**Predictive eye movements are adjusted in a Bayes-optimal fashion in response to unexpectedly changing environmental probabilities**

**Arthur, T. & Harris, D.J.**

**Abstract**

This study examined the application of active inference to dynamic visuomotor control. Active inference proposes that actions are dynamically planned according to uncertainty about sensory information, prior expectations, and the environment, with motor adjustments serving to minimise future prediction errors. We investigated whether predictive gaze behaviours are indeed adjusted in this Bayes-optimal fashion during a virtual racquetball task. In this task, participants intercepted bouncing balls with varying levels of elasticity, under conditions of higher or lower environmental volatility. Participants' gaze patterns differed between stable and volatile conditions in a manner consistent with generative models of Bayes-optimal behaviour. Partially observable Markov models also revealed an increased rate of associative learning in response to unpredictable shifts in environmental probabilities, although there was no overall effect of volatility on this parameter. Findings extend active inference frameworks into complex and unconstrained visuomotor tasks and present important implications for a neurocomputational understanding of the visual guidance of action.

**Keywords:** active inference; Bayesian; predictive coding; visuomotor; virtual reality

**Predictive eye movements are adjusted in a Bayes-optimal fashion in response to unexpectedly changing environmental probabilities**

## 1. Introduction

The visual guidance of movement involves strategic gaze shifts towards spatial locations that inform future actions (de Brouwer et al., 2021; Land, 2009; Zhao & Warren, 2015). Gaze shifts are partly driven by mental models of the scene and expectations about the location of important information (Henderson, 2017; Itti & Koch, 2001), but inherent processing delays can limit our ability to monitor dynamic and unstable visual cues. To combat these systemic shortfalls, gaze is controlled in an anticipatory manner, based on estimations of the current state and learned properties of the environment (Diaz, Cooper, & Hayhoe, 2013; Hayhoe et al., 2012; Mrotek & Soechting, 2007; Nijhawan, 2008). For instance, when hitting a ball, agents generally execute a saccade to its *predicted* future location, situating their fixation a few degrees above likely bounce positions (Mann et al., 2019). While these anticipatory gaze strategies seem to be pervasive (Diaz, Cooper, & Hayhoe, 2013; Diaz, Cooper, Rothkopf, et al., 2013), it is less clear how they are adjusted when an agent becomes uncertain about their predictions. Picture, for instance, a tennis player trying to return a ball on a rough surface, or a batsman in cricket attempting to play a ball delivered with unknown direction and degree of spin. Is the optimal strategy to stop predicting and rely on online information, or to modify predictions in line with the additional uncertainty and persist with anticipatory eye movements?

Bayesian theories of perception propose that the brain responds to uncertainty by adjusting predictions in a statistically optimal fashion (Knill & Pouget, 2004; Körding & Wolpert, 2004; Rao & Ballard, 1999). Beliefs about hidden states – such as the origin of sensory information or the likely behaviour of the tennis ball – are believed to be the result of integrating top-down expectations with bottom-up sensations (Knill & Pouget, 2004; Körding, 2007; Körding & Wolpert, 2004; Rauss & Pourtois, 2013). The influence of these informational sources is thought to be scaled according to their *precision* (i.e., inverse of the variance) such that bottom-up sensory signals will have a greater impact on posterior beliefs when predictions are weak (Friston, 2005; Shipp et al., 2013). So, when the tennis court is rough (i.e., predictions are uncertain), the player should rely less heavily on their prior expectations about post-bounce trajectory and place added weight on incoming ball motion information.

These Bayes-optimal computations are not only proposed to shape perception (as described by predictive coding; Rao & Ballard, 1999), but also learning and action policy selection (Friston et al., 2016; Parr & Friston, 2019). For both perception and action, agents are said to encode an

internal 'generative' model of the world, which simulates expected sensory data and infers the likely causes of sensations to minimise prediction errors, or *variational free-energy* (VFE; Friston et al., 2006; Friston, 2010; Rao & Ballard, 1999). Active inference extends free energy minimisation to the case of actions, where agents seek to select motor plans (or policies) that minimise future free energy, or *expected free energy* (EFE; Parr & Friston, 2019). EFE represents not only the minimisation of prediction error (i.e., information gain), but also the preference for particular outcomes associated with that action.

In addition to minimising uncertainty via overt motor actions, fixations and saccades can also be conceptualised as individual hypotheses about the state of the world that are aimed at minimising future prediction errors (also known as Bayesian surprise/surprisal; Friston et al., 2012; Itti & Baldi, 2009; Najemnik & Geisler, 2005). Under these assumptions of free energy minimisation, agents should actively 'sample' the world in a way that minimises EFE. While most empirical support for this notion comes from relatively simple perceptual and motor tasks, a small number of previous investigations have identified the importance of predictive cues in the execution of more unconstrained visuomotor skills (Abernethy et al., 2001; Gray & Cañal-Bruland, 2018; Gredin et al., 2018; Stevenson et al., 2009). These studies have suggested that unreliable prior information may be weighted less heavily by agents when executing a motor response (e.g., during interceptive baseball swings; Gray & Cañal-Bruland, 2018). Nevertheless, such processes have not been explicitly modelled as active inference and the above studies have not considered situations in which environmental statistics change dynamically over time. While a tennis player might know that roughly 5% of balls will behave erratically, this degree of uncertainty could also fluctuate as, for instance, the court gets worn or the balls get older.

In fact, previous work has discussed three main types of environmental uncertainty that can affect perception and action (Bland & Schaefer, 2012; Hein et al., 2021; Yu & Dayan, 2005). *Expected uncertainty* refers to the inherent ambiguity derived from probabilistic relationships that exist in the world, such as the outcome of a coin toss. *Estimation uncertainty* emerges from imperfect knowledge of those relationships and diminishes as a result of learning (e.g., repeatedly observing that a coin toss is unbiased). *Unexpected uncertainty* or *volatility* refers to changes in expected uncertainty over time, such as a shift in the probabilistic association between a stimulus and an outcome. Bland and Schaefer (2012) further distinguish between these latter constructs, in the sense that unexpected uncertainty is characterised by rare unforeseen changes in probabilistic relationships, while volatility typifies frequent variations that can, in effect, become expected. Although visually guided movements should theoretically account for such uncertainty and volatility statistics (Arthur et al., 2020; Beesley et al., 2015; Domínguez-Zamora et al., 2018), it remains

unclear how gaze behaviours are adjusted during complex visuomotor skills. Specifically, do agents minimise prediction error in a progressive, Bayes-optimal manner over time? Or do they show abrupt step-changes in visuomotor control under conditions of environmental uncertainty, where gaze is directed to more strategic, non-linear spatial locations?

A further consideration relating to environmental uncertainty is its effect on learning. The rate of associative learning should be enhanced for stimuli whose consequences are uncertain (Dayan & Yu, 2003), as larger and more frequent prediction errors cause the generative model to be revised. Environmental volatility also modulates the rate at which prior models are updated (Behrens et al., 2007). The assimilation of new observations with prior expectations is weighted such that strong priors resulting from lifelong learning (e.g., gravity; Zago et al., 2004)) are not easily modified. Under stable conditions, it would be sub-optimal for a single aberrant event to reshape these predictions. However, under volatile conditions it is necessary for top-down predictions to be more easily modified in light of new observations, resulting in a functionally increased learning rate. Indeed, learning from expected uncertainty and unexpected uncertainty may even be signalled via different neuromodulators, with acetylcholine and norepinephrine performing these two respective roles (Yu & Dayan, 2005). We would, therefore, expect predictive visual behaviours to be updated more rapidly in the context of more unpredictably changeable environmental probabilities. This hypothesis is yet to be empirically examined and is the objective of this paper.

Consequently, the present work sought to understand how unpredictable environmental changes affect visuomotor control during naturalistic and unconstrained interceptive actions. Further, we sought to test whether such changes approximate Bayes-optimal behaviour, as predicted by active inference accounts of perception and action (Friston et al., 2016; Parr & Friston, 2019). To do this, we studied a virtual racquetball task (see Fig 1), in which participants typically display strong prediction-driven gaze behaviours (Arthur et al., 2020; Diaz, Cooper, Rothkopf, et al., 2013; Mann et al., 2019). In line with active inference approaches, it was hypothesised that: i) performers will adjust predictive gaze behaviours between stable and volatile trials in a Bayes-optimal fashion, such that they will place less weight on top-down predictions under volatile conditions; ii) performers will show an adjusted learning rate, such that gaze behaviours will be more strongly influenced by recent context under volatile trial conditions; and iii) environmental shifts that are more unexpected will create a further increase in learning rate (i.e., for unexpected uncertainty compared to volatility; see Bland & Schaefer, 2012).

## 2. Methods

### 2.1. Experimental task and procedures

Behavioural data were primarily collected in the context of understanding how environmental uncertainty and volatility are processed in Autism Spectrum Disorder and a detailed description of all experimental procedures are provided in an accompanying manuscript (Arthur et al., 2020). As a result, no part of the study procedures or analyses were pre-registered prior to the research being conducted. We report how we determined our sample size, all data exclusions, all inclusion/exclusion criteria, all manipulations, and all measures. In short, the interception task took the form of a VR racquetball (squash) game, in which participants had to return a bouncing ball back towards a target on the wall (videos available online: https://osf.io/qjbf2/). The task was developed using the gaming engine Unity (Unity Technologies, San Francisco, CA) and presented to participants via an HTC Vive head-mounted display system (HTC Inc., Taoyuan City, Taiwan; Fig 1). Movements of the headset and hand controller were monitored at 90Hz, based on positional detection in relation to two infra-red 'lighthouse' tracking stations, while gaze was monitored at 120 Hz via an inbuilt Tobii eye-tracking system (spatial accuracy: 0.5°). A virtual racquet was animated based on the movement of the handheld controller, while the simulated court was a 15m square room with an aiming target projected onto the front wall (see Fig 1). Virtual balls visually resembled those used in real tennis, and were launched from just above the aiming target, along the midline of the room (which was 0.75m away from participants on their 'forehand' side). All balls had the same *pre-bounce* flight trajectory and speed, which were both consistent with the effects of gravity (-9.8m/$s^2$). Participants were instructed to hit balls back towards the centre of the target on the front wall.

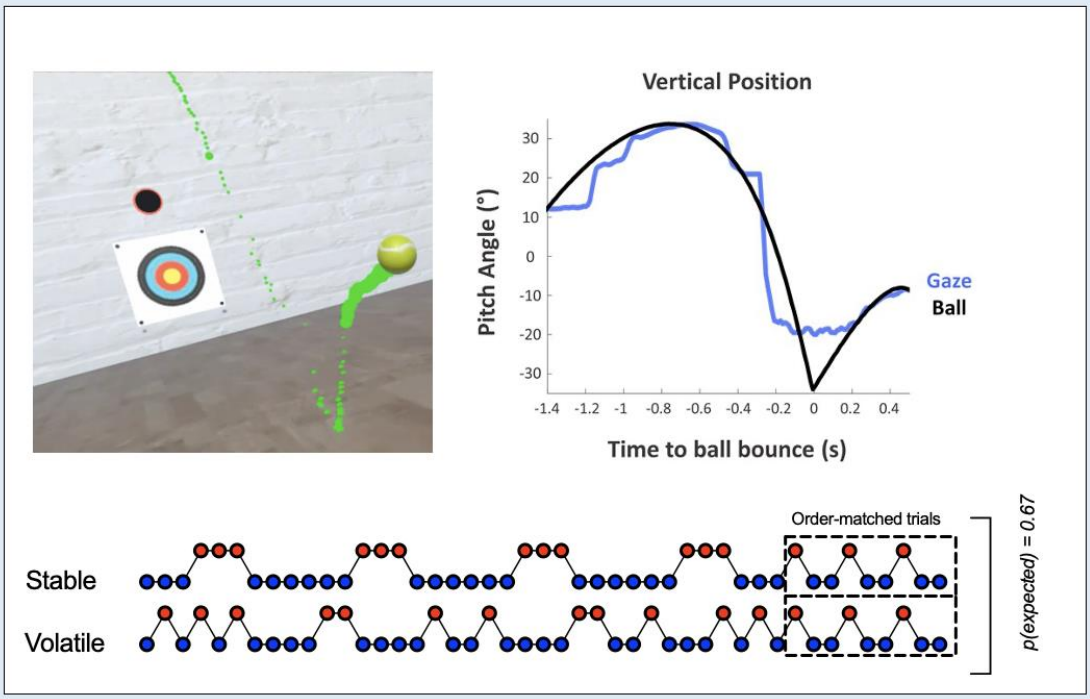


117

118 **Fig 1.** Virtual racquetball task*. Participants intercepted bouncing balls in the virtual environment (top*
119 *left; here shown with gaze point overlaid) using a tracked hand controller while wearing an HTC Vive*
120 *Head-mounted display. The presented balls had identical visual appearance, pre-bounce flight*
121 *trajectory, and speed but were given differing elasticity profiles that either corresponded with real*
122 *tennis balls (i.e., expected), or were unusually bouncy (i.e., unexpected). In the 'stable' condition,*
123 *balls were presented in a predictable serial order (cue-outcome congruency fixed at 66.67%),*
124 *whereas under 'volatile' conditions, cue-outcome probabilities were unpredictably changeable,*
125 *switching irregularly between highly- (83%), moderately- (67%) and non-predictive (50%) blocks*
126 *(bottom). Both conditions ended with nine probability- and order-matched trials. The top right panel*
127 *illustrates typical ball and eye trajectories for a single trial, as has been observed previously* (Diaz,
128 Cooper, & Hayhoe, 2013; Mann et al., 2019)*: after pursuing early-flight trajectory, gaze shifts ahead*
129 *of the ball to a location just above the expected future bounce point (150-190ms), the location of*
130 *which is sensitive to expectations of ball elasticity* (Diaz, Cooper, Rothkopf, et al., 2013)*. After the ball*
131 *has 'caught up', the eyes attempt to track the ball towards the racquet through a combination of*
132 *smooth pursuit and corrective saccadic shifts* (Diaz, Cooper, Rothkopf, et al., 2013; Land & McLeod,
133 2000; Mann et al., 2013; Mann et al., 2019; Mrotek & Soechting, 2007).

Balls presented in each trial were of two possible types – expected and unexpected – which corresponded to two different elasticity profiles. In *expected* trials, ball elasticity was set at 65%, corresponding with the normal behaviour of a real tennis ball. For *unexpected* trials, elasticity was increased to 85%; an easily detectable change in 'bounciness' and post-bounce trajectory (Arthur et al., 2020; Diaz, Cooper, & Hayhoe, 2013). Crucially, the unnatural ball elasticity profile in the unexpected trials was designed such that it deviated substantially from any prior real-world experience of ball bounciness. Since pre-bounce ball trajectories were the same for all trials, ball elasticity information could only be obtained from these distinct post-bounce ball trajectory profiles. Before performing the task, participants received no information about these experimental manipulations of ball elasticity, they were simply informed that the ball would bounce once and that they were free to hit it at any point after this event.

By manipulating the frequency of presentation of the different ball elasticity profiles we created stable and volatile conditions. Under stable conditions balls were presented in a predictable serial order with cue-outcome congruency fixed at 66.67% (i.e., two thirds of balls were expectedly bouncy and one-third unexpectedly bouncy). Under volatile conditions, cue-outcome probabilities were made unpredictably changeable (i.e., unexpected uncertainty) by switching irregularly between highly- (83%), moderately- (67%) and non-predictive (50%) trials in blocks of 6, 9 or 12 (trial order sequences available from https://osf.io/ewnh9/). Crucially, each condition contained an equivalent number of expected (n=30) and unexpected (n=15) trials, ensuring that the marginal probability was identical, and conditions differed only in environmental volatility. Each 45-trial condition took approximately 10 mins to complete and conditions were separated by a short comfort break.

To enable within-condition comparisons of different levels of uncertainty, three expected and three unexpected "test" trials were situated within each block. These trials had identical prior cue-outcome contingencies (66.67%) and identical trial histories (n–1 were all expected trials). Additionally, in order to compare environmental shifts and make learning rate comparisons, the final nine trials in each of the stable and volatile blocks were "order-matched". While these nine trials matched the cue-outcome congruency in the rest of the stable condition, they represented an unexpected shift away from the previously serial trial orders. In the volatile condition, however, they effectively continued both the probability contingencies and volatile presentation order. This allowed us to distinguish unexpected uncertainty from environmental volatility (Bland & Schaefer, 2012).

### 2.2. Participants

57 participants who did not have a history of musculoskeletal or neurological disorders completed the study (34 male, 23 female; mean age: 22.05 ± 3.51 years; 91.23% right-handed). A power analysis indicated that 50 participants were sufficient to detect effects of d=0.7 with 98% power, d=0.5 with 80% power, and d=0.3 with 65% power, given α=.05 in a two-tailed paired t-test (power curves for a range of effect sizes and sample sizes are available online at https://osf.io/9exfk/). Three participants were excluded from the study due to poor quality eye tracking and/or incomplete data, leaving a sample of *n*=54. Prior to completing the stable or volatile conditions, these participants provided written informed consent and were familiarised with the VR procedures. During this time, gaze was calibrated over five virtual locations, a process subsequently repeated upon any obvious displacement of the headset during trials. Participants then began by completing six practice trials on the interception task, before undertaking both experimental conditions in a counterbalanced order. Practice balls were projected from the target without a bounce to ensure that ball elasticity remained unknown. Participants were all naïve to the experimental aims, had no prior experience playing VR-based racquet sports, and received no visual or haptic feedback in relation to racquet-ball contact. The study received approval from the School of Sport and Health Sciences Ethics Committee (University of Exeter, UK) and Department of Psychology Ethics Committee (University of Bath, UK).

### 2.3. Simulation modelling analysis

To simulate statistically-optimal behaviour in this task, we employed a Bayesian generative model of perception derived from the Markov decision process (MDP) formulation of active inference (Da Costa et al., 2020). In general terms, this simulation modelling was used to illustrate the type of belief updating that would occur under Bayesian inference. The model takes a set of initial parameters specified by the experimenter – e.g., prior beliefs about ball bounciness and rate of learning – then iteratively updates beliefs (according to Bayes rule) following predefined observations (observed ball bounces) and determines the action choices a Bayes-optimal agent should make. The POMDP used here solves the otherwise intractable integrals required for model inversion (i.e., moving from prior to posterior beliefs) through estimating posteriors over states via an optimization routine (gradient descent) that seeks to minimise free energy in the model. This is achieved by combining categorical prior expectations about states (**D**) and transitions between states (**B**) with observed instances via the likelihood matrix (**A**), which maps the probability of hidden states given observed instances (see Table 1 for descriptions of model parameters). Consequently, the outputs of the model are estimated beliefs about the bounciness of future balls (i.e., the 'hidden state') and the action choices that agents make based on that belief. For more details on the

mathematical formalism of these models, see Friston et al. (2017), Smith et al. (2020), or an introductory review by Smith and colleagues (2021).

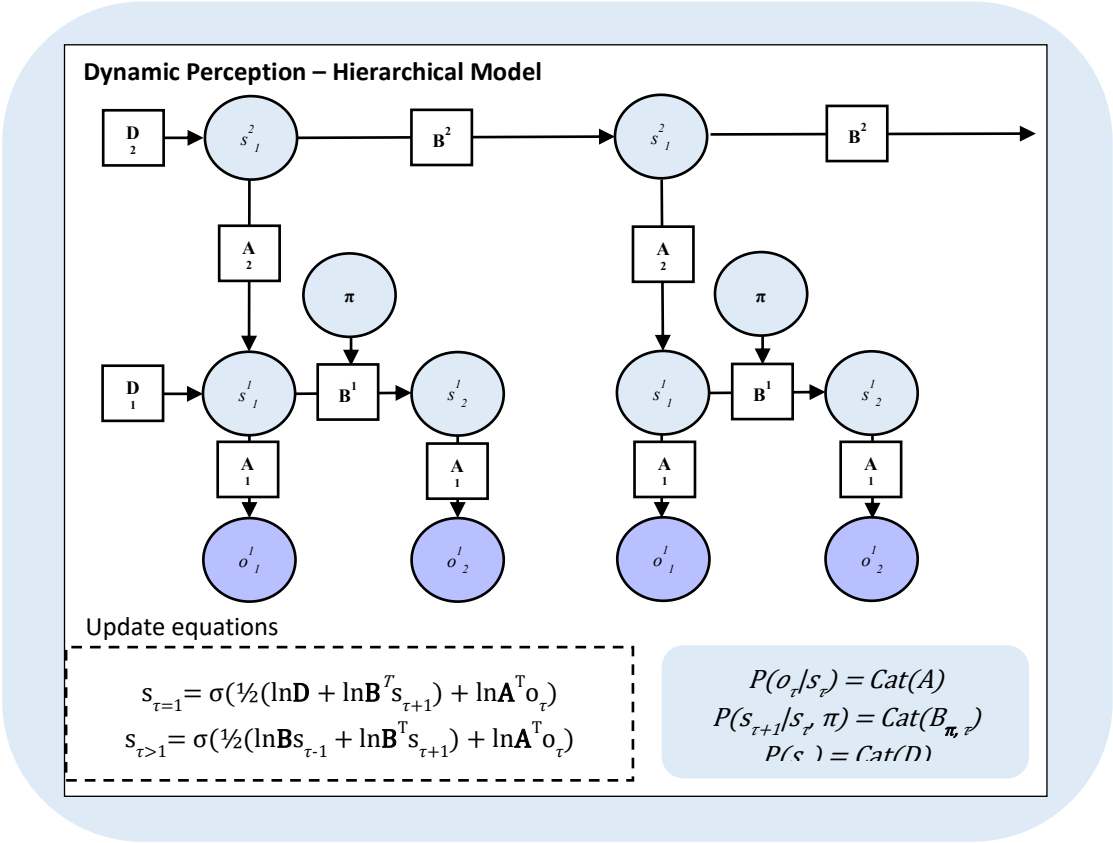


**Fig 2.** Bayesian network representation of two-level POMDP model. *This POMDP model combines categorical prior expectations about states and transitions with observed instances via the likelihood matrix. Circles ('nodes') correspond to variables: s=state, o=observations, $\pi$=action policies. Squares are factors mediating the conditional relationships and take the form of matrices in the model: A=likelihood mapping between states and outcomes (i.e., $(o_\tau|s_\tau)$), D=initial state priors (i.e., $(s_1)$), B=state transition matrices encoding beliefs about how hidden states evolve over time (i.e., $(s_{\tau+1}|s_\tau)$). Observation and state subscripts correspond to time point in a trial ($\tau$). Importantly, when $\tau$>1, the B matrix from $\tau$−1 functions as an empirical prior, playing the same role as the D vector at $\tau$=1. Arrows connecting nodes indicate dependencies between variables. Sigma ($\sigma$) in the update equations refers to a SoftMax function (normalized exponential), which allows vector values to make up a proper probability distribution. The lower level of the model (superscript 1) represents probability updating at the trial level, whereas the higher level (superscript 2) represents learning about trial sequences. Consequently, $D^1$ represents a prior over the probability of normal versus bouncy balls, while $D^2$ represent a prior over the stability of those probabilities over time.*

To specifically index uncertainty and volatility estimations, we computed a generative model with a hierarchical or 'deep temporal' structure. Here, the lower model encodes probability updating during a single trial and from one trial to the next, while the higher model encodes patterns of observations over longer trial sequences. An illustration of this process is provided in Fig 2. Observations (*o*) in the lower level POMDP model were categorical and included a start observation, the appearance of the ball, an expected ball bounce profile, and an unexpectedly high bounce profile. Although participants were required to hit the ball with their racquet, the observation of bounciness was made before the hit, hence the model effectively encodes up until hit point. The lower-level trials formally had three timesteps ($\tau$ = 1 [start], $\tau$ = 2 [ball appearance], and $\tau$ = 3 [bounce observation]) while higher-level trials operated over blocks of nine trials. The agent began in the "start" state and made the associated "start" observation, then either observed a normal bounce, or an unexpectedly high bounce. It was modelled that they then inferred a posterior distribution over states that assigned a probability to the normal/high bounciness state, informed by prior beliefs about the likelihood of observing a normal versus a bouncy ball and the mapping between observations and states. Crucially, these lower-level observations updated priors at the higher level, where expectations of environmental stability were represented. As a result, the higher-level model encodes a form of 'meta-uncertainty' (i.e., volatility) about changing patterns of uncertainty or deviation from an expected order.

These simulations reflect an instrumental use of modelling to generate qualitative predictions about Bayes-optimal behaviour, rather than a veridical representation of psychological processes. Plausible prior models for $\mathrm{p}(\mathrm{expected})$ and $\mathrm{p}(\mathrm{stability})$ were specified, but the true strength of participants' prior expectations is unknown. Lifelong learning about credible ball bounce profiles means that participants likely had a strong prior favouring the expected ball. However, the virtual nature of the task and the experimental setting could mean that despite lifelong learning about ball bounciness, expectations were somewhat weaker (e.g., see Zago et al., 2004). Therefore, a prior was selected to encode a belief that the expected ball was 99 times as likely as the unexpected ball. The distribution over this belief was set to indicate that this belief was fairly weak (i.e., low precision or confidence). Priors for categorical outcomes are represented in the model as Dirichlet distributions, a multivariate generalization of the beta distribution which is defined over a vector of values that sit on the interval [0,1], and sum to one. Therefore a 99 to 1 belief can be represented as a relatively stronger [99,1] or weaker [0.99,0.1] distribution. We specified a prior belief over bounciness of [9.9, 0.1] at the lower level and an even prior belief over stable/volatile [5, 5]. Fifty simulated participants were then modelled for our analyses, with a degree of stochasticity in action selection.

**Table 1.** *Description of computational model elements*

| Model variable | General definition | Model-specific definition |
|---|---|---|
| $\tau$ | Timepoint within a trial | At $\tau$=1 agent was modelled as waiting to observe the ball, at $\tau$=2 the ball appeared, at $\tau$=3 either a normal or bouncy ball was observed and a posterior probability about ball bounciness was inferred. |
| $o_\tau$ | Observable outcomes at time $\tau$ | 1. Start<br>2. Ball appearance<br>3. Normal bounce<br>4. High bounce |
| $s_\tau$ | Hidden states at time $\tau$ | 1. Start<br>2. Ball appearance<br>3. Normal bounciness<br>4. High bounciness |
| **A matrix ($p(o_\tau \mid s_\tau)$)** | Matrix encoding beliefs about the relationship between hidden states and observable outcomes (i.e., the likelihood) | Beliefs about the relationship between the observed bounciness and the hidden state of bounciness. In this instance the observation provides perfect evidence for the state. |
| **B matrix ($p(s_{\tau+1} \mid s_\tau)$)** | Matrix encoding how beliefs about states will evolve over time | Encodes the prior beliefs about whether a normal or bouncy state would occur on each trial |
| **D vector ($p(s_{\tau=1})$)** | Matrix encoding beliefs about initial hidden states | Initial belief |
| $\pi$ | Action policy | Action to anticipate an expected or unexpected ball (i.e., predictive gaze). |
| *Free parameters used during model fitting* | | |
| $\alpha$ | Action precision | Parameter that controls how random action selection is after a policy has been chosen. Higher values indicate deterministic behaviour. |
| $\eta$ | Learning rate (eta) | Parameter from 0-1 that scales the size of the update to beliefs at each time point. |
| **RS** | Risk seeking | Parameter encoding how strong the preference is to predict the correct outcome |
| **LA** | Loss aversion | Parameter encoding how strong the preference is for not predicting the incorrect outcome |

### 2.4. Gaze data analysis

Data extraction and cleaning procedures are described in an accompanying paper (Arthur et al., 2020). Here, participant's cyclopean gaze vector and head position (x,y,z) were recorded from the virtual environment and plotted with respect to 2D direction in space, to provide relative 'in-world' angular orientations (head-ball, gaze-head, and gaze-ball angles). Gaze fixations were detected using a spatial dispersion algorithm (Krassanakis et al., 2014) where fixation events were defined as clusters of successive gaze points within 3° for >100ms, and where gaze velocity was < 30°/s (as in Diaz, Cooper, & Hayhoe, 2013). Saccades were defined as portions of data where gaze acceleration (°/s$^2$) exceeded five times its median absolute acceleration value (as in Mann et al., 2019). To remove any potential artefacts resulting from tracking loss, an additional filter was applied whereby the velocity of saccades had to exceed 40°/s for five consecutive frames and be at least 20% greater than that of the ball. For trials where this automated acceleration criteria did not identify any anticipatory pre-bounce saccades, trials were manually inspected using a 30°/s velocity threshold (Cesqui et al., 2015). Saccade onset and offset times were determined from acceleration minima and maxima (Fooken & Spering, 2020).

The following prediction-related metrics were extracted from this data: the onset of the pre-bounce saccade; the pitch angle of the pre-bounce fixation; and an index of surprise calculated from the unexpected-expected (UE-E) gaze tracking difference. The onset of the pre-bounce saccade was used to index how early the predictive fixation was initiated. Pitch angle indicates the spatial position of the predictive fixation with elevated positions indicating the prediction of bouncier trajectories (Arthur et al., 2020; Diaz, Cooper, & Hayhoe, 2013). Since both of these outcome variables refer to pre-bounce gaze events (that occur before post-bounce sensory information can be obtained), they are driven by an agent's prior expectations about ball elasticity and environmental stability. However, to probe the effects of unexpected outcomes and their levels of associated 'surprisal' we examined the UE-E gaze tracking difference, which was calculated from a z-score of the average post-bounce gaze-ball pitch difference (vertical plane) for each participant. Mean values on expected trials were subtracted from their corresponding unexpected test trial values, such that higher scores indicated a greater difference between expected and unexpected trials (i.e., greater behavioural 'surprise'; Arthur et al., 2020; Baldi & Itti, 2010).

Statistical analysis was conducted using JASP 0.12.1 (JASP team, 2018). Univariate outliers ($p<.001$) for gaze and kinematic variables were identified and removed from the analysis. Four participants with data identified as outliers were excluded from gaze analysis (remaining $n$=50). Poor quality ball tracking was identified in a further two participants at this stage, and UE-E difference scores were consequently removed for these cases. Group comparisons were conducted using

Students t-test, or a Wilcoxon signed rank test in cases where data significantly deviated from normality. Cohen's d was used to quantify effect size for student's t-test, and the rank biserial correlation[1] ($r_{rb}$) for Wilcoxon signed rank test. Conclusions were primarily based on significance values, but Bayes factors ($BF_{10}$) were also calculated to further illustrate the evidence for the alternative versus the null. We used a symmetric Cauchy prior distribution, which was centred on zero with a width parameter of 0.707 (corresponding to an 80% probability that the effect size (*d*) lies between -2 and 2). We follow the convention that $BF_{10}$>3 indicates moderate support for the alternative model and $BF_{10}$>10 indicates strong support (van Doorn et al., 2019). All data from this experiment is freely available and can be accessed from the Open Science Framework (https://osf.io/h5nu7/).

### 2.5. Learning rate analysis

By fitting a POMDP model to real participant data, we were able to identify a set of parameters that enable model predictions to best match observed behaviour (i.e., maximising $(participant\ behavior|model)$). In contrast to the model simulations described above, which predict actions from initial model parameters, here we work backwards from observed behaviours to determine the model parameters that would best explain those actions. In this way, we can estimate the values of pre-specified free parameters (e.g., learning rate) that may vary between participants or conditions. Specifically, we estimated learning rate over the lower level of the model described in Fig 2 which denotes the trial-by-trial perception. While both flat and hierarchical models have been found to explain learning in different tasks (Heilbron & Meyniel, 2019), a simple flat model was chosen here to address the primary question about learning rate while avoiding additional assumptions about hierarchical perception. For instance, unexpected uncertainty represented at a second level could be encoded via tracking unexpected changes in marginal probabilities *or* via beliefs about the wider volatility of the environment (as in our simulations; Fig. 3; Mathys et al., 2014; Meyniel et al., 2015). A single level model negated the need for these additional assumptions which would influence model behaviour. By searching through different possible combinations of parameter values, the best combination can be found for a given participant. Parameter estimation was achieved using variational Bayes (Friston et al., 2007), which works from a set of estimated prior values over parameters and performs gradient descent on VFE in a similar manner to the POMDP (Smith et al., 2021). After obtaining parameter estimates for real participants, we checked the 'parameter recoverability' of the model, i.e., could the model accurately estimate values that were used to generate some artificial data. After simulating data from a range of learning

[1] Glass(1966) recommends that the rank biserial correlation is treated as approximate to Pearson's correlation for the interpretation of effect size.

rate values, we subsequently refit the model to the simulated data to ensure the parameter estimates converged on the known values. Model checks showed good recoverability (plots available in the online files; https://osf.io/h5nu7/).

To obtain the best fit of the model, we enabled four free parameters and compared model fits with different combinations of these parameters. The free parameters were learning rate, sensitivity to reward, sensitivity to loss, and 'action precision', a value which encodes the extent to which action choices are deterministic or random. Successive POMDP models with different free parameter combinations were compared using a Bayesian random effects model (spm_BMS.m function: Statistical Parametric Mapping 12 toolbox; Wellcome Trust Centre for Neuroimaging, London, UK, http://www.fil.ion.ucl.ac.uk/), which assesses the VFE of each model fit and returns the relative probabilities (e.g., [0.9, 0.1]) of the better fit (the protected exceedance probability; Rigoux et al., 2014). As it has been shown that the pitch angle of the bounce fixation is adjusted in line with previously observed ball bounciness (Diaz, Cooper, & Hayhoe, 2013), pitch angle was used to index beliefs about ball bounciness. Lower locations were taken to indicate a belief that $p(\text{expected})$ was more likely. Pitch angle was discretised for modelling purposes: when gaze was shifted to a higher location than on the previous trial (>1SD change) this was taken as a shift towards higher $p(\text{expected})$ and vice versa. To maintain the trial orders, participants with <15% missing values had pitch angle imputed using a linear moving average (median) imputation (Cole, 2008; Moritz & Bartz-Beielstein, 2017). Less than 15% missing values corresponded to >95% imputation efficiency, as recommended by Cole (2008). Participants with >15% missing data were excluded, resulting in 42 datasets for the learning rate analysis.

## 3. Results

### 3.1. Simulation modelling results

The results of POMDP simulations with 50 Bayes-optimal agents are presented in Fig 3. At the lower level of the model there was no difference in free energy between stable and volatile conditions as the marginal probability for $p(\text{expected})$ remained equivalent. However, at the higher level of the model (beliefs about volatility), agents exhibited larger prediction errors in the volatile compared to the stable context and shifted away from a belief in stability. During dynamic visuomotor actions this additional uncertainty should induce a higher learning rate and greater weighting of recent context. If human agents adjust beliefs in a Bayes-optimal fashion and seek to minimise free energy through their visual sampling, as predicted by active inference, a greater weighting of recent context and higher learning rate should be observed in predictive visual behaviours.

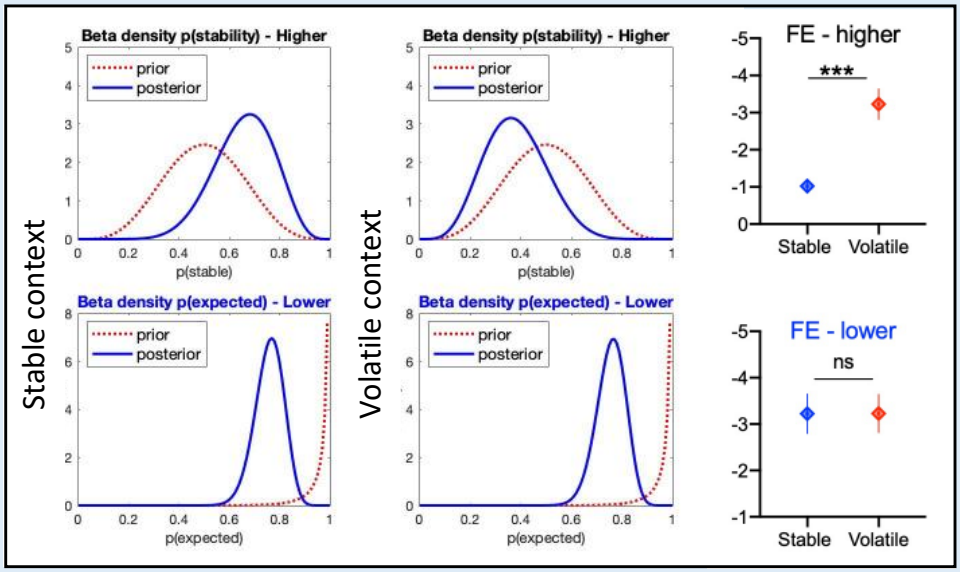


**Fig 3.** POMDP simulation results. *Figures to the left and centre are prior and posterior beta density plots for expectations about context (p(stable) for the higher level of each model) and ball elasticity (p(expected) for the lower level of each model). Plots to the right are total model free energy (means and 95% confidence intervals) over trials at the higher and lower levels for stable contexts (note: more negative indicates greater free energy). Plots indicate group.* ***p<.001, ns=non-significant.

### 3.2. Behavioural results

To investigate whether agents adjust visuomotor control in a Bayes-optimal fashion, we extracted key prediction-related gaze variables from the racquetball dataset. According to the simulated models above, Bayes-optimal agents should place less weight on top-down predictions when forming their beliefs under volatile conditions. When extended to the behaviour of real participants, this could manifest in a later onset of the pre-bounce saccades and/or smaller distinctions in post bounce tracking between expected and unexpected balls (i.e., lower UE-E difference scores). Furthermore, greater prior to posterior shifts in the volatile context (i.e., greater weighting of recent context) should be reflected in higher pre-bounce fixation positions (pitch angles), which are more frequently adjusted (i.e., highly variable) over time.

A Wilcoxon signed-rank test indicated no difference in the timing of the onset of the predictive saccade ($W$ = 606, $p$ = .32, $r_{rb}$ = 0.17), with the Bayes factor favouring the null ($BF_{10}$ = 0.27). However, analysis of the pre-bounce fixation showed that the pitch angle of gaze was different between stable and volatile conditions, in spite of their equivalent cue-outcome probabilities (Fig 4). Specifically, a paired Student's t-test indicated significantly higher pitch averages in the volatile condition ($M$ = -26.6, $SD$ = 5.5) compared to the stable condition ($M$ = -27.5, $SD$ = 4.95; $t$(49) = 2.52, $p$

= .02, *d* = 0.36), although the Bayes Factor was only weakly supportive ($BF_{10}$ = 2.67). This tendency to predictively position gaze at a higher location (under more volatile trials) seemingly affected post-bounce ball tracking responses. Here, Student's t-tests indicated a marginally significantly reduction in E-UE difference for the volatile condition (*M* = -0.30, *SD* = 1.32) compared to the stable condition (*M* = -0.56, *SD* = 1.05; *t*(47) = 2.02, *p* = .049, *d* = 0.29), but the Bayes factor was inconclusive ($BF_{10}$ = 1.01). Supplementary Analysis showed that swing kinematics were also adapted between conditions, with participants generally restricting their range of motion in more volatile trials (Supplementary Fig 1). Therefore, participants appeared to adjust both their *weighting* and *updating* of predictions in a dynamic, Bayes-optimal manner.

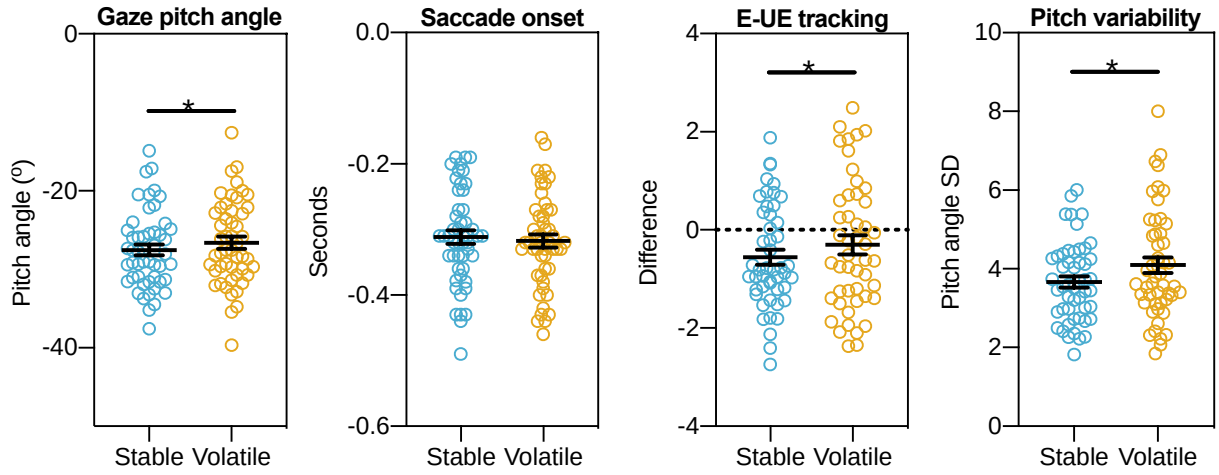


**Fig 4.** Interceptive gaze Behaviours. *Dot plots (with mean and standard error) comparing eye movement variables between stable and volatile conditions.* Note: *p<.05.

However, the degree to which these visuomotor patterns are Bayes-optimal required further scrutiny. It is possible that participants may be responding to environmental volatility using *non-linear* behavioural strategies, as opposed to context-sensitive modulations predicted by active inference. An example of this could be 'centring' strategies, where gaze is positioned mid-way between two outcome possibilities (e.g., see Heinen et al., 2005). Crucially, such a strategy is characterised by a rapid step-change in behaviour, as opposed to a more gradual prior-driven adjustment over time. Hence, in stark contrast to dynamic Bayesian updating (which should be highly variable under volatile conditions), stimulus ordering should have little effect on non-linear behavioural strategies. To rule out a centring strategy, we therefore analysed pitch angle *variability*. A paired t-test indicated significantly lower within-subject standard deviations in position for stable (*M*=3.67, *SD*=1.00) compared to volatile conditions (*M*=4.09, *SD*=1.39; *t*(49)=2.31, *p*=.025, *d*=0.33), although the Bayes Factor provided only weak support ($BF_{10}$=1.72). Therefore, under volatile

conditions, it appears that participants may have made more variable and dynamic adjustments to pitch angle, rather than just maintaining a generic, non-linear behavioural strategy.

### 3.3. Learning rate analysis

Finally, the POMDP model developed during simulations was fitted to the real participant data to estimate parameters for participants' learning rate, prediction errors, and beliefs about ball elasticity. The model was based on gaze pitch angle; a previous indicator of bounciness expectations (Diaz, Cooper, Rothkopf, et al., 2013). This estimation was achieved using variational Bayes (Friston et al., 2007), whereby model parameters are optimised for the behaviour of each individual using gradient descent. A best fitting POMDP model was subsequently identified, which contained free parameters for an overall learning rate estimate ($\eta$), an action precision parameter ($\alpha$), and a parameter encoding loss aversion. This model predicted behaviour to a high degree – the probability of the true action being the one predicted by the model was 0.83. Parameter recoverability (i.e., the ability of the model to accurately estimate, or recover, artificially imputed values) was assessed by simulating data based on known parameters, then re-estimating those parameters from the data. Recoverability was found to be moderate to good, with correlations of r = 0.4, 0.5, and 0.7 between true and estimated parameters for alpha, eta, and loss aversion (for further details on model selection and checks of fit see https://osf.io/h5nu7/).

Firstly, we examined the blocks of stable and volatile trials (excluding the 9 order-matched trials at the end) which varied in order predictability but had equivalent cue-outcome contingencies. Paired t-tests showed that there was no difference in the overall learning rate parameter ($\eta$) between stable ($M$ = 0.52, $SD$ = 0.18) and volatile blocks ($M$ = 0.53, $SD$ = 0.15; $t(41) = -0.36$, $p = .72$, $d$ = 0.06), and the Bayes factor supported the null ($BF_{10} = 0.18$). Nor was there a difference in free energy over likelihood beliefs between stable ($M$ = 0.03, $SD$ = 0.01) and volatile blocks ($M$ = 0.03, $SD$ = 0.01; $t(41) = -0.43$, $p = .67$, $d$ = 0.07, $BF_{10} = 0.18$). There was also no difference in free energy over transition beliefs between stable ($M$ = 0.01, $SD$ = 0.001) and volatile blocks ($M$ = 0.01, $SD$ = 0.001; $t(41) = -0.02$, $p = .98$, $d$ = 0.004), with the Bayes factor again supporting the null ($BF_{10} = 0.17$).

Next, we modelled the nine order-matched trials that followed stable and volatile blocks. These 'matched' trials were essentially a continuation of the volatile block (i.e., unknown order), but represented a sudden shift from the predictable serial order of the stable condition. Hence they provided a comparison of volatility and unexpected uncertainty (Bland & Schaefer, 2012). Paired t-tests indicated that the estimated learning rate parameter ($\eta$) was significantly higher following the stable condition ($M$ = 0.55, $SD$ = 0.05) than the volatile condition ($M$ = 0.55, $SD$ = 0.04; $W(41) = 507$, $p$ = .02, $r_{rb} = 0.44$), but with a weakly supportive Bayes Factor ($BF_{10} = 1.20$). Free energy over the

likelihood was also higher for post-stable ($M$ = 0.036, $SD$ = 0.001) compared to post-volatile blocks ($M$ = 0.035, $SD$ = 0.001; $t(41) = 3.33$, $p$ = .002, $d$ = 0.51, $BF_{10}$ = 17.44). Likewise, free energy over transition beliefs was significantly higher for post-stable ($M = 1.10*10^{-2}$, $SD = 2.6*10^{-4}$) compared to post-volatile blocks ($M = 1.04*10^{-2}$, $SD = 5.03*10^{-4}$; $W(41) = 645$, $p$=.02, $r_{rb}$ = 0.43, $BF_{10}$ = 18.38), indicating a larger updating of beliefs about likely sequences. Therefore, though we did not identify any overall differences in learning rate between stable and volatile conditions, there appeared to be significant between-condition differences in how gaze responses are updated in response to unexpected probability shifts (see Fig 5).

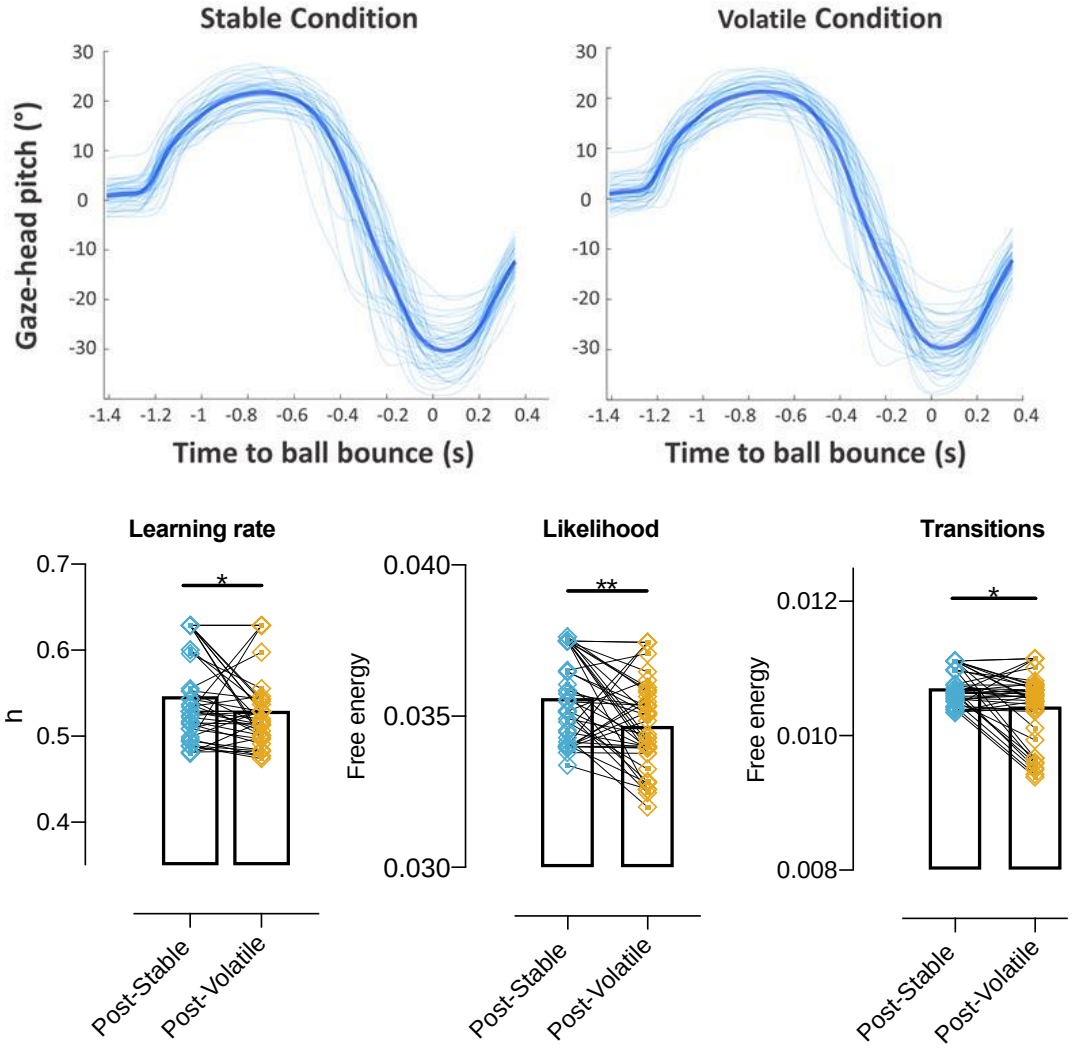


**Fig 5.** *Gaze traces from stable and volatile conditions (top), with dot plots of POMDP parameter estimates showing the differences for post-stable (unexpected uncertainty) and post-volatile (continuing volatility) order-matched trials indicative of reduced surprisal and learning rate.* Note: *p<.05, **p<.01

## 4. Discussion

Predictive coding theories propose the minimisation of prediction error as a biologically plausible governing principle for perception (Friston, 2005; Rao & Ballard, 1999). Active inference extends free energy minimisation to action planning and selection as a mechanism to minimise Bayesian surprisal (Friston et al., 2016; Parr & Friston, 2019). This scheme generates a number of predictions about how the visual guidance of action should be affected by the uncertainty which arises from noisy sensory feedback, imperfect knowledge of the world, and changing environmental states. We examined these predictions of active inference theories in relation to visual guidance and motor execution in a volatile environment. Our findings extend the active inference framework beyond simple tasks and support it as a theoretical foundation for studying unconstrained visuomotor skills.

Hierarchical POMDP simulations of the racquetball task indicated that, for a Bayes-optimal agent, the volatile context should induce elevated prediction errors and a greater reliance on recent context, with posterior beliefs $(\mathrm{p}(\mathrm{expected}))$ that were more weighted towards new observations. No change was observed in the onset of predictive saccades, suggesting that additional uncertainty arising from the volatile context did not affect temporal aspects of anticipatory gaze shifts. However, results showed that subsequent pre-bounce visual fixations were directed to a higher spatial position in the volatile context. This led to a reduction in post-bounce E-UE tracking difference (i.e., dampened surprisal), in a manner that is indicative of weaker prior expectations about 'normal' bounce trajectories. Together, these results suggests that participants used dynamic adjustments to predictive gaze behaviour to minimise tracking error from unexpected balls when they felt less certain about predicting likely outcomes.

In contrast to our predictions, our POMDP model indicated no overall difference in learning rate (indexed by shifts in gaze pitch angle) between the stable and volatile conditions. There was, however, increased *variance* in the pitch angle of the bounce fixation during the volatile condition, suggesting larger and more frequent adjustments of gaze position in the more uncertain context. The POMDP model did reveal differences in learning rate for the nine order-matched trials that followed the stable and volatile block sequences. Specifically, computationally 'surprising' trial order changes in the stable condition led to an increased learning rate and exaggerated updating of beliefs about ball transitions probabilities. Here, recent context was weighted more heavily following an unexpected shift in environmental probabilities (i.e., post-stable), compared to order matched trials that followed a period of already high volatility. These findings further support active inference accounts, and show an important difference in the response to unexpected uncertainty compared to volatility, as proposed by Bland and Schaefer (2012).

Active inference schemes suggest that optimal learning under unexpected uncertainty is achieved through a hierarchical generative model which estimates shifting probabilities (Friston et al., 2016; Meyniel et al., 2015). Our simulations (Fig. 3), for instance, encode dynamic beliefs about environmental volatility (level 2) as well as marginal probabilities of ball bounce outcomes (level 1). However, the exact nature of these context-sensitive higher-level mechanisms, and how they influence learning rate, is unclear. Indeed, there have been proposals that volatility estimations are also hierarchical in nature and that agents encode beliefs about the stability of this parameter (Mathys et al., 2014). This raises the possibility that environmental volatility could itself be considered more or less stable (see supplementary analyses). Moreover, while environmental volatility has been conceptualised as a dynamic modulator of learning (Behrens et al., 2007), changes in learning rate could be achieved through multiple hierarchical beliefs about higher level encoding of volatility, or simple single-level (flat) learning models (Bell et al., 2016; Heilbron & Meyniel, 2019; Meyniel et al., 2016; Wyart & Koechlin, 2016). Although our design did not allow us to disambiguate hierarchical and flat model origins of learning rate changes (although see Heilbron & Meyniel, 2019), the effect of the preceding sequence on learning rate (post-stable versus post-volatile) suggests an awareness of a change point, and therefore appears to support a hierarchical model of learning. Additional simulations using our hierarchical model predicted higher learning rates following a volatile sequence, in contrast to our behavioural results (see supplementary analyses for further discussion). Therefore, future research must establish precisely which higher-level computations underpin active inference behaviours and how these estimates are encoded across dynamic and hierarchical sensorimotor systems.

Our results suggest that not only did people update their expectations about likely ball trajectories in a Bayes-optimal fashion, but that those predictions were reflected in adjusted visuomotor coordination (see supplementary analyses). In addition to the gaze changes discussed above, supplementary analyses highlighted clear changes in swing kinematics between stable and volatile conditions. Specifically, ROM was significantly reduced in more volatile conditions, in a manner that is suggestive of a 'freezing' of degrees of freedom. These kinematic changes seem to reveal a regression towards a simpler movement pattern, whereby joint angles are 'fixed' as an active attempt to minimise movement uncertainty (O'Sullivan et al., 2009). However, findings are incompatible with 'step-changes' in visuomotor behaviour that would suggest deliberate strategic policies. Instead, adjustments in gaze, motor, and learning profiles appeared to follow Bayesian updating principles, whereby actions were progressively adapted in line with environmental probabilities and the frequency with which these statistics changed.

It should be noted that the statistically significant gaze and kinematic effects we report are relatively small (Cohen's *d*=0.3-0.4), particularly when contrasted with the large changes in prior belief predicted by the POMDP simulations. However, visuomotor processes are inherently noisy and impaired by imperfect sensory information (Körding & Wolpert, 2004), unlike the simulations. That Bayes-optimal changes were detected across relatively few trials in noisy visual and motor variables indicates that probabilistic context exerts an important influence on prediction in complex visuomotor behaviours. Consequently, these relatively modest statistical effects could have important theoretical and practical implications. In particular, the ability to make predictions from prior models is fundamental to various high-performance domains, such as elite sport or military combat. While theories of skilled anticipation in interceptive tasks have begun to identify the importance of probabilistic context for action planning (Gray & Cañal-Bruland, 2018; Gredin et al., 2018, 2020; Harris et al., 2021; Loffing & Cañal-Bruland, 2017; Runswick et al., 2020), the present work illustrates how an active inference framework can extend this understanding. Specifically, our findings demonstrate that unexpectedly changing probabilistic contexts will alter anticipatory behaviours and result in greater weighting of recent context. Indeed, strong predictions can become maladaptive in uncertain environments. Of note here then, is evidence that task experts (e.g., in sport) might use prior probabilistic information to greater effect than novices (Gredin et al., 2020; Runswick et al., 2020). Within an active inference framework, extensive task knowledge will inevitably generate stronger priors leading to prediction-driven behaviour. However, future work may wish to investigate whether the use of probabilistic context in real-world anticipation is entirely driven by volume of experience (i.e., more precise prior) or if 'skill' somehow relates to more judicious application of this prior knowledge.

### 4.1. Conclusions

The present work illustrated that predictive gaze behaviours, such as those made to intercept a bouncing ball, are adjusted in a Bayes-optimal fashion in response to unexpectedly changing or volatile conditions. This result extends our neurocomputational understanding of dynamic motor tasks and highlights the potential of an active inference framework for studying visually guided actions. In essence, when faced with unpredictably changing environmental conditions, such as the tennis court becoming rough or a between-set change of balls, agents will adjust their predictions in a statistically optimal fashion. This has important implications for theories of skilled action (Gredin et al., 2020; Runswick et al., 2020), which have considered the influence of probabilistic context (expected uncertainty) but are yet to outline the impact of unexpectedly changing environments (volatility) on predictive visual behaviours.

561 **References**


562 Abernethy, B., Gill, D. P., Parks, S. L., & Packer, S. T. (2001). Expertise and the Perception of
563 Kinematic and Situational Probability Information. *Perception*, *30*(2), 233–252.
564 https://doi.org/10.1068/p2872

565 Arthur, T., Harris, D., Buckingham, G., Brosnan, M., Wilson, M., Williams, G., & Vine, S. (2020).
566 *Expecting the Unexpected: An examination of active inference in autistic adults using*
567 *immersive virtual reality.* PsyArXiv. https://doi.org/10.31234/osf.io/x52ed

568 Baldi, P., & Itti, L. (2010). Of bits and wows: A Bayesian theory of surprise with applications to
569 attention. *Neural Networks*, *23*(5), 649–666. https://doi.org/10.1016/j.neunet.2009.12.007

570 Beesley, T., Nguyen, K. P., Pearson, D., & Pelley, M. E. L. (2015). Uncertainty and predictiveness
571 determine attention to cues during human associative learning. *The Quarterly Journal of*
572 *Experimental Psychology*, *68*(11), 2175–2199.
573 https://doi.org/10.1080/17470218.2015.1009919

574 Behrens, T. E. J., Woolrich, M. W., Walton, M. E., & Rushworth, M. F. S. (2007). Learning the value of
575 information in an uncertain world. *Nature Neuroscience*, *10*(9), 1214–1221.
576 https://doi.org/10.1038/nn1954

577 Bell, A. H., Summerfield, C., Morin, E. L., Malecek, N. J., & Ungerleider, L. G. (2016). Encoding of
578 Stimulus Probability in Macaque Inferior Temporal Cortex. *Current Biology*, *26*(17), 2280–
579 2290. https://doi.org/10.1016/j.cub.2016.07.007

580 Bernstein, N. A. (1967). *The control and regulation of movements.* London: Pergamon Press.

581 Bland, A. R., & Schaefer, A. (2012). Different Varieties of Uncertainty in Human Decision-Making.
582 *Frontiers in Neuroscience*, *6*. https://doi.org/10.3389/fnins.2012.00085

583 Cesqui, B., Mezzetti, M., Lacquaniti, F., & d'Avella, A. (2015). Gaze Behavior in One-Handed Catching
584 and Its Relation with Interceptive Performance: What the Eyes Can't Tell. *PLOS ONE*, *10*(3),
585 e0119445. https://doi.org/10.1371/journal.pone.0119445

Cole, J. C. (2008). *How to deal with missing data: Conceptual overview and details for implementing two modern methods.* In J. W. Osbourne (Ed.), Best practices in quantitative methods (pp. 214–238). Thousand Oaks: Sage Publications.

Da Costa, L., Parr, T., Sajid, N., Veselic, S., Neacsu, V., & Friston, K. (2020). Active inference on discrete state-spaces: A synthesis. *ArXiv:2001.07203 [q-Bio]*. http://arxiv.org/abs/2001.07203

Dayan, P., & Yu, A. J. (2003). Uncertainty and Learning. *IETE Journal of Research*, *49*(2–3), 171–181. https://doi.org/10.1080/03772063.2003.11416335

de Brouwer, A. J., Flanagan, J. R., & Spering, M. (2021). Functional Use of Eye Movements for an Acting System. *Trends in Cognitive Sciences*. https://doi.org/10.1016/j.tics.2020.12.006

Diaz, G., Cooper, J., & Hayhoe, M. (2013). Memory and prediction in natural gaze control. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *368*(1628), 20130064. https://doi.org/10.1098/rstb.2013.0064

Diaz, G., Cooper, J., Rothkopf, C., & Hayhoe, M. (2013). Saccades to future ball location reveal memory-based prediction in a virtual-reality interception task. *Journal of Vision*, *13*(1), 20–20. https://doi.org/10.1167/13.1.20

Domínguez-Zamora, F. J., Gunn, S. M., & Marigold, D. S. (2018). Adaptive Gaze Strategies to Reduce Environmental Uncertainty During a Sequential Visuomotor Behaviour. *Scientific Reports*, *8*(1), 14112. https://doi.org/10.1038/s41598-018-32504-0

Fooken, J., & Spering, M. (2020). Eye movements as a readout of sensorimotor decision processes. *Journal of Neurophysiology*, *123*(4), 1439–1447. https://doi.org/10.1152/jn.00622.2019

Friston, K. (2005). A theory of cortical responses. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *360*(1456), 815–836. https://doi.org/10.1098/rstb.2005.1622

Friston, K. (2010). The free-energy principle: A unified brain theory? *Nature Reviews Neuroscience*, *11*(2), 127–138. https://doi.org/10.1038/nrn2787

611 Friston, K., Adams, R., Perrinet, L., & Breakspear, M. (2012). Perceptions as Hypotheses: Saccades as
612 Experiments. *Frontiers in Psychology*, *3*. https://doi.org/10.3389/fpsyg.2012.00151
613 Friston, K., FitzGerald, T., Rigoli, F., Schwartenbeck, P., O□Doherty, J., & Pezzulo, G. (2016). Active
614 inference and learning. *Neuroscience & Biobehavioral Reviews*, *68*, 862–879.
615 https://doi.org/10.1016/j.neubiorev.2016.06.022
616 Friston, K. J., Parr, T., & de Vries, B. (2017). The graphical brain: Belief propagation and active
617 inference. *Network Neuroscience*, *1*(4), 381–414. https://doi.org/10.1162/NETN_a_00018
618 Friston, K., Kilner, J., & Harrison, L. (2006). A free energy principle for the brain. *Journal of*
619 *Physiology-Paris*, *100*(1), 70–87. https://doi.org/10.1016/j.jphysparis.2006.10.001
620 Friston, K., Mattout, J., Trujillo-Barreto, N., Ashburner, J., & Penny, W. (2007). Variational free energy
621 and the Laplace approximation. *NeuroImage*, *34*(1), 220–234.
622 https://doi.org/10.1016/j.neuroimage.2006.08.035
623 Glass, G. V. (1966). Note on Rank Biserial Correlation. *Educational and Psychological Measurement*,
624 *26*(3), 623–631. https://doi.org/10.1177/001316446602600307
625 Gray, R. (2020). Changes in Movement Coordination Associated With Skill Acquisition in Baseball
626 Batting: Freezing/Freeing Degrees of Freedom and Functional Variability. *Frontiers in*
627 *Psychology*, *11*. https://doi.org/10.3389/fpsyg.2020.01295
628 Gray, R., & Cañal-Bruland, R. (2018). Integrating visual trajectory and probabilistic information in
629 baseball batting. *Psychology of Sport and Exercise*, *36*, 123–131.
630 https://doi.org/10.1016/j.psychsport.2018.02.009
631 Gredin, N. V., Bishop, D. T., Broadbent, D. P., Tucker, A., & Williams, A. M. (2018). Experts integrate
632 explicit contextual priors and environmental information to improve anticipation efficiency.
633 *Journal of Experimental Psychology: Applied*, *24*(4), 509–520.
634 https://doi.org/10.1037/xap0000174
635 Gredin, N. V., Bishop, D. T., Williams, A. M., & Broadbent, D. P. (2020). The use of contextual priors
636 and kinematic information during anticipation in sport: Toward a Bayesian integration

framework. *International Review of Sport and Exercise Psychology*, *0*(0), 1–25. https://doi.org/10.1080/1750984X.2020.1855667

Harris, D., Arthur, T., Broadbent, D., Wilson, M., Vine, S., & Runswick, O. (2021). *An active inference account of skilled anticipation in sport*. PsyArXiv. https://doi.org/10.31234/osf.io/2x3kg

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

Heilbron, M., & Meyniel, F. (2019). Confidence resets reveal hierarchical adaptive learning in humans. *PLOS Computational Biology*, *15*(4), e1006972. https://doi.org/10.1371/journal.pcbi.1006972

Hein, T. P., de Fockert, J., & Ruiz, M. H. (2021). State anxiety biases estimates of uncertainty and impairs reward learning in volatile environments. *NeuroImage*, *224*, 117424. https://doi.org/10.1016/j.neuroimage.2020.117424

Heinen, S. J., Badler, J. B., & Ting, W. (2005). Timing and velocity randomization similarly affect anticipatory pursuit. *Journal of Vision*, *5*(6), 1–1. https://doi.org/10.1167/5.6.1

Henderson, J. M. (2017). Gaze control as prediction. *Trends in Cognitive Sciences*, *21*(1), 15–23. https://doi.org/10.1016/j.tics.2016.11.003

Itti, L., & Baldi, P. (2009). Bayesian surprise attracts human attention. *Vision Research*, *49*(10), 1295–1306. https://doi.org/10.1016/j.visres.2008.09.007

Itti, L., & Koch, C. (2001). Computational modelling of visual attention. *Nature Reviews Neuroscience*, *2*(3), 194–203. https://doi.org/10.1038/35058500

JASP team. (2018). *JASP (version 0.9)[computer software]. Amsterdam, NLD: University of Amsterdam.*

Knill, D. C., & Pouget, A. (2004). The Bayesian brain: The role of uncertainty in neural coding and computation. *Trends in Neurosciences*, *27*(12), 712–719. https://doi.org/10.1016/j.tins.2004.10.007

Körding, K. (2007). Decision Theory: What 'Should' the Nervous System Do? *Science*, *318*(5850), 606–610. https://doi.org/10.1126/science.1142998

Körding, K. P., & Wolpert, D. M. (2004). Bayesian integration in sensorimotor learning. *Nature*, *427*(6971), 244–247. https://doi.org/10.1038/nature02169

Krassanakis, V., Filippakopoulou, V., & Nakos, B. (2014). EyeMMV toolbox: An eye movement post-analysis tool based on a two-step spatial dispersion threshold for fixation identification. *Journal of Eye Movement Research*, *7*(1). https://doi.org/10.16910/jemr.7.1.1

Land, M. F. (2009). Vision, eye movements, and natural behavior. *Visual Neuroscience*, *26*(1), 51–62. https://doi.org/10.1017/S0952523808080899

Land, M. F., & McLeod, P. (2000). From eye movements to actions: How batsmen hit the ball. *Nature Neuroscience*, *3*(12), 1340–1345. https://doi.org/10.1038/81887

Lawson, R. P., Bisby, J., Nord, C. L., Burgess, N., & Rees, G. (2021). The Computational, Pharmacological, and Physiological Determinants of Sensory Learning under Uncertainty. *Current Biology*, *31*(1), 163-172.e4. https://doi.org/10.1016/j.cub.2020.10.043

Loffing, F., & Cañal-Bruland, R. (2017). Anticipation in sport. *Current Opinion in Psychology*, *16*, 6–11. https://doi.org/10.1016/j.copsyc.2017.03.008

Mann, D. L., Nakamoto, H., Logt, N., Sikkink, L., & Brenner, E. (2019). Predictive eye movements when hitting a bouncing ball. *Journal of Vision*, *19*(14), 28–28. https://doi.org/10.1167/19.14.28

Mann, D., Spratford, W., & Abernethy, B. (2013). The Head Tracks and Gaze Predicts: How the World's Best Batters Hit a Ball. *PLOS ONE*, *8*(3), e58289. https://doi.org/10.1371/journal.pone.0058289

Mathys, C. D., Lomakina, E. I., Daunizeau, J., Iglesias, S., Brodersen, K. H., Friston, K. J., & Stephan, K. E. (2014). Uncertainty in perception and the Hierarchical Gaussian Filter. *Frontiers in Human Neuroscience*, *8, 825.*

688 Meyniel, F., Maheu, M., & Dehaene, S. (2016). Human Inferences about Sequences: A Minimal
689 Transition Probability Model. *PLOS Computational Biology*, *12*(12), e1005260.
690 https://doi.org/10.1371/journal.pcbi.1005260
691 Meyniel, F., Schlunegger, D., & Dehaene, S. (2015). The sense of confidence during probabilistic
692 learning: A normative account. *PLOS Computational Biology*, *11(6), e1004305.*
693 Moritz, S., & Bartz-Beielstein, T. (2017). ImputeTS: Time Series Missing Value Imputation in R. *The R*
694 *Journal*, *9*(1), 207.
695 Mrotek, L. A., & Soechting, J. F. (2007). Target Interception: Hand–Eye Coordination and Strategies.
696 *Journal of Neuroscience*, *27*(27), 7297–7309. https://doi.org/10.1523/JNEUROSCI.2046-
697 07.2007
698 Najemnik, J., & Geisler, W. S. (2005). Optimal eye movement strategies in visual search. *Nature*,
699 *434*(7031), 387–391. https://doi.org/10.1038/nature03390
700 Nijhawan, R. (2008). Visual prediction: Psychophysics and neurophysiology of compensation for time
701 delays. *Behavioral and Brain Sciences*, *31*(2), 179–198.
702 https://doi.org/10.1017/S0140525X08003804
703 O'Sullivan, I., Burdet, E., & Diedrichsen, J. (2009). Dissociating Variability and Effort as Determinants
704 of Coordination. *PLOS Computational Biology*, *5*(4), e1000345.
705 https://doi.org/10.1371/journal.pcbi.1000345
706 Parr, T., & Friston, K. J. (2019). Generalised free energy and active inference. *Biological Cybernetics*,
707 *113*(5), 495–513. https://doi.org/10.1007/s00422-019-00805-w
708 Rao, R. P. N., & Ballard, D. H. (1999). Predictive coding in the visual cortex: A functional
709 interpretation of some extra-classical receptive-field effects. *Nature Neuroscience*, *2*(1), 79–
710 87. https://doi.org/10.1038/4580
711 Rauss, K., & Pourtois, G. (2013). What is Bottom-Up and What is Top-Down in Predictive Coding?
712 *Frontiers in Psychology*, *4*. https://doi.org/10.3389/fpsyg.2013.00276

Reid, M., Elliott, B., & Crespo, M. (2013). Mechanics and Learning Practices Associated with the Tennis Forehand: A Review. *Journal of Sports Science & Medicine*, *12*(2), 225–231.

Rigoux, L., Stephan, K. E., Friston, K. J., & Daunizeau, J. (2014). Bayesian model selection for group studies—Revisited. *NeuroImage*, *84*, 971–985. https://doi.org/10.1016/j.neuroimage.2013.08.065

Runswick, O. R., Roca, A., Williams, A. M., & North, J. S. (2020). A Model of Information Use During Anticipation in Striking Sports (MIDASS). *Journal of Expertise*, *3*(4), 197–211.

Shipp, S., Adams, R. A., & Friston, K. J. (2013). Reflections on agranular architecture: Predictive coding in the motor cortex. *Trends in Neurosciences*, *36*(12), 706–716. https://doi.org/10.1016/j.tins.2013.09.004

Smith, R., Friston, K., & Whyte, C. (2021). *A Step-by-Step Tutorial on Active Inference and its Application to Empirical Data*. PsyArXiv. https://doi.org/10.31234/osf.io/b4jm6

Smith, R., Kuplicki, R., Feinstein, J., Forthman, K. L., Stewart, J. L., Paulus, M. P., Investigators, T. 1000, & Khalsa, S. S. (2020). A Bayesian computational model reveals a failure to adapt interoceptive precision estimates across depression, anxiety, eating, and substance use disorders. *PLOS Computational Biology*, *16*(12), e1008484. https://doi.org/10.1371/journal.pcbi.1008484

Stevenson, I. H., Fernandes, H. L., Vilares, I., Wei, K., & Körding, K. P. (2009). Bayesian Integration and Non-Linear Feedback Control in a Full-Body Motor Task. *PLOS Computational Biology*, *5*(12), e1000629. https://doi.org/10.1371/journal.pcbi.1000629

van Doorn, J., van den Bergh, D., Bohm, U., Dablander, F., Derks, K., Draws, T., Evans, N. J., Gronau, Q. F., Hinne, M., Kucharský, Š., Ly, A., Marsman, M., Matzke, D., Raj, A., Sarafoglou, A., Stefan, A., Voelkel, J. G., & Wagenmakers, E.-J. (2019). *The JASP Guidelines for Conducting and Reporting a Bayesian Analysis* [Preprint]. PsyArXiv. https://doi.org/10.31234/osf.io/yqxfr

738 Wyart, V., & Koechlin, E. (2016). Choice variability and suboptimality in uncertain environments.
739 *Current Opinion in Behavioral Sciences*, *11*, 109–115.
740 https://doi.org/10.1016/j.cobeha.2016.07.003
741 Yu, A. J., & Dayan, P. (2005). Uncertainty, Neuromodulation, and Attention. *Neuron*, *46*(4), 681–692.
742 https://doi.org/10.1016/j.neuron.2005.04.026
743 Zago, M., Bosco, G., Maffei, V., Iosa, M., Ivanenko, Y. P., & Lacquaniti, F. (2004). Internal models of
744 target motion: Expected dynamics overrides measured kinematics in timing manual
745 interceptions. *Journal of Neurophysiology*, *91*(4), 1620–1634.
746 https://doi.org/10.1152/jn.00862.2003
747 Zhao, H., & Warren, W. H. (2015). On-line and model-based approaches to the visual control of
748 action. *Vision Research*, *110*, 190–202. https://doi.org/10.1016/j.visres.2014.10.008
749

## Supplementary Materials

### Motor Kinematics Analysis

In our main analyses, we observed that participants' gaze responses were adjusted in a Bayes-optimal fashion. Yet, active inference mechanisms not only relate to how an individual samples the world, they are also said to underpin how an agent acts upon their surrounding environment (e.g., through motor initiation, movement adjustments). Accordingly, we extracted several additional swing kinematic variables from our interceptive racquetball data (collected in the context of examining sensorimotor control in Autism, available at: https://osf.io/ewnh9/), to illustrate whether participants' motor responses were consistent with Bayes-optimal control.

As decision times are modulated according to uncertainty (Lawson et al., 2021), swing onset (the first frame in which forward motion of the racket began) and time to peak swing velocity (the square root of the sum of squared vector differentials) were calculated during the fore swing phase of movement. Here, attenuations in the use of prior knowledge might be reflected in later movement onsets and/or disrupted velocity profiles. Furthermore, swing range of motion (ROM; °) was calculated from the angular deviation of the hand controller in the transverse plane during the foreswing. Reductions in swing ROM values indicate greater 'fixing' or 'freezing' of movement degrees of freedom that are associated with poorer motor control (Bernstein, 1967; Gray, 2020). Under uncertain conditions joint stiffness can be increased and multi-effector redundancy can be restricted to reduce uncertainty (i.e., from signal-dependent noise; O'Sullivan et al., 2009)). Consequently, we hypothesised that ROM would be sensitive to changes in environmental stability, and that performers would show restricted degrees of freedom under more volatile conditions.

Paired t-tests indicated no differences between stable and volatile conditions, for either swing onset times ($t(51) = 0.86$, $p = .39$, $d = 0.12$, $BF_{10} = 0.22$) or time of peak swing velocity ($t(51) = 0.18$, $p = .86$, $d = 0.02$, $BF_{10} = 0.15$). This suggests that the volatile context did not delay the motor action or disrupt hand velocity profiles over time. There were, however, differences in ROM (Supplementary Fig 1). Specifically, paired t-tests indicated a significantly reduced ROM in the volatile ($M = 79.9$, $SD = 27.2$) compared to the stable condition ($M = 83.1$, $SD = 25.6$; $t(51) = 2.74$, $p = .008$, $d = 0.38$), with moderate Bayesian support for our alternative hypotheses ($BF_{10} = 4.32$). These motor patterns are indicative of more novice-like swing kinematics(Bernstein, 1967; Reid et al., 2013), highlighting that a regression towards simpler movement patterns (i.e., 'fixing' of joint angles) may serve as an active attempt for minimising movement uncertainty (O'Sullivan et al., 2009).

Overall, these results support our main observations, that not only did people adjust their expectations about likely ball trajectories in a Bayes-optimal fashion, but that those predictions influenced gaze behaviours *and* motor responses.

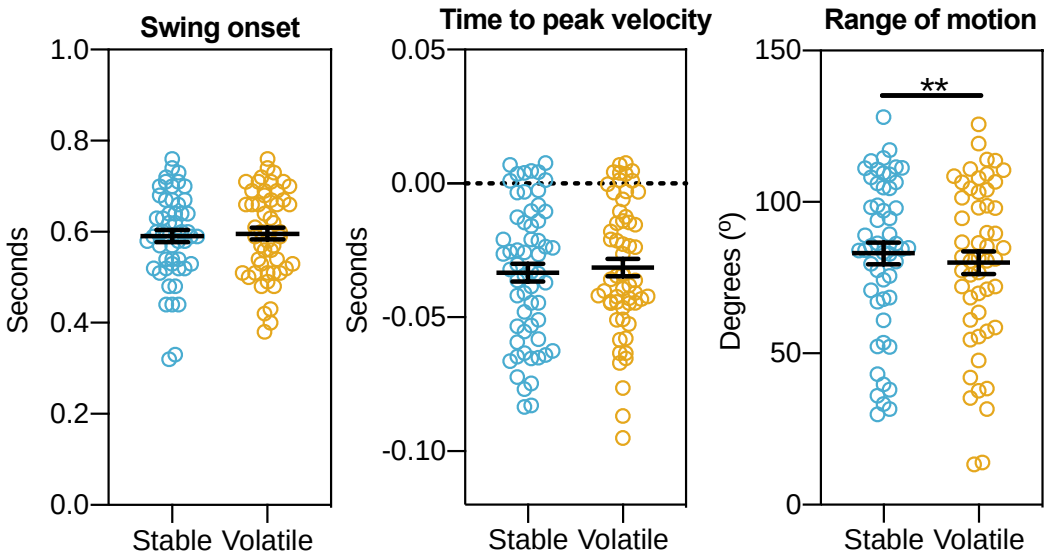


**Supplementary Fig 1.** *Movement variables. Dot plots (with mean and standard error) comparing kinematic variables between stable and volatile conditions. Note: *p<.05, **p<.01.*

**Simulations of model free energy during order-matched trials**

Within an active inference framework, the optimal learning strategy when faced with environmental volatility is to track not only trial-to-trial probabilities but also wider shifts in probabilistic relationships (Meyniel et al., 2015). These dynamic computations generally assume a hierarchical generative model of the world, whereby higher-level beliefs about hidden environmental states modulate lower-level prediction errors (Behrens et al., 2007). However, previous work has found that single-level models, or flat approximations of hierarchical models, can better explain learning in some instances (Bell et al., 2016; Heilbron & Meyniel, 2019; Wyart & Koechlin, 2016). Although a 'flat' model could potentially account for changes in learning rate during volatile trial periods (e.g., due to down-weighting of predictions relative to sensory information when trial to trial contingencies are more uncertain), single-level explanations cannot easily explain why learning rate varies when there is an *awareness* of change (Heilbron & Meyniel, 2019). Consequently, the increase in learning rate for the post-stable trials in our data is generally suggestive of hierarchical rather than single-level learning.

While our design did not allow us to fully disambiguate hierarchical and flat model origins of learning rate (but see Heilbron & Meyniel, 2019), we ran additional simulations using the POMDP

model in Fig. 2 to determine what effect our hierarchical model would predict for the order matched trials that followed stable or volatile periods. This supplementary analysis specifically focused on the size of prediction errors (VFE) in the post-stable and post-volatile order-matched trials. By studying these parameters after a *change point* in our ball bounciness trial sequences, we could understand how volatility computations affect prediction error (and learning rate) in this particular hierarchical system. Simulations showed that prediction error was markedly influenced by the preceding environmental volatility, with VFE profiles clearly contrasting between post-stable and post-volatile blocks (see Supplementary Fig 2).

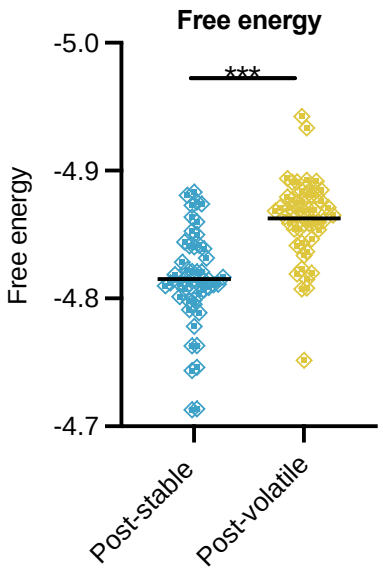


**Supplementary Fig 2.** *Dot plots (with mean) of model VFE for the 9 order-matched trials following periods of stable and volatile trial sequences. ***p<.001*

Interestingly, and in contrast to our behavioural data, these simulations indicated higher prediction errors for the *post-volatile* (*M*=-4.86, *SD*=0.04), as opposed to *post-stable* (*M*=-4.82, *SD*=0.03), trials ($t(49)=7.55$, $p<.001$, $d=1.07$) (see Supplementary Fig 2). The reason for this is that the higher level of our model encoded beliefs about stability/volatility, meaning that agents were encoding the wider unpredictability of their environment in the subsequent order-matched trials. Our behavioural observations of higher learning rates in post-stable conditions (Fig.5) may indicate that participants were actually more sensitive to the abrupt change that followed a period of stability, compared to the persistent changeability of the volatile condition. While often used interchangeably, Bland and Schaefer (2012) distinguish between the concepts of volatility and unexpected uncertainty, where unexpected uncertainty refers to rare unpredicted environmental shifts whereas high volatility is a more frequent fundamental change. Similarly, Mathys et al. (2014) outline how multi-level representations of volatility in the Hierarchical Gaussian Filter model are

827 themselves controlled by volatility parameters, such that the volatility of the environment can be
828 more or less changeable (and so on). The implication of these descriptions is that unpredictable
829 shifts can almost come to be expected. Our behavioural results might, then, be explained by
830 participants' beliefs that the volatile condition was highly unpredictable, and would continue to be
831 so, whereas the shift after the post-stable condition was less expected and worth adapting to.
832 However, this kind of expectation about the persistence of the volatility was not captured in our
833 hierarchical model. As such, it must be stressed that these interpretations remain speculative at this
834 stage and further work should seek to decipher how higher-level expectations of volatility are
835 encoded.