Cognition 210 (2021) 104598

Contents lists available at ScienceDirect

Cognition

journal homepage: www.elsevier.com/locate/cognit

The effect of uncertainty on prediction error in the action perception loop

Kelsey Perrykkad [a,*], Rebecca P. Lawson [b], Sharna Jamadar [c], Jakob Hohwy [a]

[a] Cognition and Philosophy Lab, Philosophy Department, School of Philosophy, Historical and International Studies, Monash University, Clayton, Australia
[b] Department of Psychology, University of Cambridge, United Kingdom
[c] Turner Institute for Brain and Mental Health, Monash University, Clayton, Australia

ARTICLE INFO

Keywords:
Agency
Prediction error
Action-perception loop
Uncertainty
Autism traits

ABSTRACT

Among all their sensations, agents need to distinguish between those caused by themselves and those caused by external causes. The ability to infer agency is particularly challenging under conditions of uncertainty. Within the predictive processing framework, this should happen through active control of prediction error that closes the action-perception loop. Here we use a novel, temporally-sensitive, behavioural proxy for prediction error to show that it is minimised most quickly when volatility is high and when participants report agency, regardless of the accuracy of the judgement. We demonstrate broad effects of uncertainty on accuracy of agency judgements, movement, policy selection, and hypothesis switching. Measuring autism traits, we find differences in policy selection, sensitivity to uncertainty and hypothesis switching despite no difference in overall accuracy.

A significant challenge to an agent's perceptual and decision-making processes is to distinguish between sensations that it can control, and those out of its control. For example, imagine you are working on your computer and it beeps. How do you know if you caused it, as opposed to a colleague emailing you? Influential theoretical work on predictive processing and active inference suggests that the brain relies on prediction errors to assess and test hypotheses about agency (Friston et al., 2013), but empirical evidence for this suggestion is lacking.

Inferring the relations between actions and their sensory consequences is riddled with uncertainty due to the complexities involved in deconstructing sensory evidence from the non-linear confluence of hidden causes. Sometimes when you click, the ensuing beep occurs later because the computer is updating its virus-software; other times, it happens straight away. The brain must represent this uncertainty at numerous hierarchical levels to identify when it is appropriate to attribute agency to oneself. In this example, the breadth of the distribution representing how long it takes for the beep to occur is the *variability* and the frequency of the virus-updates is the *volatility* (how often does the variability distribution change). Crucially, we do not yet know how this uncertainty changes ongoing decisions about *which* actions to perform when trying to explore and infer agency; thus, we have yet to explore how agents close the action-perception loop under uncertainty.

A *judgement of agency* is the verdict that the agent was herself the source of a sensory event – the conscious "I did that" response. It is often (but not always) based on a *sense of agency* (or a feeling of authorship) during the movement. Predictability is often investigated in sense and judgement of agency paradigms by manipulating whether or not the identity (Bednark, Poonian, Palghat, McFadyen, & Cunnington, 2015; Engbert & Wohlschlager, 2007; Hughes, Desantis, & Waszak, 2013; Kuhn et al., 2011; Majchrowicz & Wierzchoń, 2018), timing (Hughes et al., 2013; Majchrowicz & Wierzchoń, 2018) and/or presence (Moore & Haggard, 2008) of a sensory outcome meets some prediction set up by the block-wise probability of each outcome. However, very few studies consider a more continuous distribution of deviations from the expected outcome (e.g., Zalla, Miele, Leboyer, and Metcalfe (2015)) and, to our knowledge, no previous studies have considered volatility (changes to such a distribution) in an agency paradigm.

In classic agency experiments, there are so few actions available to participants that action-selection strategies (or *policies*) cannot easily change in response to changes in prediction error or uncertainty. In some designs, such as Desantis, Hughes, and Waszak (2012), specific actions trigger specific outcomes, but the participants are instructed to equally perform each action. This does not allow participants to explore and attempt to optimally vary policy-selection. In other studies, participants do have freedom to change strategy, and have online action outcome mismatches, but the dependent variables are not sensitive to these strategies and so the temporal dynamics of online decisions with respect to this error are unknown (Zama, Takahashi, & Shimada, 2017). This gap in knowledge is crucial for understanding how we distinguish self-generated and externally-caused sensations in the real world. The

* Corresponding author.
E-mail address: kelsey.perrykkad@monash.edu (K. Perrykkad).

https://doi.org/10.1016/j.cognition.2021.104598
Received 9 July 2020; Received in revised form 7 January 2021; Accepted 8 January 2021
Available online 23 January 2021
0010-0277/© 2021 The Authors. Published by Elsevier B.V. This is an open access article under the CC BY license (http://creativecommons.org/licenses/by/4.0/).

current study sought to close this gap using a novel judgement of agency task that dynamically closed the action-perception loop while independently manipulating variability and volatility.

To understand these missing components in the process of inferring agency, we turn to recent accounts of agency from predictive processing - an explanatory framework whose fundamental claim is that the brain's function is to minimise the long-term average error between its expected and actual sensory input (*prediction error*) (Clark, 2015; Friston, 2010; Hohwy, 2013). In doing so it reduces uncertainty by refining models of the hidden causes of sensory input in the environment and in the agent itself.

Prediction error can be minimised by updating expectations while passively receiving sensory input (*perceptual inference; perception*). Another way to minimise prediction error is through action, by selectively sampling sensory input to satisfy beliefs about sensory input in future states of the world and the agent's own body, given certain actions (*active inference; action*) (Friston, 2017). Previous agency research has focused on perceptual inference in the context of agency, and has not interrogated the ongoing process of active inference.

Under an active inference account, agency attribution would occur by minimising the divergence between the predicted outcomes of available policies for action and the most probable future sensory states; in other words, when there is a belief that goals can be reached from the agent's current state (Friston et al., 2013; Friston, Samothrakis, & Montague, 2012; Hohwy, 2015). Thus, precision (i.e., the inverse of uncertainty) of these inferences is important (Friston et al., 2013) and lead us to investigate the effect of such variability on actions, prediction error and inferred agency.

According to active inference, the very purpose of action is then to minimise expected prediction error. To understand how this plays out in the action-perception loop it is then essential to reveal the interplay between action selection and the magnitude of prediction error at a given time, under a given policy. For the critical case of agency attribution, it is not known how an agent infers policies that may help reduce uncertainty about agency; this is mainly because thus far its magnitude has been under the control of experimenters, not participants themselves. Here, rather than dictating the magnitude of prediction error and measuring effects on behaviour and neural processes, we instead measure the prediction error itself and allow participants to control it with their actions.

The most straightforward expectation for active interrogation in an action-perception loop is that, where possible, policies are inferred which minimise prediction error. Part of the difficulty in testing this prediction is finding an appropriate way to measure prediction error. Here, we operationalise prediction error using eye position to calculate the evolving divergence between hand-movement and stimulus trajectories. Eye-tracking indicates moment-to-moment beliefs about agency which can be tested by mouse-movement. We predict that variability and volatility will have independent effects on movement patterns and policy selection, as well as on prediction error minimisation and subsequent judgements of agency. Specifically, high variability allows less precise representation of control states, which predicts more repetitive policy selections (Perrykkad & Hohwy, 2020a), more prediction error and less accurate judgements of agency. High volatility suggests potentially discoverable interfering hidden causes, predicting more policy exploration and more variance in prediction error which could aid accurate inference of agency. Independent of accuracy, we expect a positive correlation between agency-driven prediction error minimisation and judgements of agency, partly based on active inference theory and partly on prior literature on the role of prediction-expectation mismatch for agency reports.

It is instructive to consider how prediction error minimisation might differ in clinical or subclinical populations because such comparisons help reveal how the prediction error mechanisms work. We focus here on predictive processing accounts of autism, according to which autistic individuals have difficulty abstracting causal rules to higher statistical levels, and thus classify more uncertainty as irreducible. This has been theorised to be due to weaker priors, weightier prediction errors or hyper-flexible estimates of volatility, which all result in a higher learning rate in autism (for review and details, see Palmer, Lawson, & Hohwy, 2017). Manipulation of uncertainty in tasks that rely on perceptual inference has been shown to change performance in autistic populations (Lawson, Mathys, & Rees, 2017). Characteristic differences in action in autism, such as restricted and repetitive behaviours, may indicate differences in active inference in variable environments (Palmer, Lawson, & Hohwy, 2017). Previous research, not framed in terms of predictive processing, has used basic versions of the task we use here, and found no difference between groups of autistic and non-autistic participants (Grainger, Williams, & Lind, 2014; Russell & Hill, 2001; Williams & Happé, 2009), however, we predict the relationship between autism traits and agency attribution should be specific to interactions with uncertainty in the environment as the action-perception loop is dynamically closed (cf. Zalla et al. (2015)). This in turn speaks to underexplored topics in autism research relating to the sense of self and agency (Perrykkad & Hohwy, 2020b). Hence, here we additionally measured autistic traits in our sample and we predict that uncertainty will differently affect policies for movement and prediction error minimisation for participants along this scale.

1. Methods

1.1. Participants

Fifty neurotypical adult participants were recruited. Ten participants were excluded: five participants were removed for technical errors in recording, two for poor quality eye-tracking data (>35% lost trials) and three for poor accuracy (<45%). The final sample of 40 participants were primarily undergraduate students (55%, the remainder had completed tertiary degrees) with an overall mean age of 22.8 years (SD: 3.65, range: 18–34) and included 24 female participants. None of the participants reported neurological conditions, taking medications which affect cognition, nor a history of drug abuse. One participant reported a diagnosis of depression, and one of ADHD, removing these participants did not affect the primary results of interest (see supplementary materials). Two participants reported previously suffering a blow to the head that rendered them unconscious. All participants were fluent in English, had normal or corrected-to-normal vision and 95% were right handed. This study was approved by Monash University Human Research Ethics Committee (Project Number 11396). The experiment was conducted in accordance with the relevant guidelines and regulations, and all participants signed informed consent documents upon commencing the protocol.

1.2. Autism quotient

None of the participants were previously diagnosed with Autism Spectrum Disorder or its nominal variants. All participants completed the Autism Quotient questionnaire (Baron-Cohen, Wheelwright, Skinner, Martin, & Clubley, 2001) to quantify autistic traits. The mean AQ score was 21.43 (SD: 5.89, range: 12–38).

1.3. Experimental task design and procedure

For a schematic diagram of the experimental set up, task and experimental manipulation, see Fig. 1 and a video of the task is available at https://figshare.com/s/fd2742b897e21d901dd0 (DOI: 10.26180/5eabbfb9a8aa4).

Testing was conducted in a quiet, darkened room. Participants were seated at a table with a chin rest set to a comfortable height, 84 cm from the screen, and approximately 55 cm from eye to eye-tracking camera. The task was completed using a computer mouse in the participant's dominant hand which was hidden in a curtained box (base dimensions: 32 cm wide x 30 cm deep). Their opposite hand gave judgement of

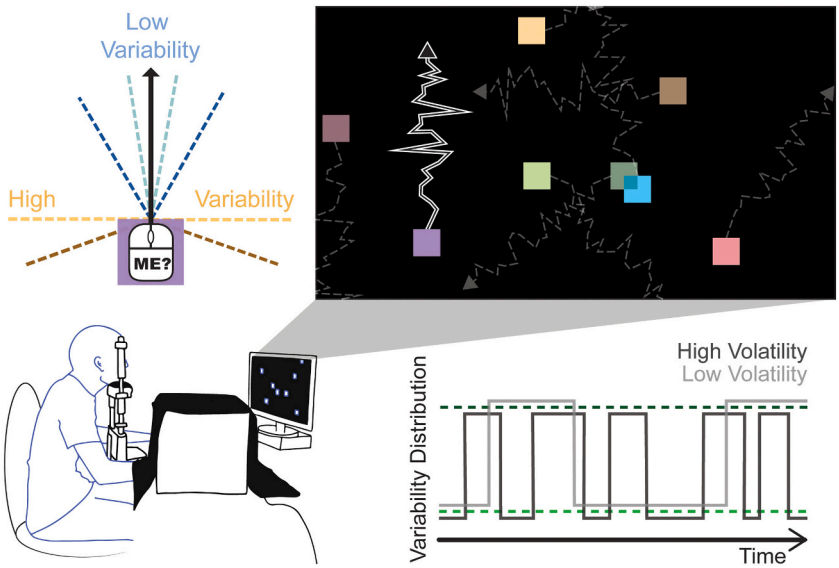

Fig. 1. Task and manipulation.
The participants' task involved using a hidden mouse to control eight squares on the screen. The mean of the target square's movement was the participants' movement, and distractor squares moved at a random angle offset from mouse movement. Jitter was added to the motion of all squares depending on the condition. In *low variability* blocks, the 95% confidence interval from which jitter was sampled switched between ±10° (light blue) and ± 30° (dark blue), and for *high variability,* it switched between ±90° (yellow) and ± 110° (brown) around the input movement or offset. In *low volatility* the distribution changed three random times per trial (light grey), and in *high volatility* it changed ten times (dark grey). See also https://figshare.com/s/fd2742b897e21d901dd0 (DOI: 10.26180/5eabbfb9a8aa4) for a video demonstration. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

agency responses using the numbers on a keyboard. Participants had self-timed breaks between blocks.

We implemented a variant of *the Squares Task* (Grainger et al., 2014; Russell & Hill, 2001; Williams & Happé, 2009), presented using Psychtoolbox-3.0.14 version beta in Matlab 2017b (Mathworks, Natick, Massachusetts) on a 1920 × 1080 screen (60 Hz refresh rate). Eight randomly coloured squares (100px^2) appeared in an array at the beginning of each trial. All the squares moved when the mouse was moved and all the squares stopped when the mouse stopped, so participants had to move in order to accurately complete the task. Participants were given 15 s to identify the target square which they controlled. Distracter squares moved at a random angle offset from the vector of mouse movement, and this angle was also independently and randomly changed (and smoothly transitioned) five times in each trial. This means that each distracter square appeared to turn five times when the participant did not initiate a turn, breaking any illusion of control resulting from motor adaptation. Other than these turns, because the distracter squares took mouse input as part of determining their trajectory (mouse movement + angular offset + variability), the structural features of the motion of the target square (mouse movement + variability) and distracter squares were identical. There were also less frequent *no-control* trials in which all the eight squares were distracter squares. After the 15 s, all squares froze and were numbered, and prompted an unspeeded numerical response from participants indicating which square they controlled or '0' if they thought they controlled none of the squares.

There were four uncertainty conditions in a 2 × 2 design (variability × volatility). Some jitter was added to all squares (*variability*), such that depending on the condition, there was a range (95% CI) of random noise around the mean angle input by the mouse (or the mouse angle + distractor offset for distractors). This specified range also changed throughout the trial; the number of these changes was specified by the *volatility*. In the *low variability* condition, the distribution switched between a 10° and 30° 95% confidence interval on either side of the mean, and for *high variability,* it switched between 90° and 110°. The volatility manipulation decided how frequently the variability changed between these two distributions. In the *low volatility* condition, the variability changed three times, while in *high volatility*, there were 10 changes (pseudo-randomly timed with at least 50 frames between). Each trial's starting distribution was randomly selected. So, for example, in blocks of the *low variability, low volatility* condition, if the distribution started with 10°, the target square was initially jittering within a 95% confidence interval spanning 10° either side of the input mouse angle, and this distribution randomly changed three times during the trial (so widened to ±30°, then narrowed to ±10°, then back to ±30°) with a minimum of 833 ms between these changes in variability (see Fig. 1). There were two blocks of each condition (variability-volatility pair) with 18 trials per block (16 agentive trials, 2 no-control trials) and block order was randomized for each participant.

Prior to completing the task blocks, participants engaged in an interactive instruction demonstration. During the instruction period, participants were given control of a square and the features of the experiment were slowly introduced with text explanations. For instance, on the first page of the instructions, participants were shown one square with no variability added, and the text read: "This is you. Try moving around the screen…". On subsequent screens features such as wrapping around the screen edge, random colours between trials, variability ("Sometimes, the square will not perfectly match your movements, but they are pretty close! Try this…"), distracter squares, and the typical trial structure ("In each trial, you will have 15s to determine which square is you. The colours and starting positions are random, so do not rely on them! You only get information about which one is you by moving around the screen.") including how to respond were slowly introduced. Participants then completed a practice block containing sixteen total trials consisting of all trial types, which was excluded from all analyses. Participants received feedback following practice trials and summative feedback following the practice block. No feedback was given in the main task.

At the end of the experiment there was a short motor control task. In this task, participants were asked to move a perfectly controllable square along a white path as fast and as accurately as possible. There were 10 predesignated paths ranging in length and complexity. This task allowed us to quantify participants' ability to execute motor intentions.

1.4. Analysis

1.4.1. Behaviour

Performance on the motor control task was summarised by multiplying average area traversed outside the white path by average reaction time. This index accounts for the speed-accuracy trade-off, where low scores indicate better motor performance.

In the 'squares' task, *accurate* trials were those in which participants either correctly identified the target square, or correctly identified that there was no such square (*no-control* trials). Accuracy was the primary measure of overall task performance.

The *time spent moving* on each trial was calculated in seconds. This

served as a proxy for environmental sampling, as participants were given freedom to start and stop moving as they pleased though only got task-relevant information by moving.

The *speed* of movement was calculated as the average pixels moved per frame, *acceleration* as change in speed per frame, and *jerk* as the change in acceleration per frame. Derivatives to the level of jerk were analysed to investigate the minimum jerk hypothesis of motor control (Wolpert, 1997) and for its possible relationship to movement trajectories in autism and its traits (Palmer, Paton, Hohwy, & Enticott, 2013; Palmer, Paton, Kirkovski, Enticott, & Hohwy, 2015).

On each frame, the participant's angle of motion was discretised into one of eight cardinal directions. These were plotted for visual inspection. Participants were found to primarily move in the cardinal directions (up, down, left, right), with smaller peaks at the diagonal midpoints. These plots, in combination with observation of trial replays, informed subsequent policy definition. A *turn* was defined as any change in direction which was preceded by at least three frames of one direction and sustained for at least three frames. More than simply sampling, which also occurred in straight movements, turning involves participant induced intervention on expected stimuli direction. These turns were further grouped into types, which were taken to indicate the participant's policy. These are pictorially and algorithmically defined in the Supplementary Materials. In brief, six policy types were identified: 1) Horizontal, 2) Vertical, 3) Perpendicular-Cardinal, 4) Non-Cardinal, 5) Hesitant-straight and 6) Circle. Note that rounded corners and circles were redefined for analysis as one turn each as they are taken as a unified intent of intervention by the participant.

While none of these policies has an a priori advantage over any other for task performance, we were interested in how flexible each participant was in switching between policies. For each policy, we created a mean percentage of turns that were of that type, across all conditions. We defined a participants' *dominant policy* as the policy which had the highest percentage of turns across the entire experiment. This allowed us to look at the number of turns in each trial which were of the participants' dominant policy as compared to alternative policies, as a proxy for exploratory behaviour (i.e., more dominant policy use as exploitative policy selection, less dominant policy use as exploratory). The number of turns on each trial which fell into a participants' dominant policy were used for this analysis.

1.4.2. Eyetracking

Binocular eyetracking data was collected using the SR Research Eyelink 1000 system. For each participant, binocular thirteen-point calibration was conducted; where calibration was unsuccessful using both eyes, one eye was used. The screen x and y coordinates were preprocessed for analysis. Preprocessing involved removing any values outside of the screen bounds, interpolating eyeblinks (as defined by pupil size outside of 1.5 standard deviations below to 2 standard deviations above participant average pupil size), applying a Hanning window of 15 samples (93% overlap) to smooth the eyetrace, and replacing temporarily lost values in one eye with valid data from the other eye (including for whole trials if one eye was excessively noisy). Data was then epoched into trials, and downsampled to match the stimuli framerate for alignment with behavioural data. Trials with poor signal were defined as those with more than 30% of the samples interpolated in both eyes, or whose recorded behavioural data was outside of two standard deviations above or below the participants mean recorded trial length (as the source of these outliers could not be identified). For the final sample of participants, there were a maximum of 65 poor-signal trials (mean = 33.4). Poor-signal trials were removed from all analyses (including behavioural only dependent variables above).

The square the participant was looking at was determined by a novel biased-nearest-object method (see Supplementary Materials), which assumes that at a given moment the participant is looking at one square. While this assumption is nearly always correct, observation of the replay of many trials during development of this method (see also task demonstration video: https://figshare.com/s/fd2742b897e21d901dd0 (DOI: 10.26180/5eabbfb9a8aa4)) showed occasionally participants did not fixate or smoothly pursue one square at a time, instead seemingly relying on peripheral vision. This was likely the participants initially and temporarily tracking multiple squares to narrow down their next hypothesis, consistent with multiple object tracking literature (Fehd & Seiffert, 2008). Times of *hypothesis switch* from one square to another were defined as any change in the looked-at square that lasts longer than one frame.

The Euclidean distance between the expected location (had the stimuli followed the mouse) and the actual location of the hypothesised square was calculated as a proxy for *prediction error*. This means that the prediction error is contingent on how quickly the participants move (the error is higher if they move faster). Due to the manipulation, low variability trials accrue less prediction error on average than high variability trials. The prediction error is also impacted by the magnitude of the distracter's angular offset when the hypothesised square is not the target, so the quality of the hypothesis will affect prediction error. The *average prediction error* for each participant was calculated across each trial. The *slope of prediction error*, representing the rate of prediction error minimisation, was the slope of the line of best fit of the average prediction error at each time point in each condition for each participant (see Fig. 5a and c). As such, negative values here represent prediction error minimisation, while positive values represent accumulating prediction error.

Finally, given the temporal resolution of our prediction error measure, we were interested in the pattern of prediction error around key temporal events – namely hypothesis switches and changes to the variability distribution (due to volatility). We call these analyses *event-related prediction errors* (ERPE). A one second epoch was centred on the event of interest (time zero) and prediction errors were averaged for each participant in each condition to create an average pattern of activity around the event. Means over five 200 ms time bins for each participant were taken for statistical analysis (bin number three is centered on the event onset, see Fig. 6a). There was no effect of time-bin in the volatility ERPE analysis, hence these are reported in the Supplementary Materials.

1.5. Statistics

All statistical analyses were conducted as Mixed Linear Models (MLM) using Jamovi version 1.1.4 and the GAMLj module (Gallucci, 2019; R Core Team, 2018; The Jamovi Project, 2019). Trial-wise data was used for all dependent variables except prediction error slopes and ERPEs for which condition-wise data was used. Variability and volatility were modelled as simple fixed effect factors, and AQ score was modelled as a continuous fixed effect. All interactions between fixed effects were included. By-participant random intercepts were included to address the non-independence of subject-level observations across trials and capture individual variability in task performance. Compared to traditional methods, this approach affords more sophisticated handling of missing and outlying data, thus improving the accuracy, precision, and generalisability of fixed effect estimates (Singmann & Kellen, 2020). See Table 1 for additional covariates for each model. Degrees of freedom are reported as estimated by the Satterthwaite method. *Post-hoc* analyses were conducted with Bonferroni correction for multiple comparisons and post-hoc *p*-values are reported with this correction. For ease of interpretation, post-hoc tests for interactions with AQ were simple effects contrasting participants with three levels of autism traits: low (<Mean-1SD = 16, $n = 6$), within one standard deviation from the mean, and high (>Mean + 1SD = 27, n = 6) scores.

1.6. Data availability

The dataset used for Results, Table 1, and Figs. 2–6 is freely available at https://figshare.com/s/77dececaa2b966db4cf7 (DOI:10.26180/5ed0708f103a2) (Perrykkad, Lawson, Jamadar, & Hohwy, 2020).

Table 1
Significant results summary.

	Dependent Variable	Additional Covariate	M.E. Variability	M.E. Volatility	M.E. AQ	Var*Vol	Var*AQ	Vol*AQ	Var*Vol*AQ
Task	Accuracy		*** low>high			** effect of var. is stronger in high vol			
Movement and Strategy	Time Spent Moving	Time to Movement	*** low<high						*** only for high AQ in high var.: low vol > high vol
	Speed		*** low>high			*** low var. high vol > both low vol & high high; low low>high high			
	Acceleration		*** low<high		* decreasing with AQ				
	Jerk								
	Turning Behaviour		*** low>high			* effect of var. is stronger in low vol			
	Dominant Policy Use	Number of Turns						*** in low AQ low vol < high vol. In high AQ, low vol > high vol	* the low AQ*vol diff. is lost in high variability
	Hypothesis Switches		*** low>high			** in low var. only low vol < high vol			
Prediction Error	Average Prediction Error	Accuracy (and all interactions)	*** low<high				*** var. difference is smaller in high AQ		
	Condition-wise Slope		*** low<high[1]	* low>high					
			M.E. Agency	M.E. Accuracy	M.E. AQ	Accuracy*Agency	Accuracy*AQ	Agency*AQ	Accuracy*Agency*AQ
	Agency-wise Slope		*** not>agent			*** in no agency judgement there is no diff. between correct and incorrect			* High AQ there is no diff. between accuracy; Low AQ when judge agency correct<incorrect, reverse when no-agency and no diff. between agency in incorrect
			M.E. Variability	M.E. Volatility	M.E. AQ	Var*Vol	Var*AQ	Vol*AQ	Var*Vol*AQ
	Volatility ERPE	Average Prediction Error	* low<high						
	Hypothesis Switch ERPE	Average Prediction Error; Hypothesis Switches	*** low>high	* low<high		** only in low var., low vol < high vol			
			…	M.E. Time Bin	AQ*Time Bin	Var*Time Bin	…		
				*** T2 and T3 (event) are>T1, T4, T5 which are equal	*** In T3, AQ is negatively associated with prediction error	*** Only in low var. T2 > T1, bigger var. diff. at T3			

For variables, M.E. = Main Effect, * = Interaction, For results, * = $p \leq 0.05$, ** = $p \leq 0.01$, *** = $p \leq 0.001$ (Post-hoc values Bonferroni corrected for multiple comparisons), Var = Variability, Vol = Volatility, AQ = Autism Quotient, T1-5 = Time bins 1–5, diff. = difference, [1]See Supplementary Materials.

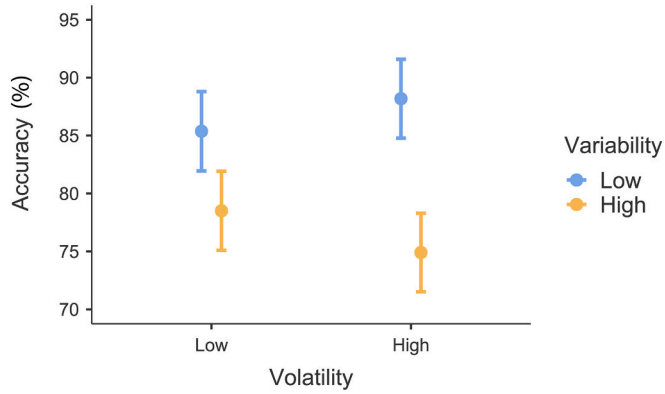

Fig. 2. Accuracy.
Proportion of trials where participants chose the correct square. Participants were more accurate in low variability (blue) than high variability (orange), and this difference was more pronounced under high volatility (right) than low (left). Error bars are 95% CI.

2. Results

In this section, we summarise all statistical models in three sections, first, covering overall task performance, second, the movement and strategy measures, and last, prediction error measures. For each section, we describe the effect of uncertainty on the dependent variables followed by AQ results (though all statistical models included all fixed factors as above). For brevity, we report only main effects of uncertainty in the movement and policy variables. Full statistical reporting is included in Supplementary Materials. See also Table 1 for a summary of all significant results. Performance on the motor control task did not significantly correlate with AQ ($r = 0.07$, $p = 0.65$) or overall accuracy ($r = -0.21$, $p = 0.20$), and so was not included as a random effect in any mixed model.

2.1. Overall task performance: judgement of agency

Average accuracy in the judgement of agency task (Fig. 2) was moderately high across conditions ($\mu = 81.0\%$, $\sigma = 9.12\%$). MLM results show a significant main effect of variability ($F(1,4664) = 85.07$, $p < 0.001$) such that accuracy was approximately 10% higher in the low variability condition than in the high variability condition. Additionally, there was a significant interaction between variability and volatility ($F(1,4665) = 8.62$, $p = 0.003$). Post-hoc analysis revealed significant differences in all comparisons between the four conditions ($z = 4.41$–8.66, $p < 0.001$) except between low and high volatility when variability remained constant (low/low vs low/high $p = 0.421$, high/low vs high/high $p = 0.115$). This result indicates that while volatility does not make a significant difference to accuracy on its own, the effect of variability on accuracy was stronger under high volatility. There was no significant effect of AQ on accuracy.

2.2. Movement characteristics and policy selection

Participants moved for an average of 13.7 s per trial ($\sigma = 1.55$, Range = 3.95–15.1). An MLM comparing the average duration of each trial spent moving across conditions (with the additional fixed effect of time to movement on each trial to account for possible confound) found a significant main effect of variability (Fig. 3a; $F(1,4660) = 727.71$, $p < 0.001$). Participants moved for longer in high variability conditions compared to low variability conditions by an average of 801 ms.

An MLM analysis on average speed of movement revealed a significant main effect of variability (Fig. 3b; $F(1,4661) = 36.42$, $p < 0.001$) such that participants moved faster in the low variability condition compared to high ($z = 6.03$, $p < 0.001$). An MLM on acceleration showed a main effect of variability (Fig. 3c; $F(1,4664) = 12.68$, $p < 0.001$), with faster average acceleration in the high variability trials, compared to low ($z = 3.56$, $p < 0.001$). An MLM on jerk showed no significant results (Fig. 3d).

On average, each trial contained 35 turns (Fig. 3e; $\sigma = 13.9$, Range = 6–107). An MLM on turn count showed a significant main effect of

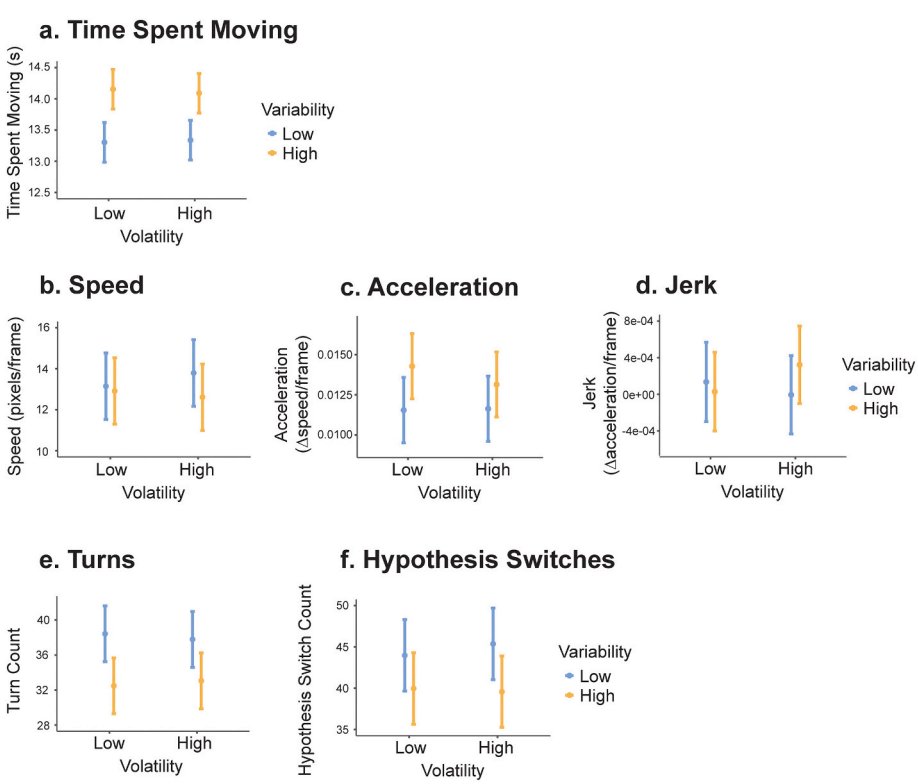

Fig. 3. Movement and strategy.
These graphs depict movement and strategy variables (except dominant policy use, see Fig. 4) across all participants. Volatility is along the x-axis for each graph. Orange bars represent high variability, blue bars represent low variability. Error bars are 95% CI. a) shows mean duration of each trial spent moving, controlling for time to movement onset on each trial; b) shows average speed of movement, c) average acceleration and d) average jerk; e) shows average turn count on each trial; f) shows the average number of hypothesis switches on each trial, when the participant moves their eyes from one square to another. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

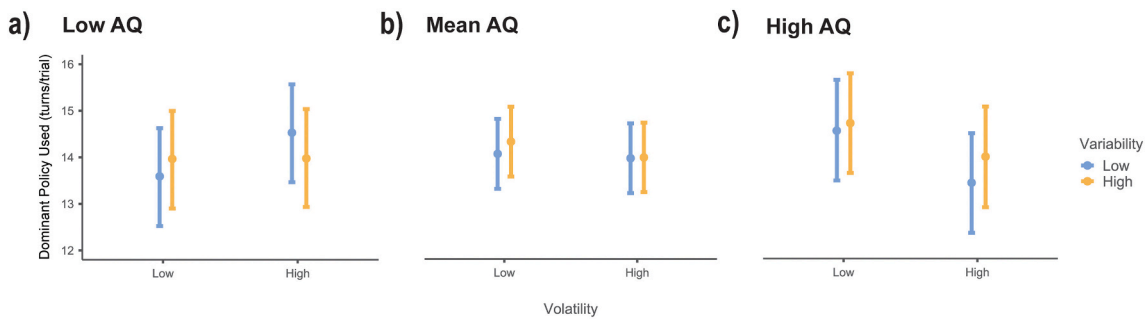

Fig. 4. Dominant policy use.
The turns participants made were categorised into types. This figure shows the number of turns in the participants' own dominant strategy, controlling for total number of turns. For participants with low AQ (<16, panel a), only for low variability trials (blue), participants used their dominant policy more in high volatility (right) than low (left). For participants with AQ scores within one standard deviation of the mean (panel b), there was no difference between volatility conditions (left/right). In both variability conditions (blue and orange), participants with high AQ (>27, panel c) used their dominant policy more in low volatility (left) than high (right). Error bars are 95% CI. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

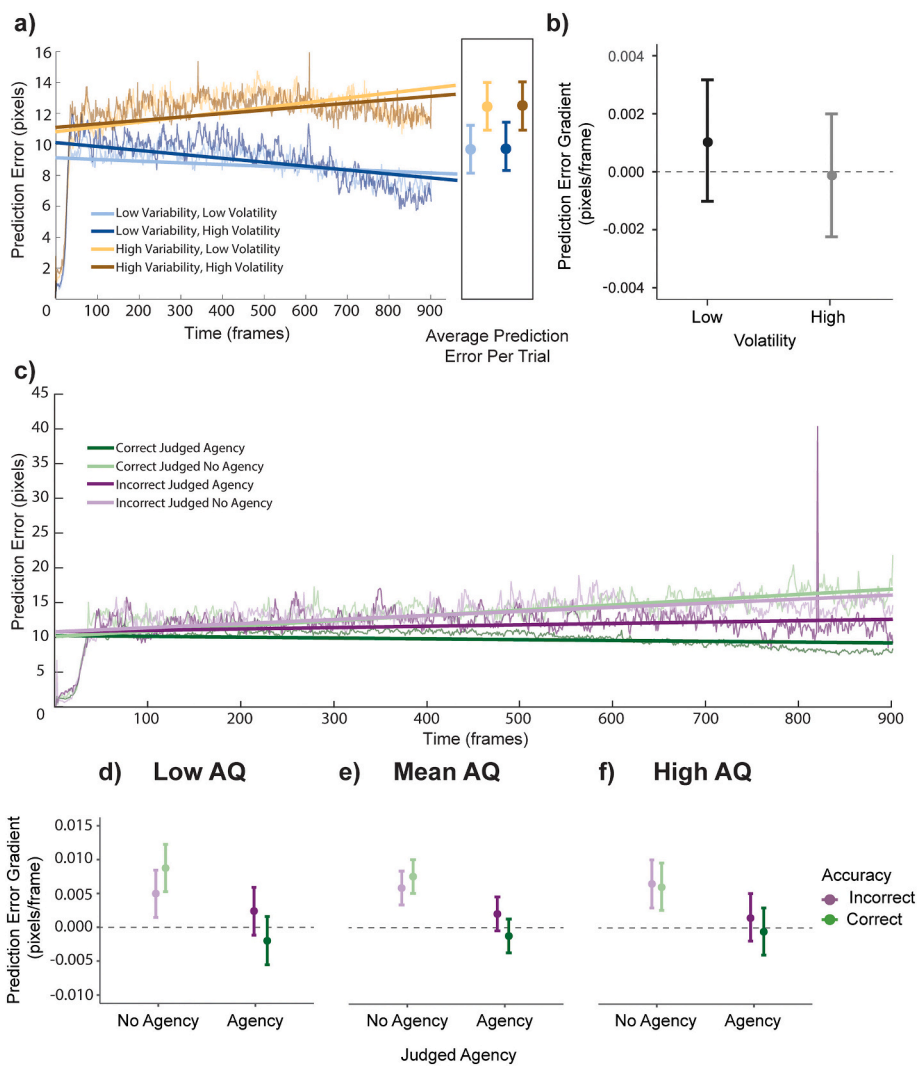

Fig. 5. Prediction error average and gradient.
Panel a) shows the grand average prediction error across the trial split by condition with lines of best fit for each. The box at the end of the graph shows the average prediction error across trials in each condition. Panel b) shows the mean gradient or slope for the lines of best fit for each participant under different levels of volatility. Data used for the box at the end of panel a) is adjusted to account for the influence of accuracy. Panel c) shows the grand average prediction error across the trial split by correct (green) and incorrect (purple) trials and whether the participants chose a square (Judged Agency, dark colours) or said that it was a no-control trial (Judged No Agency, light colours) with lines of best fit for each. Panels d-f show the mean gradient or slope for the lines of best fit for each participant in each combination, split by AQ score. Error bars are 95% CI. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

variability ($F(1,4661) = 346.22$, $p < 0.001$) such that participants turned more frequently in low variability than high variability trials ($z = 18.61$, $p < 0.001$). The dominant turn-types (*policies*) in order of frequency across participants were Non-Cardinal ($n = 22$), Hesitant-straight ($n = 14$), Horizontal ($n = 3$) and Circle ($n = 1$). On average, in each trial, participants used their dominant policy 39.3% of the time ($\sigma = 16.0$, Range = 0–100), and within each participant, the average percent of turns on each trial that were of their dominant policy ranged from 30.1% to 51.9%. For the MLM on dominant policy turn count for each trial, the additional covariate of absolute number of turns on each trial was included to account for this confound. There were no significant main effects.

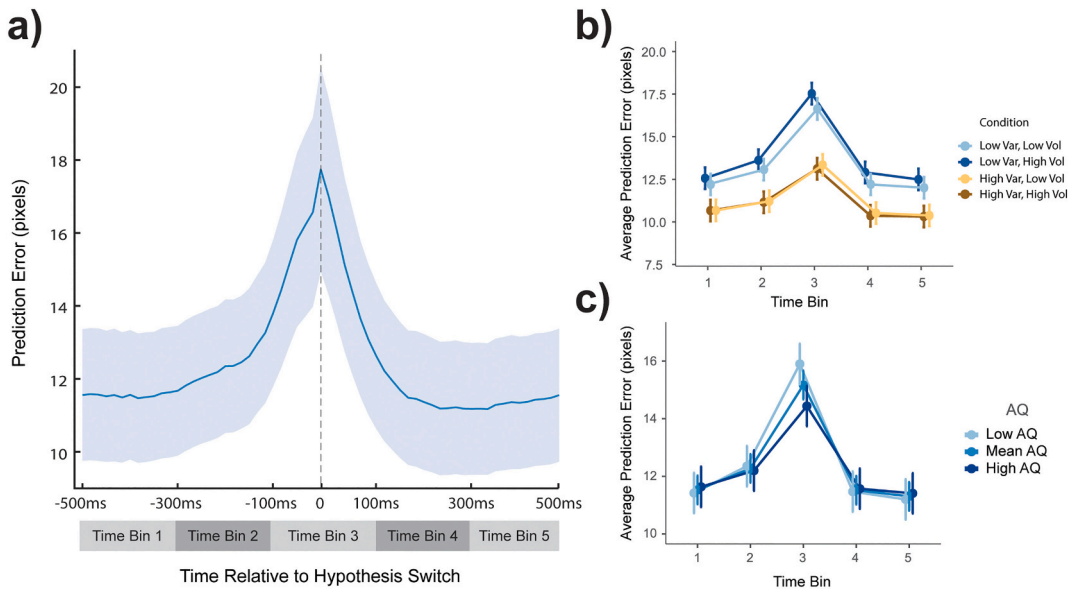

Fig. 6. Hypothesis switch event-related prediction error (ERPE).
Panel a) shows the grand average (blue line) prediction error across participants in a one second epoch centered on hypothesis switches. Time bins used for statistical models are represented in grey shaded bars below. Data used in statistical models, and therefore in panels b) and c) is adjusted for average prediction error differences between conditions and average number of hypothesis switches. Panel b) shows average prediction error in each time bin for each condition. There is more prediction error in this epoch in low variability (blue) conditions than high (orange). This difference is greatest in time bin three, at the time of the event. In low variability (blue), low volatility (light blue) conditions showed less prediction error in this epoch than high (dark blue). Time bin three has the greatest prediction error, followed by time bin two, and none of the others are significantly different from each other. The increase from time bin one to two is only significant in low variability (blue). Panel c) shows the data split by AQ score - lower AQ scores (lightest blue) are associated with greater prediction error at the time of the event (time bin three). Error bars and shading are 95% CI. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

On average participants switched hypotheses 42.2 times per trial (Fig. 3f; $\sigma = 13.48$, Range = 6–134). An MLM on hypothesis switch counts in each trial showed a main effect of variability ($F(1,4661) = 195.91$, $p < 0.001$) such that participants switch hypotheses more when variability is low than when it is high ($z = 14.00$, $p < 0.001$).

These findings suggest that participants' movement was strongly affected by increased environmental variability, causing participants to move more, move slower but accelerate more quickly, and switch hypotheses less often.

2.3. Autism traits and movement and policy

For the dependent variable of time spent moving, there was a significant three-way interaction between AQ, variability and volatility ($F(1,4660) = 11.37$, $p < 0.001$). Post-hoc tests showed that for participants with high AQ only, under high variability only, participants moved for an average of 200 ms longer in low volatility than high volatility conditions. Additionally, the model considering acceleration showed a main effect of AQ ($F(1,38) = 5.73$, $p = 0.022$), such that mean acceleration decreased with increasing AQ ($R^2 = -0.011$, $p < 0.001$). There were no significant findings relating to autism traits across other movement characteristics.

Considering how participants across the AQ range changed their policies in response to uncertainty, the model for dominant policy use showed a significant interaction between AQ and volatility ($F(1,4660) = 19.17$, $p < 0.001$), and a significant three way interaction between AQ, variability and volatility (Fig. 4; $F(1,4661) = 4.27$, $p = 0.039$). Post-hoc analyses for the two-way interaction showed that for low AQ (Fig. 4a) participants used their dominant policy more in the high volatility condition ($z = 2.14$, $p = 0.032$), but only when variability was low ($z = 2.92$, $p = 0.004$), otherwise volatility made no difference ($z = 0.10$, $p = 0.918$). For high AQ (Fig. 4c) participants used their dominant policy more in the low volatility condition ($z = 4.05$, $p < 0.001$), regardless of the variability (high: $z = 2.20$, $p = 0.028$; low: $z = 3.53$, $p < 0.001$).

These findings suggest that different levels of autism traits were associated with differences in the quantity of sampling behaviour, differences in fine-grained movement qualities and differences in the flexibility of policy-selection itself.

2.4. Prediction error

Across all participants, the average calculated prediction error per trial was 10.5 pixels per frame (Fig. 5a; $\sigma = 4.93$, Range = 0.43–50.1). An MLM analysis with the addition of accuracy and all of its interactions with the other fixed factors revealed that, as expected, average prediction error across each trial was significantly associated with the variability condition ($F(1,4653) = 284.05$, $p < 0.001$).

Comparing the slope of prediction error in each condition, an MLM with the addition of accuracy as a random effect revealed a significant main effect of variability ($F(1,114) = 58.15$, $p < 0.001$), which indicated that there was more prediction error minimisation in the low variability condition (lower gradient) than the high (See Fig. 5a; $t = 7.63$, $p < 0.001$). However, this main effect may be explained by a confound of the effect of accuracy which could not be modelled as a fixed effect due to high correlation between the effect of variability and the effect of accuracy on this dependent variable (see Supplementary Materials for a model including accuracy as a random effect in which the significant main effect is removed). There was also a marginally significant main effect of volatility ($F = 3.96$, $p = 0.049$), which showed steeper prediction error minimisation in high volatility compared to low volatility (See Fig. 5a and b, $t = 1.99$, $p = 0.049$).

To investigate the relationship between prediction error minimisation and the participant's judgements, we performed an MLM with

a different structure. For each participant, a linear fit to prediction error across trials with the same accuracy and agency judgement served as the dependent variable. AQ score, accuracy and agency were included as fixed effects, and participant as a random intercept. This MLM showed a main effect of agency ($F(1,113) = 82.89$, $p < 0.001$) and an interaction between agency and accuracy (Fig. 5c; $F(1,113) = 12.79$, $p < 0.001$). Agency was associated with increased prediction error minimisation ($t(113) = 9.10$, $p < 0.001$). Post-hoc tests for the interaction showed that only when participants judge that they did not control any of the stimuli was there no difference in prediction error minimisation between correct and incorrect trials ($t(113) = 1.74$, $p = 0.51$). When participants judge that they did have agency, there is more prediction error minimisation when they are correct than incorrect ($t(113) = 3.31$, $p = 0.007$). However, when considering either only the correct or incorrect trials, prediction error minimisation was steeper when participants judged that they had agency (correct: $t(113) = 9.00$, $p < 0.001$; incorrect: $t(113) = 3.89$, $p = 0.001$;) which confirms that the judgement of agency was associated with steeper prediction error minimisation regardless of accuracy. Numerically, the mean slope of the prediction error was only negative (indicating successful prediction error minimisation) when participants were both accurate and judged that they had agency.

To look at the effect of uncertainty and AQ score on dynamics of prediction error and hypothesis testing, we performed an MLM on the ERPE centered on hypothesis switches. In addition to the standard MLM, we included time-bin as an additional fixed effect of interest and average prediction error and average number of hypothesis switches in each condition as fixed-effect covariates. Fig. 6a shows the timeseries for the average prediction error across conditions and participants in the analysed epoch. There were significant main effects of variability ($F(1,512) = 125.10$, $p < 0.001$), volatility ($F(1,719) = 6.14$, $p = 0.013$) and time-bin ($F(4,719) = 252.29$, $p < 0.001$) and two-way interactions between variability and volatility (Fig. 6b; $F(1,729) = 10.76$, $p = 0.001$) and variability and time-bin ($F(4,719) = 17.94$, $p < 0.001$). Time bins one, four and five were not significantly different from one another ($t(720) = 0.06–1.55$, $p = 1.00$) but the others were all significantly different from one another ($t(720) = 5.24–26.79$, $p < 0.001$), indicating a significant increase before the hypothesis switch starting at least 300 ms before, and a drop after back to the initial level of prediction error. Post-hoc analyses into the main effect of variability showed that low variability conditions had *greater* prediction error around the time of a hypothesis switch than did high variability conditions ($t(519) = 11.09$, $p < 0.001$), which is the inverse of the pattern for average prediction error across the whole trial. Post-hoc analysis of the interaction between time-bin and variability showed that the difference between variability conditions held across all time bins surrounding the hypothesis switch ($t(697) = 6.15—13.81$, $p < 0.001$), but that this difference was greater during time bin three (3.89 pixels, greater than other bin averages by at least 1.69 pixels). Further, only in low variability is there a significant increase from time bin one to two ($t(720) = 6.15$, $p < 0.001$), indicating the increase may occur closer to the event in high variability conditions. While overall, low volatility was associated with less prediction error than high volatility around the time of hypothesis switches ($t(721) = 2.48$, $p = 0.013$), post-hoc analysis of the interaction between variability and volatility showed that this only holds when variability was low ($t(727) = 4.07$, $p < 0.001$). The main effects of this analysis also hold when data is restricted to only incorrect trials, suggesting this result is not driven by an artefact of trial accuracy (see Supplementary Materials for full statistical model and figure).

These findings suggest that increased prediction error minimisation is associated with increased volatility and correctly and positively inferring agency. We have also shown that hypothesis switches function to reduce rising prediction error, and that the dynamics of minimising prediction error in this way is affected by environmental uncertainty at the levels of both variability and volatility.

2.5. Autism traits and prediction error

The model considering the effect of uncertainty and autism traits on average prediction error across a trial showed a significant interaction between variability and AQ ($F(1,4653) = 10.58$, $p = 0.001$). Post-hoc analyses of the variability × AQ interaction showed that the difference between variability conditions decreases as AQ increases (though they are still significantly different across all AQ scores; $z = 9.01–14.44$, $p < 0.001$).

Additionally, the MLM considering agency, accuracy and AQ showed a three-way interaction between these variables ($F(1,113) = 5.69$, $p = 0.02$). Post-hoc analyses showed no difference between agency judgements for incorrect trials for participants with a low AQ score (Fig. 5d; $F(1,113) = 3.51$, $p = 0.064$), but otherwise, when participants judged that they had agency over one of the stimuli, the slope of their prediction error was lower, indicating that they were more effective at minimising prediction error ($t(113) = 9.10$, $p < 0.001$) (both the mean and high AQ groups, and when correct in low AQ). Further, while low AQ participants' prediction error was maximally sensitive to accuracy (lower slopes when correctly judging agency than incorrectly doing so, $F(1,113) = 10.06$, $p = 0.002$; and lower slopes when incorrectly denying agency than when correctly doing so, $F(1,113) = 7.75$, $p = 0.006$); high AQ participants' prediction error was not sensitive to accuracy at all (Fig. 5f; $F(1,113) = 0.11–2.29$, $p = 0.13–0.74$). Participants with a mean AQ showed the appropriate difference only when they judged that they had agency ($F(1,113) = 10.99$, $p = 0.001$).

Looking at the prediction error dynamics limited to the epoch around hypothesis switches showed a significant interaction between AQ and time-bin (Fig. 6c; $F(4,719) = 12.16$, $p < 0.001$). Post-hoc analysis showed a significant difference only in time-bin three (the time of the event) depending on the AQ score ($F(1,50) = 8.58$, $p = 0.005$). A further Pearson's correlation test of AQ by prediction error in this time-bin showed that as AQ increased, the prediction error at the time of a hypothesis switch decreased ($r = -0.21$, $p < 0.001$).

These findings suggest that uncertainty in the environment differentially affects participants' prediction error depending on measured autism traits, including the relationship between prediction error minimisation and judgement of agency, and propensity to switch hypotheses in response to increasing prediction error.

3. Discussion

In this experiment, we closed the action-perception loop to investigate how uncertainty in self-caused sensations influences successive choices about which actions to perform to infer agency. Unlike many previous studies, these actions were freeform and temporally contiguous with ongoing sensory consequences. We showed that action selection changes depending on uncertainty in the mapping between actions and sensory outcomes. We also demonstrate that agency inferences reflect the temporal dynamics of prediction error.

One of the most significant advances of this study on previous designs is the ability to measure and interrogate the temporal dynamics of prediction error, and how this relates to participant behaviour. Using this proxy for prediction error there were particularly interesting findings in the behavioural pattern around hypothesis switches and prediction error minimisation for trials with different judgements of agency. We will now discuss each of these in turn.

Our eye-tracking analysis indicates a hypothesis switch when the participant moves from looking at one square to another and is indicative of a change in the moment to moment beliefs about agency with respect to the candidate square. For action to occur under the active inference account, prediction error comes first, and the action is performed to resolve it. This is consistent with the increasing prediction error leading to a hypothesis switch in our task, indicated by the significant peak in prediction error at the time of the hypothesis switch. The current agential hypothesis is abandoned when the prediction error

is too high – there is decreasing evidence that one can achieve one's expected state with the available actions under the current hypothesis, which leads to a switch that alleviates prediction error. This finding is uniquely consistent with predictive processing (Friston, 2017).

Environmental uncertainty influences this pattern too; after removing trial-wise average prediction error, low variability conditions have a higher prediction error in the hypothesis switch epoch. Also, only in these low variability conditions is there a significant increase from time bin one to time bin two, preceding the switch. Both of these findings suggest that when variability is low, prediction error is allowed to increase for a comparatively longer period of time before the participant decides to switch. This may reflect more reliance on priors in such environments, which allow stable accumulation of evidence for a given hypothesis, and a reluctance to abandon hypotheses in the face of sensory evidence to the contrary.

By looking at the relationship between participants' agency reports and the trend in prediction error over time, our results suggest that participants could be using these trends to inform their judgement of agency. Agency judgements, whether correct or not, were associated with a more negative prediction error slope. Under the predictive processing account, a correct judgement of agency should be associated with a negative trend in prediction error, and a correct judgement of no-agency should not be associated with prediction error minimisation, as the participant cannot effectively control the stimuli to reduce prediction error. These hypotheses were fully borne out for participants with low AQ scores – when participants correctly judged that they had no agency, the slope of the prediction error was more positive (i.e. failed prediction error minimisation) than when they incorrectly said that they had no agency.

Traditionally, internal representations of agency have been explained using a comparator model. In this model, upon movement, the neural system creates an efference copy of motor commands, which predicts "future states of the motor system and the sensory consequences of movement" (Moore & Obhi, 2012)p. 549). This is then compared with incoming sensory information. In both the comparator and predictive processing accounts, agency is associated with small prediction error, or a match between expected and actual outcomes of actions. The comparator however focuses on net retrospective prediction error and cannot account for hypothesis switches in the face of accumulating prediction error or other changes in future action based on inferences of agency (see also Zaadnoordijk, Besold, and Hunnius (2019)). The predictive processing account positions agency in a broader theory of action and policy selection. So, if the projected reliability of policy-outcome mappings over time under a particular hypothesis (occurrent agency) changes, this account is consistent with a threshold in accumulating prediction error after which the agent switches hypotheses and is especially well equipped if this threshold is sensitive to environmental volatility. Our hypothesis switch ERPE suggests that hypothesis switching is sensitive to volatility when variability is low, with more prediction error around a hypothesis switch when volatility is high.

These results provide a reminder that agents' ability to discern, and make judgements about, agency arises as they *actively* close the action-perception loop, not just in passive perceptual processes. The results also offer an indication of how agents do this, namely through exploratory titration of prediction error, in a pattern that is sensitive to variability and volatility. It may be that affording agents the opportunity for exploration of the action-perception loop is critical for agency inference and judgement.

Comparing the two levels of uncertainty manipulated here, changes to variability caused the most broad-reaching effects. Under high variability, participants were less accurate but spent longer sampling the environment, moved slower but accelerated more quickly, switched hypotheses less frequently and turned less, compared to the low variability conditions. The finding that participants move more under increased variability is consistent with the findings by Wen and Haggard (2020) in a similar judgement of agency paradigm.

While volatility was expected to have effects independent from variability, most of the significant effects for volatility were interactions with variability; volatility only showed two main effects. The first main effect indicated that prediction error was reduced more quickly under high volatility. In our manipulation, the timing of volatile switches was unpredictable, so this effect is likely due to an increased vigilance or sensitivity to incoming information manifesting as an increased learning rate under high volatility (Mathys, Daunizeau, Friston, & Stephan, 2011). The second main effect of volatility indicated that higher volatility was related to higher prediction error in the epoch surrounding hypothesis switches, however this was only true when variability was low. In two further cases, the effect of volatility was only seen in low variability; specifically that participants move faster and switch hypotheses more in high volatility than low. This could reflect an attempt to garner more evidence about the current state of the world before it changes. Lower volatility also magnified the effect of variability on the number of turns made during each trial. Higher volatility, on the other hand, increased the effect of variability on task accuracy. Future studies should consider ways of highlighting changes in volatility to enhance the potential effect of higher order uncertainty, such as making them large enough to stand out more saliently to the participant.

It is important to keep in mind too that our analyses of prediction error were limited to a behavioural proxy (combining eye-tracking and mouse movement) for prediction error that does not directly reflect changing internal representations of environmental uncertainty. This also affects what conclusions we can draw about the relationship between certain kinds of uncertainty and prediction error where the uncertainty strongly affected accuracy. For example, in our statistical model of the rate of prediction error minimisation split by uncertainty conditions, the effect of accuracy on the slope of prediction error was nearly identical to the effect of decreased variability on the slope of prediction error. This means that our results cannot distinguish between steeper prediction error minimisation due to an easier task or lower variability (see Supplementary Materials for statistical models including and removing the random effect of accuracy). To address the question of the effect of variability (independent of accuracy) on the rate of prediction error minimisation in this kind of task, difficulty across conditions could be titrated for each participant by adjusting other features which may affect task difficulty (such as distracter similarity or number, as in Williams and Happé (2009) and Grainger et al. (2014)). Including more levels of variability (rather than simply high and low), may also help to statistically distinguish the effect of accuracy and variability for future experiments. Future research should also consider using neural estimates of prediction error or computational modelling that appropriately changes priors with uncertainty.

Here, we found no difference in accuracy of judgement of agency between healthy participants across a range of autism traits, consistent with previous research comparing autistic and healthy participants on similar measures (David et al., 2008; Grainger et al., 2014; Russell & Hill, 2001; Williams & Happé, 2009; Zalla et al., 2015). As previously noted by Perrykkad and Hohwy (2020b) and Zalla and Sperduti (2015), this is in contrast to sense of agency in autism being shown to be reduced under typical experimental paradigms (Sperduti, Pieron, Leboyer, & Zalla, 2014; van Laarhoven, Stekelenburg, Eussen, & Vroomen, 2019). Our study also shows no main effects of AQ on other outcomes, except for a negative association with acceleration.

To our knowledge, Zalla et al. (2015) is the only other case where variability of a similar kind (which they labelled '*turbulence*') was added in a judgement of agency task pertaining to autism, in their case contrasting participants with and without an autism spectrum diagnosis. Their results demonstrated that the accuracy of autistic participants' agency judgements was less sensitive to differences in variability than the neurotypical group. This study supports our hypothesis that the addition of uncertainty has a distinctive effect on judgement of agency related to autistic traits. While we do not show any significant interactions with AQ in accuracy, our results showed participants with

high autism traits were less sensitive to differences in variability in their average prediction error. Since this measure is behavioural, this suggests that participants with high AQ were moving (that is, exploring the environment) in a way that did not reflect underlying differences in variability. Further, AQ was negatively associated with prediction error in the 200 ms window surrounding hypothesis switches. This suggests participants with high AQ are switching hypotheses earlier than participants with low AQ, or tolerating less uncertainty before abandoning their current hypothesis (see also Lawson et al. (2017)).

By additionally manipulating volatility, we could demonstrate further effects of uncertainty dependent on AQ. Participants with high AQ were more sensitive to differences in volatility such that only for this group was increased volatility associated with more time spent moving (if only in high variability) and more flexibility in policy selection. This might reflect less consistent or shallower internal models (Perrykkad & Hohwy, 2020b), which leads to less precision over all policies in high volatility, and so the selection of one over another fluctuates more frequently. This pattern is the opposite of the low AQ group, where high volatility was associated with more dominant policy use (but only in low variability). This is also consistent with Lawson et al. (2017), who showed that autistic participants update their learning in response to volatility more readily than neurotypical participants.

Of note, our findings with respect to autism are limited to scores on a trait-based measure, which may not generalise to diagnosed autistic populations. Our sample had a high average AQ score compared to what is expected in the general population (Baron-Cohen et al., 2001), so our results for "low" AQ may actually be more representative of "average" AQ individuals. While overall the sample size in post-hoc analyses is low, the omnibus interactions were based on modelled trends in the full dataset of continuous AQ scores. Nevertheless, environmental uncertainty might be particularly relevant to action selection for different levels of autistic traits and we do show interactions between uncertainty and AQ. These are worth following up in future studies in diagnosed samples.

In summary, this suggests autistic traits are related to 1) subtle differences in more abstract action policies, which are more sensitive to volatility, 2) smaller differences in prediction error between variability conditions, and 3) a greater propensity to switch hypotheses at a lower prediction error threshold when inferring agency. Notably, despite these differences, there was no significant effect of AQ on overall number of hypothesis switches or on accuracy.

4. Conclusion

This experiment shows that uncertainty in the mapping between actions and their outcomes changes not only how effectively participants can identify which stimuli they have control over, but also changes the actions they make and the overall strategies they employ. These changes have downstream impacts on the prediction error which can be used to inform their next action, and their overall response in each trial. In addition, our data illuminates subtle differences in this perception-action loop dependent on autism traits.

Acknowledgements

The authors appreciate the time and effort of their participants without whom this research would not be possible. The authors would also like to thank members of the Cognition and Philosophy Lab, particularly Andrew Corcoran, Jonathan Robinson, Stephen Gadsby, and Andrew McKilliam for their helpful comments on an earlier version of this paper. Thanks also to Chase Sherwell and Bryan Paton for analysis advice.

Perrykkad is supported by an Australian Government Research Training Program (RTP) Scholarship. Hohwy supported this project using an Australian Research Council Discovery Grant (DP160102770, DP190101805). Lawson is supported by a Wellcome Trust Henry Dale Fellowship (206691/Z/17/Z) and an Autistica Future Leaders Award (7265). Jamadar is supported by an Australian National Health and Medical Research Council Fellowship (APP1174164) and the Australian Research Council Centre of Excellence for Integrative Brain Function (CE1400100007).

Appendix A. Supplementary data

Supplementary data to this article can be found online at https://doi.org/10.1016/j.cognition.2021.104598.

References

Baron-Cohen, S., Wheelwright, S., Skinner, R., Martin, J., & Clubley, E. (2001). The autism-spectrum quotient (AQ): Evidence from Asperger syndrome/high-functioning autism, Malesand females, scientists and mathematicians. *Journal of Autism and Developmental Disorders, 31*(1), 5–17. https://doi.org/10.1023/A:1005653411471.

Bednark, J. G., Poonian, S., Palghat, K., McFadyen, J., & Cunnington, R. (2015). Identity-specific predictions and implicit measures of agency. *Psychology of Consciousness: Theory, Research and Practice, 2*(3), 253.

Clark, A. (2015). *Surfing uncertainty: Prediction, action, and the embodied mind*. Oxford University Press.

David, N., Gawronski, A., Santos, N. S., Huff, W., Lehnhardt, F.-G., Newen, A., & Vogeley, K. (2008). Dissociation between key processes of social cognition in autism: Impaired Mentalizing but intact sense of agency. *Journal of Autism and Developmental Disorders, 38*(4), 593–605. https://doi.org/10.1007/s10803-007-0425-x.

Desantis, A., Hughes, G., & Waszak, F. (2012). Intentional binding is driven by the mere presence of an action and not by motor prediction. *PLoS One, 7*(1), e29557. https://doi.org/10.1371/journal.pone.0029557.

Engbert, K., & Wohlschlager, A. (2007). Intentions and expectations in temporal binding. *Consciousness and Cognition, 16*(2), 255–264. https://doi.org/10.1016/j.concog.2006.09.010.

Fehd, H. M., & Seiffert, A. E. (2008). Eye movements during multiple object tracking: Where do participants look? *Cognition, 108*(1), 201–209. https://doi.org/10.1016/j.cognition.2007.11.008.

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

Friston, K. (2017). The variational principles of action. In J.-P. Laumond, N. Mansard, & J.-B. Lasserre (Eds.), *Geometric and numerical foundations of movements* (pp. 207–235). Cham: Springer International Publishing.

Friston, K., Samothrakis, S., & Montague, R. (2012). Active inference and agency: Optimal control without cost functions. *Biological Cybernetics, 106*(8), 523–541. https://doi.org/10.1007/s00422-012-0512-8.

Friston, K., Schwartenbeck, P., Fitzgerald, T., Moutoussis, M., Behrens, T., & Dolan, R. (2013). The anatomy of choice: Active inference and agency. *Frontiers in Human Neuroscience, 7*(598). https://doi.org/10.3389/fnhum.2013.00598.

Gallucci, M. (2019). GAMLj: General analyses for linear models [jamovi module]. Retrieved from https://gamlj.github.io/.

Grainger, C., Williams, D., & Lind, S. E. (2014). Online action monitoring and memory for self-performed actions in autism Spectrum disorder. *Journal of Autism and Developmental Disorders, 44*, 1193–1206.

Hohwy, J. (2013). *The predictive mind*. Oxford University Press.

Hohwy, J. (2015). Prediction, agency, and body ownership. In *The Pragmatic Turn:: Toward Action-Oriented Views in Cognitive Science* (pp. 109–120). The MIT Press.

Hughes, G., Desantis, A., & Waszak, F. (2013). Mechanisms of intentional binding and sensory attenuation: The role of temporal prediction, temporal control, identity prediction, and motor prediction. *Psychological Bulletin, 139*(1), 133.

Kuhn, S., Nenchev, I., Haggard, P., Brass, M., Gallinat, J., & Voss, M. (2011). Whodunnit? Electrophysiological correlates of agency judgements. *PLoS One, 6*(12), e28657. https://doi.org/10.1371/journal.pone.0028657.

van Laarhoven, T., Stekelenburg, J. J., Eussen, M. L., & Vroomen, J. (2019). Electrophysiological alterations in motor-auditory predictive coding in autism Spectrum disorder. In *Autism research*.

Lawson, R. P., Mathys, C., & Rees, G. (2017). Adults with autism overestimate the volatility of the sensory environment. *Nature Neuroscience*. https://doi.org/10.1038/nn.4615.

Majchrowicz, B., & Wierzchoń, M. (2018). Unexpected action outcomes produce enhanced temporal binding but diminished judgement of agency. *Consciousness and Cognition, 65*, 310–324. https://doi.org/10.1016/j.concog.2018.09.007.

Mathys, C., Daunizeau, J., Friston, K. J., & Stephan, K. E. (2011). A Bayesian foundation for individual learning under uncertainty. *Frontiers in Human Neuroscience, 5*, 39. https://doi.org/10.3389/fnhum.2011.00039.

Moore, J., & Haggard, P. (2008). Awareness of action: Inference and prediction. *Consciousness and Cognition, 17*(1), 136–144. https://doi.org/10.1016/j.concog.2006.12.004.

Palmer, C. J., Lawson, R. P., & Hohwy, J. (2017). Bayesian approaches to autism: Towards volatility, action, and behavior. *Psychological Bulletin, 143*(5), 521–542. https://doi.org/10.1037/bul0000097.

Palmer, C. J., Paton, B., Hohwy, J., & Enticott, P. G. (2013). Movement under uncertainty: The effects of the rubber-hand illusion vary along the nonclinical autism spectrum. *Neuropsychologia, 51*(10), 1942–1951.

Palmer, C. J., Paton, B., Kirkovski, M., Enticott, P. G., & Hohwy, J. (2015). Context sensitivity in action decreases along the autism spectrum: A predictive processing perspective. *Proceedings of the Royal Society of London B: Biological Sciences, 282* (1802), Article 20141557.

Perrykkad, K., & Hohwy, J. (2020a). Fidgeting as self-evidencing: A predictive processing account of non-goal-directed action. *New Ideas in Psychology, 56*, Article 100750. https://doi.org/10.1016/j.newideapsych.2019.100750.

Perrykkad, K., & Hohwy, J. (2020b). Modelling me, modelling you: The autistic self. *Review Journal of Autism and Developmental Disorders, 7*, 1–31. https://doi.org/10.1007/s40489-019-00173-y.

Perrykkad, K., Lawson, R. P., Jamadar, S., & Hohwy, J. (2020). *Perrykkad, Lawson, Jamadar, Hohwy - Action-Perception Loop Under Uncertainty Dataset*. Figshare. https://doi.org/10.26180/5ed0708f103a2.

R Core Team. (2018). R: A Language and environment for statistical computing. Retrieved from https://cran.r-project.org/.

Russell, J., & Hill, E. L. (2001). Action-monitoring and intention reporting in children with autism. *Journal of Child Psychology and Psychiatry, 42*(3), 317–328. https://doi.org/10.1111/1469-7610.00725.

Singmann, H., & Kellen, D. (2020). An introduction to mixed models for experimental psychology. In D. H. S. E. Schumacher (Ed.), *New methods in neuroscience and cognitive psychology* (pp. 4–31). New York, NY: Routledge.

Sperduti, M., Pieron, M., Leboyer, M., & Zalla, T. (2014). Altered pre-reflective sense of agency in autism spectrum disorders as revealed by reduced intentional binding. *Journal of Autism and Developmental Disorders, 44*(2), 343–352.

The Jamovi Project. (2019). jamovi (Version 0.9). Retrieved from https://www.jamovi.org.

Wen, W., & Haggard, P. (2020). Prediction error and regularity detection underlie two dissociable mechanisms for computing the sense of agency. *Cognition, 195*, Article 104074. https://doi.org/10.1016/j.cognition.2019.104074.

Williams, D., & Happé, F. (2009). Pre-conceptual aspects of self-awareness in autism spectrum disorder: The case of action-monitoring. *Journal of Autism and Developmental Disorders, 39*(2), 251–259.

Wolpert, D. M. (1997). Computational approaches to motor control. *Trends in Cognitive Sciences, 1*(6), 209–216.

Zaadnoordijk, L., Besold, T. R., & Hunnius, S. (2019). A match does not make a sense: On the sufficiency of the comparator model for explaining the sense of agency. *Neuroscience of Consciousness, 2019*(1), Article niz006.

Zalla, T., Miele, D., Leboyer, M., & Metcalfe, J. (2015). Metacognition of agency and theory of mind in adults with high functioning autism. *Consciousness and Cognition, 31*, 126–138. https://doi.org/10.1016/j.concog.2014.11.001.

Zalla, T., & Sperduti, M. (2015). The sense of agency in autism spectrum disorders: A dissociation between prospective and retrospective mechanisms? *Frontiers in Psychology, 6*, 1278. https://doi.org/10.3389/fpsyg.2015.01278.

Zama, T., Takahashi, Y., & Shimada, S. (2017). The effects of trajectory and endpoint errors in a reaching movement on the sense of agency. *Psychology, 8*(14), 2321.

Moore, J. W., & Obhi, S. S. (2012). Intentional binding and the sense of agency: a review. *Consciousness and Cognition, 21*(1), 546–561.