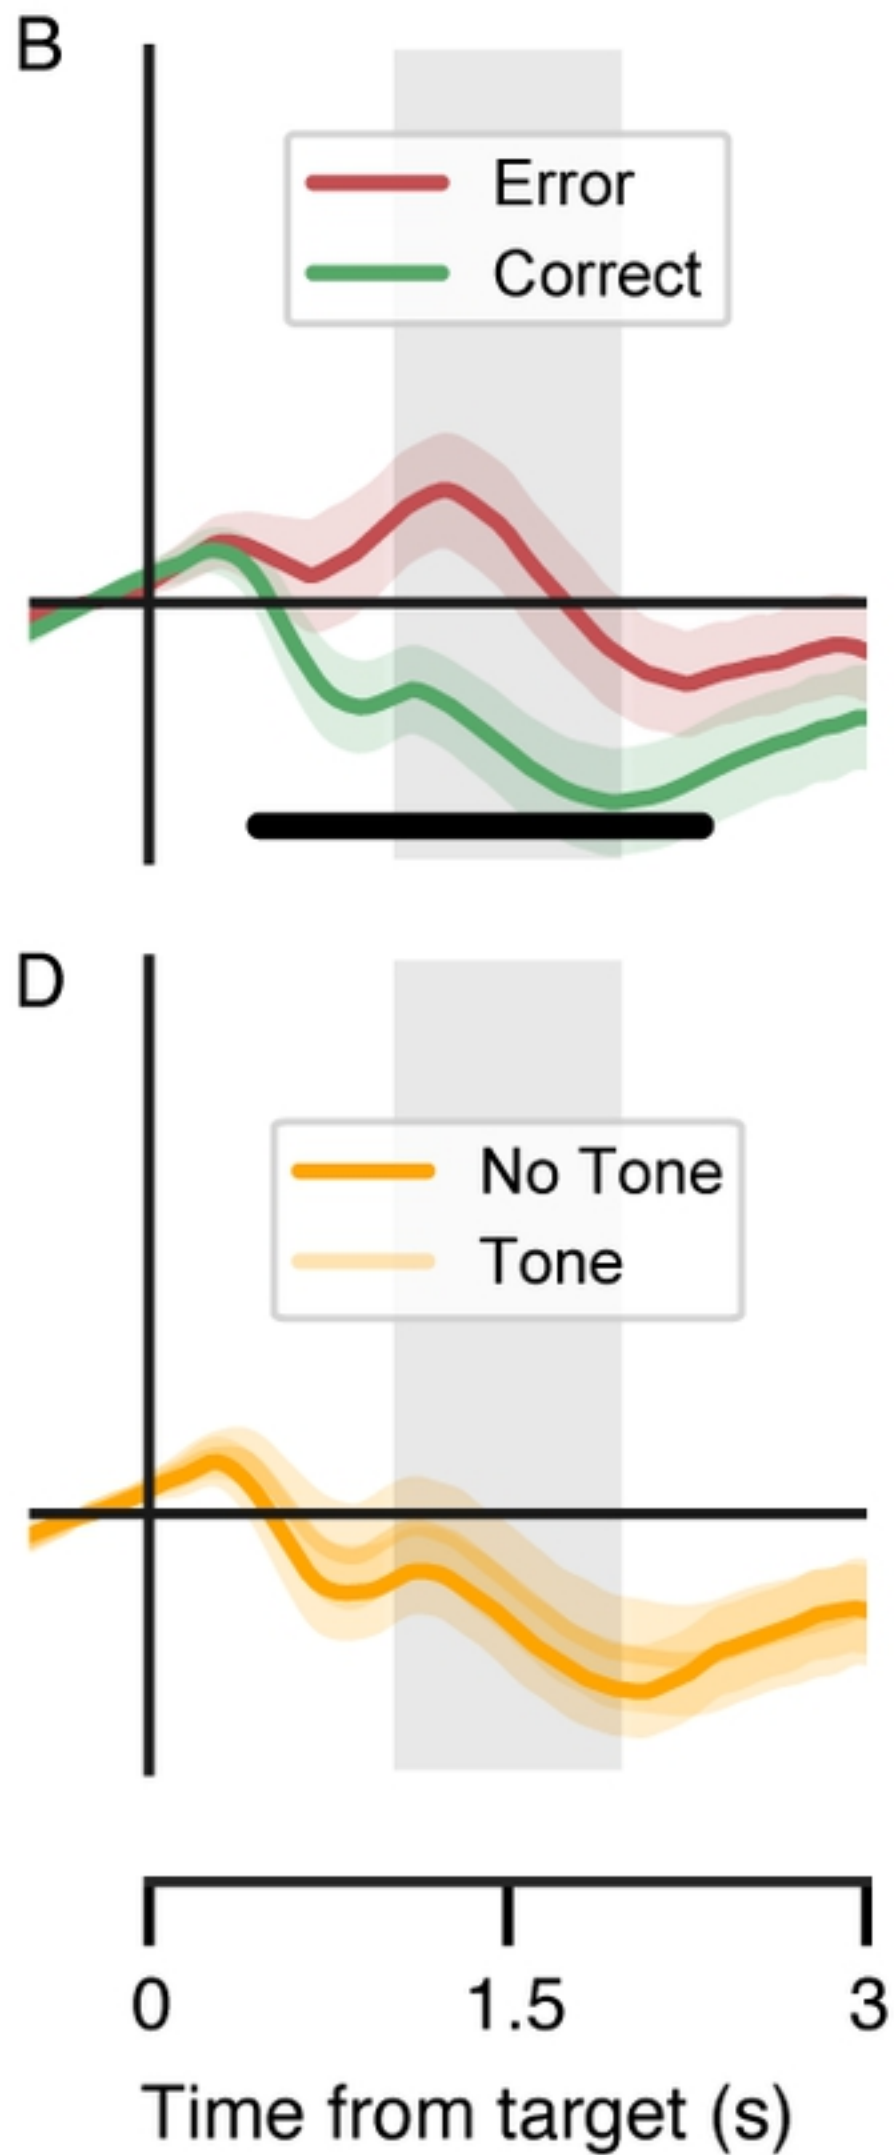
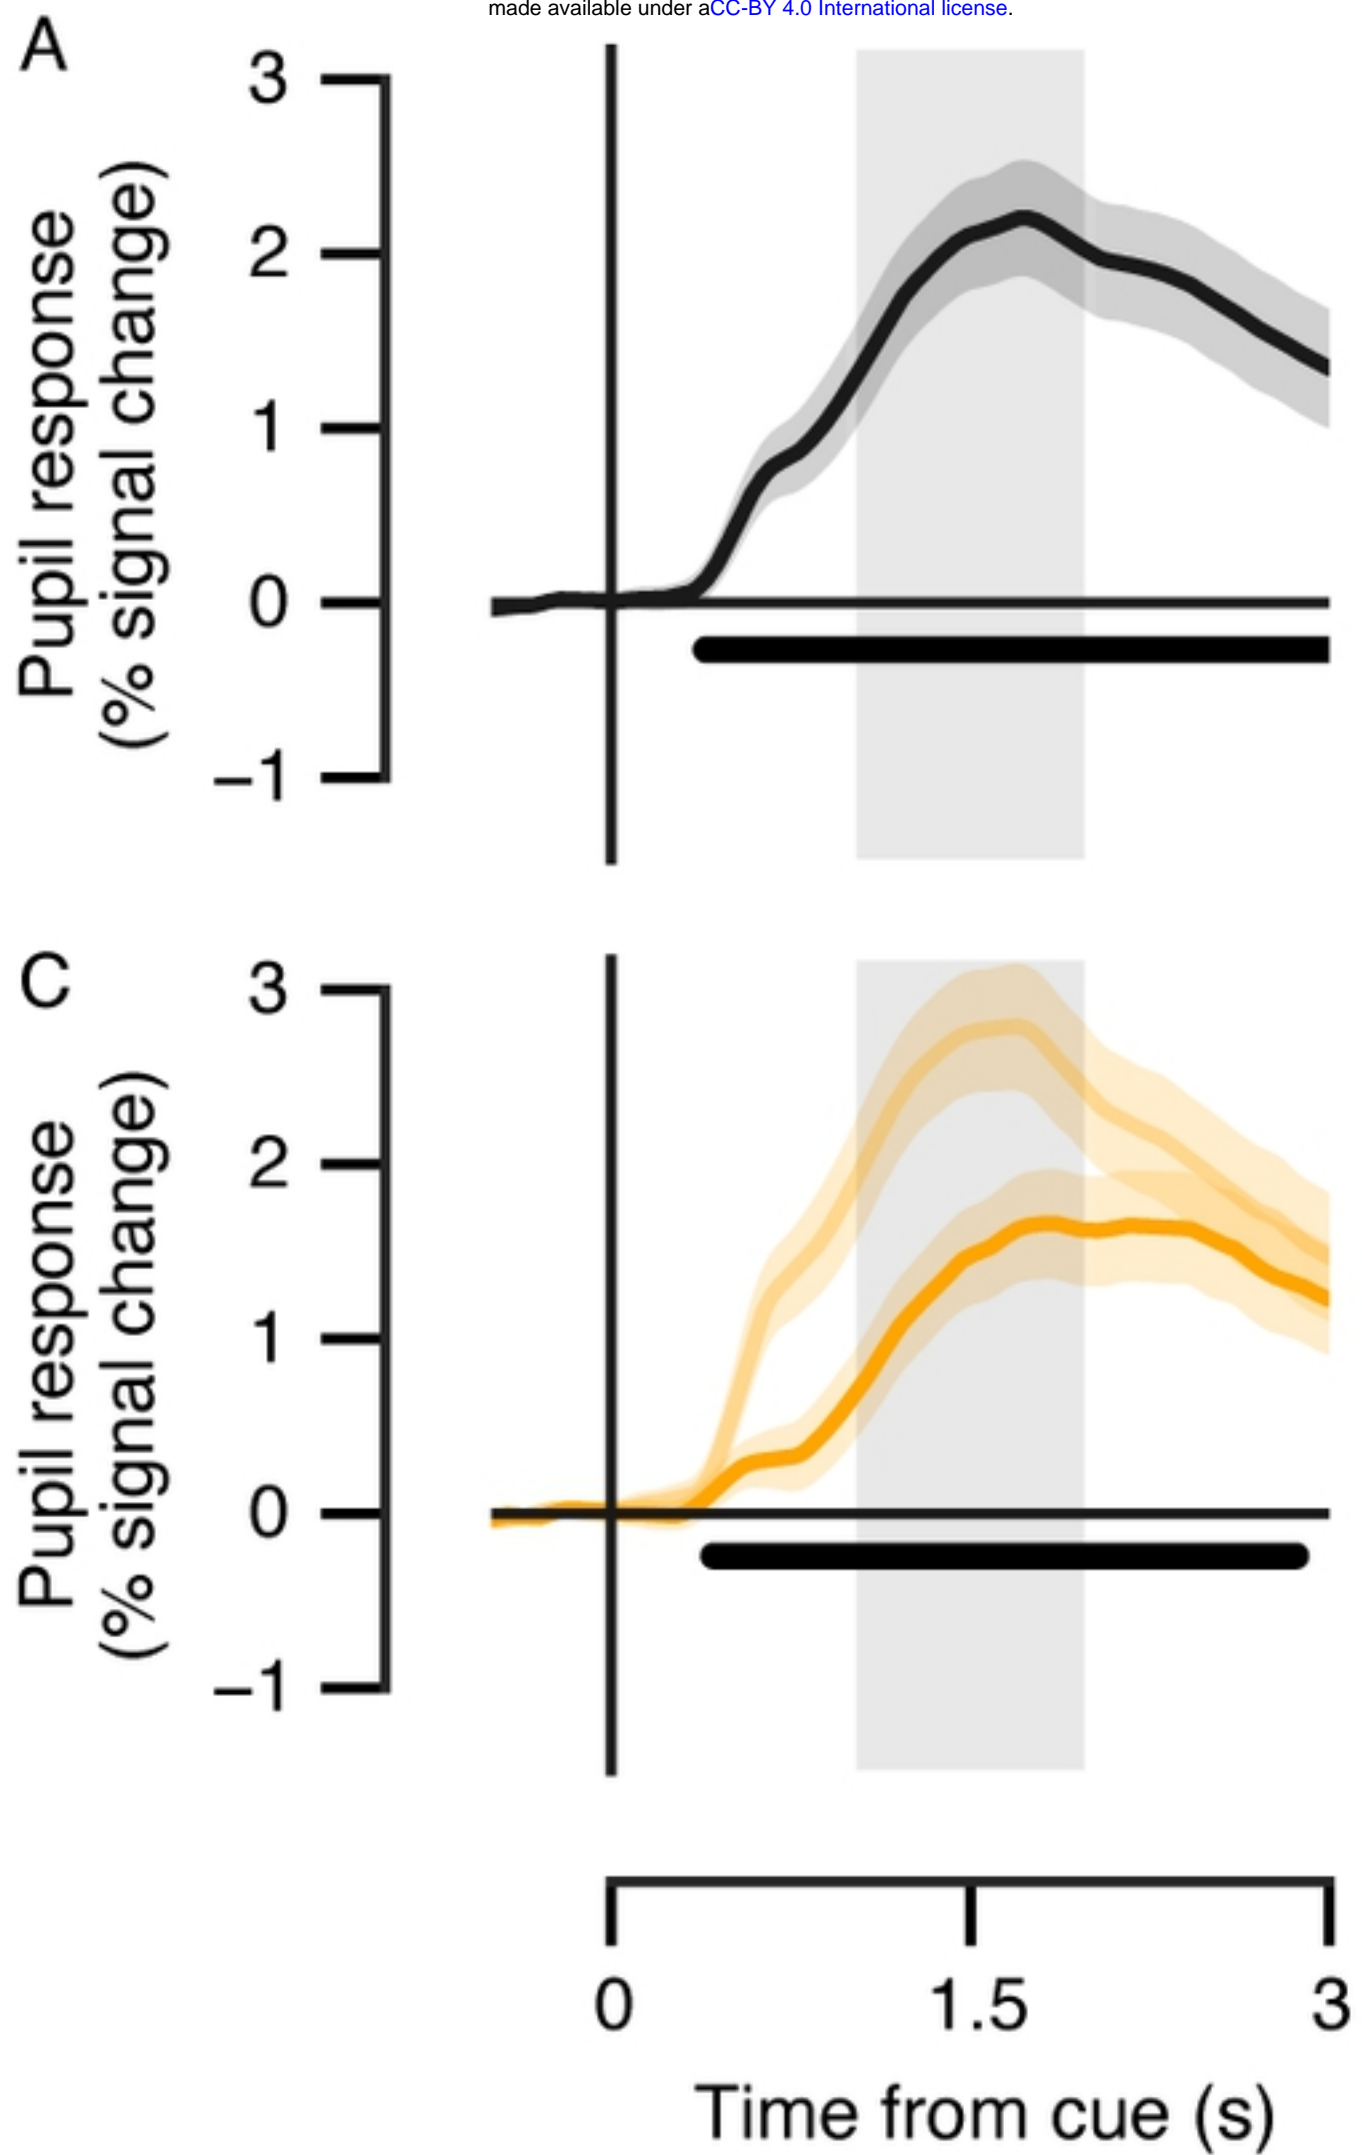
782

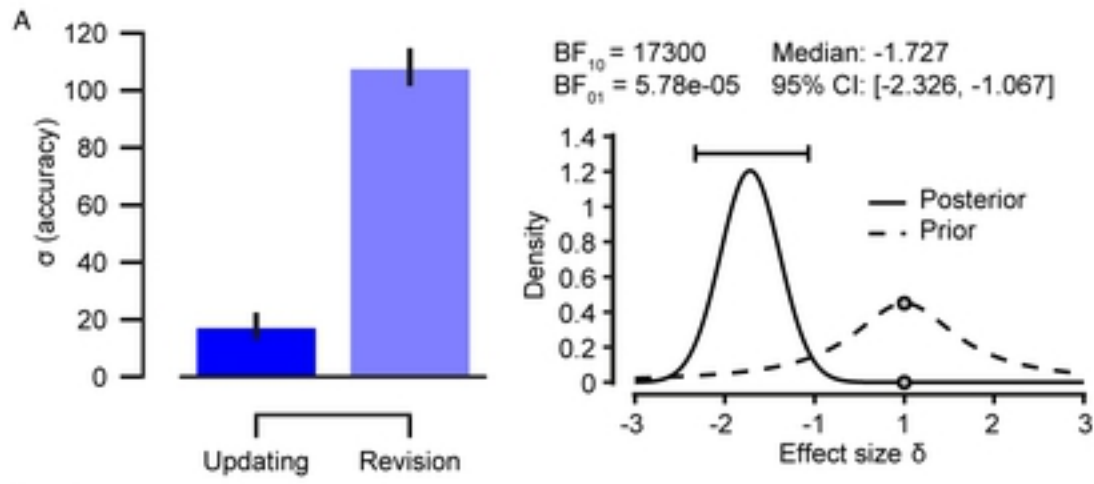
783 **Prediction error in pupil dilation explored as the 4-way interaction between accuracy,**
 784 **cue-target mapping, auditory cue, and phase**

785 If the full scope of the prediction error space was measurable in the target-locked pupil dilation,
786 a 4-way interaction between the factors: accuracy (error vs. correct), and cue-target mapping
787 (M1 vs. M2), auditory cue (tone vs. no tone), and phase (updating vs. revision). For instance, a
788 *correct* response on a *low frequency* (20%) trial may reflect a wrong button press on the part of
789 the participant but will elicit a substantial prediction error due to the unexpected outcome once
790 the contingencies are learned. The size of the two-way interaction term defined by the factors
791 accuracy and cue-target mapping should change over time in accordance with the auditory-cue
792 rules in the updating and revision phases, representing the deviation between expectation and
793 outcome in the form of the target orientation.

794 However, we did not proceed with analysis due to too many missing cases across the 16
795 conditions determined by the 4-way interaction (max. 14/30). These missing cases were due to
796 the rare occurrence of certain conditions, such as correct and low frequency trials. Therefore,
797 we could not make any inference on the potential dynamics of this 4-way interaction in the
798 current task design.







bioRxiv preprint doi: <https://doi.org/10.1101/2022.06.15.496278>; this version posted June 17, 2022. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY 4.0 International license.

