bioRxiv preprint doi: https://doi.org/10.1101/2023.01.17.524436; this version posted January 20, 2023. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC 4.0 International license.

Metacognitive Judgments during Visuomotor Learning Reflect the Integration of Error History

Christopher L. Hewitson[1], Naser Al-Fawakhiri[1], & Samuel D. McDougle[1,2]

1. Department of Psychology, Yale University, New Haven CT 06520
2. Wu Tsai Institute, Yale University, New Haven CT 06520

Running head: Confidence during Visuomotor Learning Reflect Error History

Correspondence: Christopher L. Hewitson
christopher.hewitson@yale.edu
Department of Psychology
Yale University
2 Hillhouse Avenue, New Haven, CT, 06520

AUTHOR EMAILS
christopher.hewitson@yale.edu
naser.al-fawakhiri@yale.edu
samuel.mcdougle@yale.edu

AUTHOR CONTRIBUTIONS
C.L.H and S.D.M designed the paradigm; C.L.H collected data. C.L.H and N.A analyzed data. C.L.H, N.A and S.D.M prepared figures, drafted manuscript, edited, and revised manuscript.

AUTHOR FUNDING
Christopher Louis Hewitson was funded by Yale University's Seesel Postdoctoral Fellowship.

Pages: 30
Figures: 4
Tables: 2
Supplemental: 1 table

bioRxiv preprint doi: https://doi.org/10.1101/2023.01.17.524436; this version posted January 20, 2023. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC 4.0 International license.

Metacognitive Judgments during Visuomotor Learning Reflect the Integration of Error History

Christopher L. Hewitson[1], Naser Al-Fawakhiri[1], & Samuel D. McDougle[1,2]

1. Department of Psychology, Yale University, New Haven CT 06520
2. Wu Tsai Institute, Yale University, New Haven CT 06520

Correspondence: Christopher L. Hewitson
Department of Psychology
Yale University
Email: *christopher.hewitson@yale.edu*

ABSTRACT

People form higher-level, metacognitive representations of their own abilities across a range of tasks. Here we ask how metacognitive confidence judgments of performance during motor learning are shaped by the learner's recent history of errors. Across two motor adaptation experiments, our computational modeling approach demonstrated that people's confidence judgments are best explained by a recency-weighted averaging of observed motor errors. Moreover, in the formation of these confidence estimates, people appear to re-weight observed motor errors according to a subjective cost function. Finally, confidence judgments appeared to incorporate recent motor errors in a manner that was sensitive to the volatility of the learning environment, integrating a shallower history when the environment was more volatile. Our study provides a novel descriptive model that successfully approximates the dynamics of metacognitive judgments during motor learning.

NEW & NOTEWORTHY

This study examined how, during visuomotor learning, people's confidence in their movement decisions is shaped by their recent history of errors. Using computational modeling, we found that confidence judgments incorporated recent error history, tracked subjective error costs, and were sensitive to environmental volatility. Together, these results provide a novel model of metacognitive judgments during motor learning that could be applied to future computational and neural studies at the interface of higher-order cognition and motor behavior.

bioRxiv preprint doi: https://doi.org/10.1101/2023.01.17.524436; this version posted January 20, 2023. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC 4.0 International license.

INTRODUCTION

Humans have the ability to monitor qualities of their own performance in a task, a capacity often referred to as "metacognition." Metacognitive processes have been observed across a range of tasks, including simple perceptual decision-making (1), reinforcement learning (2), social cognition (3), and memory (4). Over a century of research has shown that people's metacognitive judgements (such as their explicitly reported confidence in their choices/abilities) often closely track behavioral metrics like accuracy and response time (5).

In one of the most studied laboratory models of metacognition and confidence – perceptual decision-making – researchers have used computational models to uncover strong links between one's confidence in a choice (e.g., 'those dots are mostly moving left') and the perceptual evidence they have accumulated for that choice over its competitors (6–8). Moreover, researchers have even discovered certain neural populations that simultaneously encode both accumulated evidence and decision confidence (9). Here, we turn to a domain that has been less well studied with respect to metacognition – sensorimotor learning.

Unlike making discrete, independent decisions about incoming sense data, learning requires integrating information over protracted periods. Thus, metacognitive awareness of your state of learning requires tracking your progress across time. Consider practicing your tennis serve over a series of attempts: Your metacognitive judgment of your current ability (e.g., your confidence in any given serve attempt) should, in principle, take into account your recent history of feedback (i.e., your errors). But how does one's state of confidence integrate these errors, especially when they are in a continuous domain (i.e., like most motor learning tasks)? And how does confidence relate to second order statistics of learning, like the volatility of the environment (e.g., a particularly windy day on the courts) (10)?

There has been some recent research on confidence and learning in nonmotor domains. One recent study (11) used a perceptual decision task in which participants reported their estimate of the transition probabilities between two visual or auditory stimuli as well as their confidence in this report. The results indicated that participants not only learn a statistical model of transition probabilities over time, but also that their confidence ratings closely track this statistical inference. This work demonstrates that in a perceptual decision-making context, people's

bioRxiv preprint doi: https://doi.org/10.1101/2023.01.17.524436; this version posted January 20, 2023. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC 4.0 International license.

confidence judgments closely correlate with their performance in tracking stochastic variables over time (12).

Other work from the reinforcement learning domain has suggested that confidence in one's choices during learning evolves along with learned latent value representations, and is subject to value-driven biases (2). Moreover, volatility in an environment, a second order statistic tracked over many trials, induces uncertainty in an agent, and agents tend to operate with a faster learning rate in these uncertain environments (13). This work suggests that higher-order variables like confidence may also correspond to the statistical uncertainty that underpins the learning process itself (14, 15).

Subjective confidence in the domain of motor learning has been less well studied, but some work has attempted to capture the role of continuous motor errors on subjective evaluations of confidence. For instance, one recent study (16) showed that individuals are able to predict their future performance, and also leverage their confidence in their future performance to maximize future rewards. Another recent study (17) demonstrated that subjective confidence tracks precision in a continuous temporal estimation task. Some work on motor sequence learning has looked at a more 'zoomed-out' form of confidence – block- and day-level judgments of one's own ability (18). Lastly, recent computational work has shown that individuals might utilize information about their prior motor variability to make confidence judgments of their motor precision (19).

While these works suggest that confidence in a motor context integrates prior history of performance (perhaps in a Bayesian manner), they do not directly address metacognitive dynamics during the protracted adaptation of motor commands (20), the context of interest here. During motor adaptation, does confidence simply reveal a metacognitive readout of performance error at a given point in time, or does it represent the integration of a history of errors? And how do aspects of the learning context, such as the volatility of the environment, mediate the relationship between confidence and motor error? Addressing these questions can shed light on the psychological processes involved in motor learning, computationally isolate higher-level metacognitive variables for investigating in future neural studies, and perhaps be useful for increasing people's motivation to learn in clinical and non-clinical settings.

bioRxiv preprint doi: https://doi.org/10.1101/2023.01.17.524436; this version posted January 20, 2023. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC 4.0 International license.

Here, we used a motor adaptation task that involves modulating movement kinematics (i.e., reaching directions), and asked how motor errors affect subjective confidence. We address this via two experiments and descriptive computational modeling. We specify a model of confidence during motor learning where trial-by-trial subjective confidence judgments are approximated by a simple linear dynamic system that tracks subjectively-weighted errors made during motor learning. This straightforward model outperforms other model variants that do not incorporate error history, and also reveals that the dynamics of confidence ratings during motor adaptation are sensitive to environmental volatility. Together, these results set the stage for future computational and neural investigations of people's higher-level metacognitive representations of the state of their own sensorimotor learning processes.

METHODS

Participants

A total of 38 neurologically healthy participants (Experiment 1: N = 18 ; Age = 21±5 years; Gender: 10 identified as Female, 1 preferred not to answer; Handedness: 16 right-handed [>40 on the Edinburgh handedness inventory (21). Experiment 2: N = 20; Age = 20±5 years; Gender: 13 identified as Female; Handedness: 19 right-handed) from the Yale University subject pool participated in this study. They received monetary compensation or course credit for their participation. Written informed consent was obtained from all participants before testing and the experimental protocol was approved in accordance with Yale University's Institutional Review Board. No subjects were excluded from any of our analyses.

Apparatus

Participants sat on a height-adjustable chair facing a 24.5-in. LCD monitor (Asus VG259QM; display size: 543.74 mm x 302.62 mm; resolution: 1920 x 1080 pixels; frame rate 280 Hz; 1 ms response time), positioned horizontally ~30 cm in front of the participant above the table platform, thus preventing direct vision of the hand (Fig. 1A). In their dominant hand they held a stylus embedded within a custom modified paddle, which they could slide across a digitizing tablet (Wacom PTH860; active area: 311 mm x 216 mm). Hand position was recorded from the tip of the stylus sampled by the tablet at 200 Hz. Stimulus presentation and movement recording were controlled by a custom-built Octave script (GNU Octave v5.2.0; Psychtoolbox-3 v3.0.18; Ubuntu 20.04.4 LTS). Aiming and confidence ratings were controlled by the non-dominant hand and entered on a USB keyboard (Fig. 1B).

bioRxiv preprint doi: https://doi.org/10.1101/2023.01.17.524436; this version posted January 20, 2023. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC 4.0 International license.

Task

Typical task trials consisted of an aiming phase, a confidence rating phase, and then a reaching phase (Fig. 1B). To briefly summarize the phases: During the aiming phase, participants were instructed to “position the aiming reticle where you intend to move your hand”; during the confidence reporting phase, participants were instructed to “rate how confident you are that where you aimed is correct”; and during the reach phase, subjects made rapid reaches to the displayed target.

During the reaching phase, participants performed center-out-reaching movements from a central start location in the center of the monitor to one of 8 visual targets (0°, 45°, 90°, 135°, 180°, 225°, 270°, and 315°) arranged around an invisible circle with a radius of 10 cm. The target location for each trial was pseudo-randomly selected. Participants were instructed to move the stylus as quickly as possible from the start location in the direction of the displayed target and “slice through it.” The start location was marked by a filled white circle of 7 mm in diameter. The target locations were marked by filled green circles of 10 mm in diameter. Online visual feedback was given by a cursor (filled white circle, radius 2.5 mm). If the reach duration exceeded 400 ms, a text prompt appeared on the monitor reminding participants to “please speed up your reach,” and the trial was repeated but with a new target location.

During the aiming phase, a white crosshair 7 mm in diameter was overlaid on the target (Figure 1B). Its movement was constrained to follow the arc of the invisible circle with a radius of 10 cm from the start location. The aiming crosshair’s location was adjusted with the left hand using the left and right arrow keys, which drove crosshair movements to the counterclockwise and clockwise directions, respectively. When participants were satisfied with the match between their intended reach direction and the aiming crosshair’s position, they then registered their aim with the ‘enter’ key.

During the confidence rating phase, which directly followed the aiming phase, a rating bar (20mm x 40mm) was displayed 15˚ counterclockwise of the target. A white line, representing the participant’s confidence rating was initialized in the middle of the bar (50% confidence). As confidence increased towards 100%, the bar’s color changed from yellow to green. As confidence decreased towards 0%, the bar’s color changed from yellow to red. Participants

bioRxiv preprint doi: https://doi.org/10.1101/2023.01.17.524436; this version posted January 20, 2023. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC 4.0 International license.

reported their confidence level with their left hand using the up and down arrows, and registered their confidence rating with the 'enter' key.

Experiment 1 included reach baseline, report practice, adaptation, and washout blocks (Figure 1C). Baseline consisted of 24 trials (3 trials per target) with veridical online cursor feedback provided for the duration of the reach. Report practice consisted of 48 trials (6 trials per target) partitioned into the 'aim', confidence report,' and 'reach' phases. Veridical visual feedback was provided throughout all reaches, save the washout phase in Experiment 1. Adaptation consisted of 240 trials (30 trials per target) that included all three trial phases (Figure 1C). Crucially, during the reach phase the cursor was rotated by 30˚ (with CW/CCW rotations evenly counterbalanced across participants). Washout (Experiment 1 only) consisted of 48 reach trials (6 trials per target) with no cursor feedback provided for the duration of the reach, analogous to the baseline phase.

Experiment 2 included 16 baseline trials (2 per target), 16 report practice trials (2 per target), 208 adaptation trials (26 per target) but no washout trials (Figure 1D). The adaptation trials differed from Experiment 1 only in terms of the rotation perturbation applied to the cursor. In Experiment 2, rotation angles of -60˚, -45˚, -30˚, -15˚, 15˚, 30˚, 45˚, and 60˚ were pseudo-randomly applied across 24 blocks of 8 trial mini-blocks (3 x 8 trial mini-blocks per rotation angle, thus 192 rotation trials total). Four additional mini blocks consisting of 4 trials of 0˚ rotation were interleaved throughout adaptation. No specific rotation angle or sign was repeated consecutively (Figure 1D).

Statistical Analysis

Primary dependent variables were confidence judgements and recorded hand angles on every trial. Since participants were instructed to always adjust the confidence bar by at least one unit, all trials where the confidence rating remained at the initial 50% were removed (Exp. 1: 353 out of 5,184 trials [7.39%]; Exp. 2: 390 out of 4,480 trials [8.71%]). Data was analyzed using Matlab (Mathworks, Inc. Version 2022a). Model fits were computed using Matlab's *fmincon* function, minimizing the SSE between our confidence models and the confidence report data. Violin plots were generated using the *Violinplot* function in Matlab (22). Data and analysis code can be accessed at https://github.com/ilestz/confidence_analysis.

We validated model parameter optimization through parameter recovery and found we could achieve stable parameter fits throughout. To do this, we fit model-predicted confidence reports

bioRxiv preprint doi: https://doi.org/10.1101/2023.01.17.524436; this version posted January 20, 2023. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC 4.0 International license.

using the model that initially generated these predictions and found that model parameters were recovered to 100% accuracy within 2 iterations of fitting. R^2 values were computed through linear regression of model predictions and data. Reported ΔAIC (23) values reflect differences in summed AIC values between each model and the winning model.

RESULTS

Computational Modeling

The goal of our study was to examine how motor learning relates to one's metacognitive judgments of their movement decisions. Participants performed a standard sensorimotor adaptation task while also reporting their confidence in each of their movements (i.e., chosen reach directions; Figure 1). We constructed computational models with the goal of predicting these subjective confidence reports on each trial. All four models characterized confidence reports (Equation 1) as deviations from a maximum confidence 'offset' that is proportional to an estimate of previous sensorimotor error(s):

$$(1) \qquad Confidence = Confidence_{max} - \eta * Error$$

Where η represented an error scaling parameter. Here, "error" denotes the experienced "target error" (henceforth TE), the absolute angular error of the cursor relative to the target (though see later *Results* sections for alternatives). The first class of models we tested predict confidence based on the *current state of learning*. Specifically, these models relate confidence reports directly to the most recent error signal, representing a local "one-trial-back" (OTB) update rule. Within this model class, we tested two model variants. In the objective-error one-trial-back model (OTB_{obj}), the true (i.e., actually observed) absolute error of the cursor relative to the target was used to compute confidence. However, previous work has shown that the cost of target errors are scaled subjectively via an approximate power-law (24). Thus, we also fit a subjective-error one-trial-back model (OTB_{subj}), which scaled all target errors by an exponential free parameter, γ:

$$(2) \qquad Confidence_t = Confidence_{max} - \eta * TE_{t-1}^{\gamma}$$

An exponent $\gamma > 1$ suggests that large target errors are perceived as relatively more salient (costly), and thus drive sharper decreases in confidence versus small errors in a manner that is

bioRxiv preprint doi: https://doi.org/10.1101/2023.01.17.524436; this version posted January 20, 2023. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC 4.0 International license.

non-linearly proportional to the veridical error. In contrast, $0 < \gamma < 1$ suggests that large errors are discounted relative to their veridical magnitude and thus drive weaker decreases in confidence than would be predicted by the objective error model. Finally, $\gamma = 1$ reduces to the objective error case, where errors are not subjectively scaled, and confidence is linearly proportional to the veridical target error magnitude.

The second class of models involved retaining a sort of *memory* via estimating a running average of target errors across trials. This estimate is subsequently used to generate predicted confidence reports. We designed these "error-state-space" models to act as simple linear dynamical systems that update an estimate of the current error "state" on every trial through a canonical delta rule:

$$\widehat{TE}_t = \widehat{TE_{t-1}} + \alpha \cdot \delta_{t-1} \quad (3)$$

$$\delta_{t-1} = TE_{t-1} - \widehat{TE_{t-1}} \quad (4)$$

In effect, this learning rule constructs a recency-weighted average of the error state across trials and is similar to the learning rule employed in instrumental learning contexts for learning the predictive value of a given stimulus (25) and echoes state-space models used to model adaptation itself (26). The learning rate α reflects the degree to which errors on previous trials are incorporated into the estimate, with high α values (i.e., close to 1) reflecting a high degree of forgetting and low α values (i.e., close to 0) reflecting a more historical memory of error across trials. Whenever the observed target error on the previous trial was greater than (less than) the estimated target error, the estimated target error would increase (decrease) by an amount proportional to this "metacognitive" prediction error.

Within this error-state-space (ESS) model class, there were again two distinct variants: The objective-error state-space model (ESS_{obj}) computed the estimated error using the true veridical target errors, while the subjective-error state-space model (ESS_{subj}) computed the estimated error using the subjectively scaled error with exponent γ:

$$\delta_{t-1} = TE_{t-1}^{\gamma} - \widehat{TE_{t-1}} \quad (5)$$

bioRxiv preprint doi: https://doi.org/10.1101/2023.01.17.524436; this version posted January 20, 2023. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC 4.0 International license.

Where the estimated error tracks a history of subjective errors, instead of objective errors. Nonetheless, both history (ESS) models used the same equation to generate predicted confidence reports (Equation 1), now using an evolving estimate of error state:

$$(6) \qquad Confidence = Confidence_{max} - \eta * \widehat{TE}_t$$

Altogether, our 4 models share 2 free parameters – the maximum confidence offset and the error sensitivity scaling parameter η. Moreover, both models with nonlinear subjective error cost functions share the γ parameter. Finally, both models with error state-space tracking share an additional learning rate free parameter (α) relative to the one-trial-back models.

[Figure 1 here]

Experiment 1

We sought to explore how participants generate subjective judgements of their confidence in a motor learning task (Figure 1). Prior to performing a center out reaching motion, participants reported their intended reach direction and rated their confidence in that decision on a continuous scale (Figure 1B). After a brief baseline phase with veridical cursor feedback, a sensorimotor perturbation of 30° was applied (Figure 1C), which subjects rapidly learned to compensate for. Reach directions compensated well for the applied rotation, with a mean cursor error over the last 50 rotation trials of 0.51° (SD: 3.4°) (Fig. 2A). On average, reach directions compensated for 90% of the perturbation after ~8 trials.

[Figure 2 here]

During the unperturbed baseline phase of the experiment, confidence reports remained relatively stable but sharply decreased when the perturbation was applied, as expected. Unsurprisingly, all of our models of confidence were able to account for this decrease. Following the initial decrease in confidence, all participants gradually restored confidence to near baseline levels as their reaching errors decreased (Fig. 2B), which was also captured by all models. These expected observations provide initial support for the general form of Equation 1, where confidence is proportional to error.

To get a better picture of the dynamics of subjects' metacognitive judgments, we now turn to model comparisons. To reiterate the models tested: one class of models, the "one-trial-back"

bioRxiv preprint doi: https://doi.org/10.1101/2023.01.17.524436; this version posted January 20, 2023. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC 4.0 International license.

models (Equation 2), predicted confidence reports on a given trial based on the current state of learning. The objective-error one-trial-back (OTB_{obj}) model predicted confidence based on veridical absolute cursor errors relative to the target, and the subjective-error one-trial-back (OTB_{subj}) model predicted confidence based on errors which were scaled by a power-law. The second class of models, the "error-state-space" models, kept track of an estimate of the average error on recent trials and used this average error to compute predicted confidence reports. Again, this class of models either kept track of objective errors (ESS_{obj} model) or subjectively scaled errors (ESS_{subj} model).

Of the four confidence models we developed, the ESS_{subj} best explained the variance in confidence reports (R^2=0.41±0.23 [mean±SD]). At the individual level, 14 out of the 18 total subjects were better fit by the winning model versus the second-best model (Fig. 2C). Moreover, the ESS class of models robustly outperformed the OTB class (Fig. 2C; Table 1). This suggests that metacognitive judgments during sensorimotor learning incorporate a continually updated history of recent errors, rather than simply acting as a "read-out" of the current state of learning. Furthermore, confidence was better explained by a subjective error term rather than an objective one: The ESS_{subj} model better fit the confidence versus the ESS_{obj} model (Table 1). All model comparisons were robust, with AIC differences relative to the best-fitting model all exceeding 500.

While the results of Experiment 1 clearly favored the ESS_{subj} model, some limitations remained. First, because the rotation was of a single value (30˚) and was fixed throughout the adaptation phase, the task was relatively easy. Thus, it was important to test if our modeling results generalized to a more complex environment, one where both errors and confidence reports would be more variable. Moreover, because of the nature of the task in Experiment 1, both learning curves and confidence reports monotonically increased together; a more variable environment would thus also help us rule out potential coincidental similarities in autocorrelation structure between our winning model and subjects' learning curves as the key factor. To that end, in Experiment 2 we implemented a pseudo-randomly varying perturbation schedule. This allowed us to control for the aforementioned limitations, while also testing a novel question – are the dynamics of metacognitive confidence judgments during sensorimotor learning affected by environmental uncertainty?

[Table 1 here]

bioRxiv preprint doi: https://doi.org/10.1101/2023.01.17.524436; this version posted January 20, 2023. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC 4.0 International license.

Experiment 2

Experiment 2 involved perturbations that fluctuated every few trials (i.e., the perturbation changed size and direction every 4 or 8 trials, see Figure 3A and *Methods*). This allowed us to perform a more strict test of our modeling approach, and to examine if and how environmental uncertainty affected subjective confidence reports. Specifically, we predicted that the ESS_{subj} model would best account for subjective confidence ratings in this context, replicating Experiment 1. Moreover, we also hypothesized that while the fundamental process of confidence ratings would remain the same (i.e., the ESS_{subj} would again best explain behavior), the learning rate parameter of that model would increase in response to the increase in environmental uncertainty such that it would incorporate a more recency-biased history of errors (13).

Despite the more volatile nature of the perturbation schedule, participants were able to adapt their reach directions to account for the rotations (Fig. 3A). Excluding the transition trials where the rotation abruptly changes, subjects' average cursor error in the last two perturbation blocks was only 6.3° (SD: 6.5°).

[Figure 3 here]

As in Experiment 1, confidence remained relatively stable during the unperturbed baseline but sharply decreased after the onset of the first perturbation (Fig. 3B). Confidence also tended to sharply decrease at the start of each new perturbation. Throughout the experiment, some 4-trial zero-rotation blocks were introduced, and these blocks tended to coincide with high confidence reports (Fig. 3B).

Once again, the ESS_{subj} model best predicted confidence reports in this experiment (R^2=0.42±0.20 [mean±SD]), and the one-trial-back models were unable to account for the large fluctuations in confidence reports and performed significantly worse (Table 2). At the individual level, 18 out of the 20 total subjects were better fit by the winning model versus the second-best model (Fig. 3C). Both history (ESS) models again tracked confidence reports more accurately than the other class of models (OTB). Thus, our model comparison results closely replicated those of Experiment 1 (and again were robust; lowest AIC difference: 498). This further suggests that metacognitive judgements of sensorimotor learning incorporate a gradually

bioRxiv preprint doi: https://doi.org/10.1101/2023.01.17.524436; this version posted January 20, 2023. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC 4.0 International license.

changing history of (subjectively scaled) errors. We do note that none of the four models were able to fully capture the unusually high confidence ratings seen during the zero-rotation blocks (see *Discussion*).

[Table 2 here]

Comparing model parameters across experiments

While the variance in confidence reports was best explained in both experiments by the subjective-error history model, parameter values in each experiment were not necessarily the same. We explored how best fitting parameters in the ESS_{subj} model changed across tasks (Fig. 4), with one key prediction that the metacognitive learning rate (α) would increase in Experiment 2 versus Experiment 1 due to the increase in environmental volatility (13).

[Figure 3 here]

We first looked at the maximum confidence offset parameter (Fig. 4A). This parameter should not necessarily differ across experiments, as it reflects an individual's maximum level of confidence in the task on a somewhat arbitrary scale that should largely be independent of the dynamics of the perturbation schedule. In fact, regardless of experiment, this parameter should be close to maximal (i.e., 100) if participants are using the full range to make their confidence reports. Consistent with this hypothesis, this parameter was not significantly different across experiments (Wilcoxon rank sum test, Z=-0.63, p=0.53). The mean value of this parameter was 89±11 in Experiment 1 and 91±14 in Experiment 2 (mean±SD).

The subjective scaling of errors varied across both experiments (Fig. 4B). In Experiment 1, the mean exponent γ was 1.9 (SD: 0.95), indicating increased sensitivity to large errors. However, γ values in Experiment 2 were significantly lower at 0.46±0.23 (Wilcoxon Rank Sum, Z=-4.98, p= 6.2×10^{-7}). The large difference in exponent values across experiments is expected, and likely reflects the fact that participants in Experiment 2 become habituated to large errors due to the volatility of the perturbation schedule, and thus likely learned to blunt the effect of these errors on their confidence reports. In contrast, in Experiment 1 errors were consistently very small, leading to the opposite effect. Thus, subjects appeared to alter how a subjective cost function of error shaped their confidence reports according to the distribution of errors they experienced. Consistent with an inherent trade-off between the exponent parameter and the sensitivity

bioRxiv preprint doi: https://doi.org/10.1101/2023.01.17.524436; this version posted January 20, 2023. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC 4.0 International license.

parameter η (Equation 2), we also observed a significant change in η between experiments (Wilcoxon Rank Sum, Z=-4.49, p=7.2x10^{-6}) (Figure 4B). Specifically, η was on average more than 10 times larger in Experiment 2 than in Experiment 1 (Exp. 1: 1.4±2.9; Exp. 2: 15±14). (We note here that this parameter should not be over-interpreted, as it is primarily a scaling factor used to map error units onto confidence report units.)

Finally, consistent with our hypothesis and with previous findings in reinforcement learning that learning rates in volatile environments are larger than those in stable environments (13), we observed that the error state-space learning rate α more than doubled in Experiment 2 relative to Experiment 1 (Exp. 1: 0.18±0.15; Exp. 2: 0.41±0.24; Wilcoxon Rank Sum, Z = -3.08, p = 0.002) (Figure 4D). Thus, the amount of "error history" that was incorporated into people's metacognitive judgments was modulated based on second-order statistics of the learning environment. These parameter changes across experiments reflect the ESS_{subj} model's flexibility in explaining fluctuations in confidence in both stable and volatile environments, and over different dynamic ranges of error. Taken together, our between-experiment parameter results (Fig. 4) suggest that subjects adapted the dynamic range and memory span of their confidence reports in a manner that reflected the statistics of the environment.

Fluctuations in confidence primarily track target errors

In motor learning, compensation for errors often reflects two distinct processes, one explicit and one implicit (27–29). The explicit process is thought to primarily reflect cognitive aiming strategies meant to deliberately reduce motor errors (29, 30). In contrast, the implicit process is thought to instead reflect gradual adjustments to an internal model, which proceed largely outside of conscious awareness (29). So far, the models discussed have used target error to predict confidence reports. Importantly, target error itself reflects the consequences of both implicit and explicit learning processes (29, 31). Because our task design had us ask subjects about their confidence in an explicitly-reported movement plan, we could use a simple subtractive method to dissociate explicit and implicit learning components (29). In order to determine whether confidence reports may have been specifically sensitive to explicit or implicit motor adaptation processes rather than target error alone, we performed an additional model fitting analysis that used distinct error terms related to each component (see supplemental table T1).

bioRxiv preprint doi: https://doi.org/10.1101/2023.01.17.524436; this version posted January 20, 2023. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC 4.0 International license.

In order to isolate the effect of the explicit component, we used the pre-reach aim reports (29). The explicit error component was quantified as the discrepancy between reported aim and the aim required to fully compensate for the rotation. The implicit error component can be isolated as the discrepancy between the true reach angle and the reported aim.

In both experiments, confidence model fits to only the implicit error component were significantly worse than those fit to target error (Exp. 1: R^2=0.15±0.19, t(17)=-5.23, p=6.2x10^{-5}; Exp. 2: R^2=0.27±0.23, t(19)=-5.74), p=1.6x10^{-5}; Supplemental Table T1). Additionally, in Experiment 1, the model fit using the explicit component was significantly worse than the model using target error (R^2=0.27±0.26, t(17)=-3.59, p=0.002). However, in experiment 2, the model fit on the explicit component was not significantly different to the model using target error (R^2=0.37±0.21, t(19)=-1.61, p=0.12). This makes sense – due to the volatile nature of the perturbation's sign and magnitude in Experiment 2, very little consistent implicit learning can accrue, meaning that the explicit error component is similar to the target error. As expected, confidence models using the explicit error component captured significantly more variance in confidence reports than confidence models dependent on the implicit error component in both experiments (Exp. 1: t(17)=2.82, p=0.01; Exp. 2: t(19)=2.34, p=0.03). Taken together, these additional analyses support the reasonable conclusion that the actually observed performance state – the target error – determines the dynamics of metacognitive judgments during sensorimotor learning.

DISCUSSION

This current study is the first, to our knowledge, to examine the relationship between subjective confidence judgments and motor errors in the context of sensorimotor adaptation. We investigated this relationship via two sensorimotor learning experiments that differed with respect to the environmental volatility (i.e., the perturbation schedule applied). We constructed computational models with the goal of predicting subjective confidence reports on each trial. We specified a set of models where trial-by-trial subjective confidence tracked only the current learning state (i.e., the most recent performance error), and another set of models where confidence judgments are approximated by a simple linear dynamic system that tracks a recency-weighted history of errors made during learning.

In Experiment 1, an error history model that used a subjective error term – the ESS_{subj} model – was best able to account for the confidence data in the context of a fixed perturbation schedule.

bioRxiv preprint doi: https://doi.org/10.1101/2023.01.17.524436; this version posted January 20, 2023. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC 4.0 International license.

The ESS_{subj} model had greater numerical agreement with confidence judgments over a veridical error model (ESS_{obj}) model throughout the experiment. In control analyses (see supplemental table T1), parsing the relative contributions of implicit and explicit error components indicated that, in a static learning context, metacognitive judgements primarily track the observed performance state (target error).

In Experiment 2 subjects learned in a volatile context, and the ESS_{subj} model was best able to account for the large fluctuations in confidence reports we observed. The ESS_{subj} model again had greater numerical agreement with confidence judgments over the ESS_{obj} model throughout the experiment, replicating the results reported for Experiment 1. Taken together, these findings demonstrate that confidence reports during sensorimotor adaptation are well approximated by a running average of recent (subjectively-scaled) performance errors. To wit, these findings suggest that when people make metacognitive judgments of their own state of sensorimotor learning, they incorporate a recent history of errors rather than just taking a snapshot of their current performance state.

Comparing model parameters between experiments provides key insights into the dynamics of metacognitive judgements of performance during sensorimotor adaptation. Although the range of confidence ratings are relative to individual participants, the similarity in the maximum confidence offset parameter between experiments suggests that participants operate within a comparable confidence range (Fig. 4A). The significant difference in the subjective error scaling exponents and error sensitivity parameters between experiments (Fig. 4b-c) was expected given the differences in perturbation schedules between experimental contexts. That is, participants scaled the subjective cost functions that shaped their confidence reports according to the distribution of errors they experienced. A large exponent in Experiment 1 indicated a non-linear increase in sensitivity to large errors when the environment was stable and errors were generally small. In contrast, a small exponent in Experiment 2 indicated that participants down-weighted large errors, likely due to an increased range of errors. These results show that the relationship between motor errors and confidence judgements was sensitive to the range of errors experienced. A number of previous studies that investigated the cost-function associated with errors in the context of sensorimotor control and learning have shown that people apply a non-linear cost-function that increases quadratically for small errors and significantly less than quadratically for large errors (24). We see a similar trend in our results (Experiment 2).

bioRxiv preprint doi: https://doi.org/10.1101/2023.01.17.524436; this version posted January 20, 2023. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC 4.0 International license.

Learning rates in our confidence model were significantly larger in Experiment 2, consistent with our hypothesis that they would be higher in the more volatile environment. These results comport with previous work in reinforcement learning, showing that a higher learning rate is more useful in a volatile environment because history is less informative (13). Our findings are consistent with these results, extending them to the dynamics of metacognitive judgments during motor learning. Future studies that parametrically alter aspects of the learning environment (e.g., consistency and variability, (10)) and measuring their effects on metacognitive judgements could be useful for developing more detailed models of confidence during motor learning.

What are the psychological mechanisms that track the error-state used for metacognitive judgments? Although our models are straightforward and principally descriptive, they do constrain the time-scale at which error-signals are integrated, hinting at a role for memory in the process of subjective confidence formation. We speculate that working memory is likely important in the formation of confidence judgments during motor learning (32, 33). That is, participants may track the quality of their performance by storing recent outcomes in working memory and integrating them into an estimate of the “state” of their performance. If this is correct, one prediction is that disrupting working memory may alter the relationship between confidence and recent errors. Future studies, perhaps using dual tasks, could test this prediction.

The prospect that higher-level metacognitive judgements accurately track lower-level sensorimotor properties (e.g., visual error magnitudes) compels the search for overlapping neural correlates. In the context of sensorimotor learning, confidence can be defined as a higher-order variable that corresponds to the uncertainty that underpins the learning process (14, 15). Multiple neural regions capable of representing sensory uncertainty have been proposed, including the orbitofrontal cortex (OFC) (34, 35), midbrain (36), anterior cingulate cortex (ACC) (37), insula (38, 39), and prefrontal cortex (PFC) (40, 41). In terms of representations of confidence, activity in the rostrolateral and dorsolateral PFC (rlPFC/DLPFC) is purported to be central to the processing of explicit confidence judgements in decision making (42–45). It may be that some of these regions, in addition to areas involved in working memory, could show functional correlations to the variables we have modeled here. That is, studies leveraging our model (or similar ones) could attempt to track or disrupt neural correlates of

bioRxiv preprint doi: https://doi.org/10.1101/2023.01.17.524436; this version posted January 20, 2023. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC 4.0 International license.

metacognitive variables during motor learning – such as the estimated error state (Equation 3) or metacognitive prediction errors (Equation 4) – using techniques like fMRI and TMS.

We note several limitations in our study. First, simple adaptation tasks should not be conflated with true motor skill learning (20); measuring confidence judgments in more complex motor skill learning tasks will be essential for asking if our models generalize. Second, many models of confidence take a Bayesian approach (12), explicitly modeling sensory uncertainty as a key component of confidence. We took a simpler approach here by focusing on the overall dynamics of confidence judgments during visuomotor learning. Future studies could also incorporate uncertainty in other forms (e.g., sensory feedback, increased motor noise, etc.) to further develop our models in a more probabilistic framework. Third, in Experiment 2 we noticed some surprisingly high-confidence moments that were not easily captured by our models, suggesting that there are likely other biases affecting confidence (46). Modeling these biases will be an important future step as well.

In conclusion, here we show that a simple, straightforward Markovian learning rule was able to capture people's confidence ratings as they adapted to a novel sensorimotor perturbation (Figs. 2-3; Tables 1-2). Our model showed that people's metacognitive judgment of their motor performance, operationalized as explicit confidence judgments in their movement intentions, appeared to incorporate a recency-weighted history of subjectively scaled sensorimotor errors. This model was robust to different learning environments and altered how observed errors influenced metacognition based on the specific statistics of the learning environment (Fig. 4). Our findings provide a foundation for future studies to investigate sensorimotor confidence during more real-world learning tasks, and to localize its correlates in the brain.

Acknowledgements

We thank Ryan Morehead, Jordan Taylor, Huw Jarvis, and the ACT lab for helpful discussions. C.L.H was funded by Yale University's Seesel Postdoctoral Fellowship.

Supplemental Table

[Table S1 here]

bioRxiv preprint doi: https://doi.org/10.1101/2023.01.17.524436; this version posted January 20, 2023. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC 4.0 International license.

599

bioRxiv preprint doi: https://doi.org/10.1101/2023.01.17.524436; this version posted January 20, 2023. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC 4.0 International license.

Bibliography

1. **Pollack I**, **Decker LR**. Confidence Ratings, Message Reception, and the Receiver Operating Characteristic. *J Acoust Soc Am* 30: 286–292, 1958. doi: 10.1121/1.1909571.
2. **Lebreton M**, **Bacily K**, **Palminteri S**, **Engelmann JB**. Contextual influence on confidence judgments in human reinforcement learning. *PLOS Comput Biol* 15: e1006973, 2019. doi: 10.1371/journal.pcbi.1006973.
3. **Heyes C**, **Bang D**, **Shea N**, **Frith CD**, **Fleming SM**. Knowing Ourselves Together: The Cultural Origins of Metacognition. *Trends Cogn Sci* 24: 349–362, 2020. doi: 10.1016/j.tics.2020.02.007.
4. **Koriat A**. How do we know that we know? The accessibility model of the feeling of knowing. *Psychol Rev* 100: 609–639, 1993. doi: 10.1037/0033-295X.100.4.609.
5. **Henmon VAC**. The relation of the time of a judgment to its accuracy. *Psychol Rev* 18: 186–201, 1911. doi: 10.1037/h0074579.
6. **Yeung N**, **Summerfield C**. Metacognition in human decision-making: confidence and error monitoring. *Philos Trans R Soc B Biol Sci* 367: 1310–1321, 2012. doi: 10.1098/rstb.2011.0416.
7. **Sanders JI**, **Hangya B**, **Kepecs A**. Signatures of a Statistical Computation in the Human Sense of Confidence. *Neuron* 90: 499–506, 2016. doi: 10.1016/j.neuron.2016.03.025.
8. **Fleming SM**, **Daw ND**. Self-evaluation of decision-making: A general Bayesian framework for metacognitive computation. *Psychol Rev* 124: 91–114, 2017. doi: 10.1037/rev0000045.
9. **Kiani R**, **Shadlen MN**. Representation of Confidence Associated with a Decision by Neurons in the Parietal Cortex. *Science* 324: 759–764, 2009. doi: 10.1126/science.1169405.
10. **Gonzalez Castro LN**, **Hadjiosif AM**, **Hemphill MA**, **Smith MA**. Environmental Consistency Determines the Rate of Motor Adaptation. *Curr Biol* 24: 1050–1061, 2014. doi: 10.1016/j.cub.2014.03.049.
11. **Meyniel F**, **Schlunegger D**, **Dehaene S**. The Sense of Confidence during Probabilistic Learning: A Normative Account. *PLOS Comput Biol* 11: e1004305, 2015. doi: 10.1371/journal.pcbi.1004305.
12. **Meyniel F**, **Sigman M**, **Mainen ZF**. Confidence as Bayesian Probability: From Neural Origins to Behavior. *Neuron* 88: 78–92, 2015. doi: 10.1016/j.neuron.2015.09.039.
13. **Behrens TEJ**, **Woolrich MW**, **Walton ME**, **Rushworth MFS**. Learning the value of information in an uncertain world. *Nat Neurosci* 10: 1214–1221, 2007. doi: 10.1038/nn1954.
14. **Nassar MR**, **Wilson RC**, **Heasly B**, **Gold JI**. An Approximately Bayesian Delta-Rule Model Explains the Dynamics of Belief Updating in a Changing Environment. *J Neurosci* 30: 12366–12378, 2010. doi: 10.1523/JNEUROSCI.0822-10.2010.
15. **Payzan-LeNestour E**, **Bossaerts P**. Risk, Unexpected Uncertainty, and Estimation Uncertainty: Bayesian Learning in Unstable Settings. *PLOS Comput Biol* 7: e1001048, 2011. doi: 10.1371/journal.pcbi.1001048.
16. **Moskowitz JB**, **Gale DJ**, **Gallivan JP**, **Wolpert DM**, **Flanagan JR**. Human decision making anticipates future performance in motor learning. *PLOS Comput Biol* 16: e1007632, 2020. doi: 10.1371/journal.pcbi.1007632.
17. **Frömer R**, **Nassar MR**, **Bruckner R**, **Stürmer B**, **Sommer W**, **Yeung N**. Response-based outcome predictions and confidence regulate feedback processing and learning. *eLife* 10: e62825, 2021. doi: 10.7554/eLife.62825.
18. **Simon DA**, **Bjork RA**. Metacognition in motor learning. *J Exp Psychol Learn Mem Cogn* 27: 907–912, 2001. doi: 10.1037/0278-7393.27.4.907.

bioRxiv preprint doi: https://doi.org/10.1101/2023.01.17.524436; this version posted January 20, 2023. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC 4.0 International license.

19. **Fassold ME**, **Locke SM**, **Landy MS**. Feeling lucky? Prospective and retrospective cues for sensorimotor confidence. Neuroscience.
20. **Krakauer JW**, **Hadjiosif AM**, **Xu J**, **Wong AL**, **Haith AM**. Motor Learning. *Compr Physiol* 9: 613–663, 2019. doi: 10.1002/cphy.c170043.
21. **Oldfield RC**. The assessment and analysis of handedness: the Edinburgh inventory. *Neuropsychologia* 9: 97–113, 1971.
22. **Bechtold B**, **Fletcher P**, **Seamusholden**, **Gorur-Shandilya S**. bastibe/Violinplot-Matlab: A Good Starting Point. Zenodo: 2021.
23. **Akaike H**. A new look at the statistical model identification. *IEEE Trans Autom Control* 19: 716–723, 1974. doi: 10.1109/TAC.1974.1100705.
24. **Körding KP**, **Wolpert DM**. The loss function of sensorimotor learning. *Proc Natl Acad Sci* 101: 9839–9842, 2004. doi: 10.1073/pnas.0308394101.
25. **Rescorla RA**, **Wagner AR**. A Theory of Pavlovian Conditioning: Variations in the Effectiveness of Reinforcement and Nonreinforcement. In: *Classical conditioning II: current research and theory*. New York: Appleton-Century-Crofts, 1972, p. 64–99.
26. **Thoroughman KA**, **Shadmehr R**. Learning of action through adaptive combination of motor primitives. *Nature* 407: 742–747, 2000. doi: 10.1038/35037588.
27. **Mazzoni P**, **Krakauer JW**. An Implicit Plan Overrides an Explicit Strategy during Visuomotor Adaptation. *J Neurosci* 26: 3642–3645, 2006. doi: 10.1523/JNEUROSCI.5317-05.2006.
28. **McDougle SD**, **Ivry RB**, **Taylor JA**. Taking Aim at the Cognitive Side of Learning in Sensorimotor Adaptation Tasks. *Trends Cogn Sci* 20: 535–544, 2016. doi: 10.1016/j.tics.2016.05.002.
29. **Taylor JA**, **Krakauer JW**, **Ivry RB**. Explicit and implicit contributions to learning in a sensorimotor adaptation task. *J Neurosci Off J Soc Neurosci* 34: 3023–3032, 2014. doi: 10.1523/JNEUROSCI.3619-13.2014.
30. **McDougle SD**, **Taylor JA**. Dissociable cognitive strategies for sensorimotor learning. *Nat Commun* 10, 2019. doi: 10.1038/s41467-018-07941-0.
31. **Morehead JR**, **Xivry J-JO de**. A Synthesis of the Many Errors and Learning Processes of Visuomotor Adaptation. .
32. **Smith JD**, **Coutinho MVC**, **Church BA**, **Beran MJ**. Executive-attentional uncertainty responses by rhesus macaques (Macaca mulatta). *J Exp Psychol Gen* 142: 458–475, 2013. doi: 10.1037/a0029601.
33. **Coutinho MVC**, **Redford JS**, **Church BA**, **Zakrzewski AC**, **Couchman JJ**, **Smith JD**. The interplay between uncertainty monitoring and working memory: Can metacognition become automatic? *Mem Cognit* 43: 990–1006, 2015. doi: 10.3758/s13421-015-0527-1.
34. **Tobler PN**, **O'Doherty JP**, **Dolan RJ**, **Schultz W**. Reward value coding distinct from risk attitude-related uncertainty coding in human reward systems. *J Neurophysiol* 97: 1621–1632, 2007.
35. **Critchley HD**, **Mathias CJ**, **Dolan RJ**. Neural activity in the human brain relating to uncertainty and arousal during anticipation. *Neuron* 29: 537–545, 2001.
36. **Preuschoff K**, **Bossaerts P**, **Quartz SR**. Neural differentiation of expected reward and risk in human subcortical structures. *Neuron* 51: 381–390, 2006.
37. **Ma WJ**, **Jazayeri M**. Neural Coding of Uncertainty and Probability. *Annu Rev Neurosci* 37: 205–220, 2014. doi: 10.1146/annurev-neuro-071013-014017.
38. **Schultz W**, **Preuschoff K**, **Camerer C**, **Hsu M**, **Fiorillo CD**, **Tobler PN**, **Bossaerts P**. Explicit neural signals reflecting reward uncertainty. *Philos Trans R Soc B Biol Sci* 363: 3801–3811, 2008.
39. **Huettel SA**, **Song AW**, **McCarthy G**. Decisions under uncertainty: probabilistic context influences activation of prefrontal and parietal cortices. *J Neurosci* 25: 3304–3311, 2005.
40. **Vilares I**, **Howard JD**, **Fernandes HL**, **Gottfried JA**, **Kording KP**. Differential

bioRxiv preprint doi: https://doi.org/10.1101/2023.01.17.524436; this version posted January 20, 2023. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC 4.0 International license.

representations of prior and likelihood uncertainty in the human brain. *Curr Biol* 22: 1641–1648, 2012.

41. **Ting C-C**, **Yu C-C**, **Maloney LT**, **Wu S-W**. Neural mechanisms for integrating prior knowledge and likelihood in value-based probabilistic inference. *J Neurosci* 35: 1792–1805, 2015.
42. **Cortese A**. Metacognitive resources for adaptive learning. *Neurosci Res* 178: 10–19, 2022. doi: 10.1016/j.neures.2021.09.003.
43. **Fleming SM**, **Weil RS**, **Nagy Z**, **Dolan RJ**, **Rees G**. Relating introspective accuracy to individual differences in brain structure. *Science* 329: 1541–1543, 2010.
44. **Lau HC**, **Passingham RE**. Relative blindsight in normal observers and the neural correlate of visual consciousness. *Proc Natl Acad Sci* 103: 18763–18768, 2006.
45. **Rounis E**, **Maniscalco B**, **Rothwell JC**, **Passingham RE**, **Lau H**. Theta-burst transcranial magnetic stimulation to the prefrontal cortex impairs metacognitive visual awareness. *Cogn Neurosci* 1: 165–175, 2010.
46. **Lebreton M**, **Langdon S**, **Slieker MJ**, **Nooitgedacht JS**, **Goudriaan AE**, **Denys D**, **van Holst RJ**, **Luigjes J**. Two sides of the same coin: Monetary incentives concurrently improve and bias confidence judgments. *Sci Adv* 4: eaaq0668, 2018. doi: 10.1126/sciadv.aaq0668.

bioRxiv preprint doi: https://doi.org/10.1101/2023.01.17.524436; this version posted January 20, 2023. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC 4.0 International license.

Captions

Figure 1. Experimental Design. (A) Experimental apparatus. Participants made reaching movements over a digitizing tablet while holding an air hockey paddle. (B) Schematic of two example trials. Participants first moved a crosshair to specify their intended reach direction ("Aim"), then rated their confidence in that decision ("Report"), and finally executed their reach with feedback ("Reach"). (C) Perturbation schedule in Experiment 1. (D) Perturbation schedule in Experiment 2.

Figure 2. Learning curves and model fitting for Experiment 1 (N=18). (A) Learning curve. Participants adapted their reach angle in response to the 30° perturbation and showed significant aftereffects in a washout phase. (B) Mean confidence reports (gray) and winning model fit (green). (C) *Top*: AIC differences between best and second-best model are shown for each subject (negative values represent better fits for the winning model, ESS_{subj}). *Bottom*: Summed AIC values relative to the summed AIC value for the winning model. All error shading = 1 S.E.M.

Figure 3. Learning curves and model fitting for Experiment 2 (N=20). (A) Learning curve. Participants adapted their reaches (green) to account for the volatile perturbation schedule (gray). (B) Mean confidence reports (gray) and winning model fit (green; ESS_{Subj}). (C) *Top*: AIC differences between best and second-best model are shown for each subject (negative values represent better fits for the winning model, ESS_{Subj}). *Bottom*: Summed AIC values relative to the summed AIC value for the winning model. All error shading = 1 S.E.M.

Figure 4. Parameter comparison across Experiments 1 and 2. Green shaded regions reflect the distribution of participant parameter values. White dots indicate median parameter values and gray bars the interquartile range (IQR) between the first and third quartiles. Whiskers extend to 1.5 times the IQR. (A) Maximum confidence offset parameters were not different across experiments and were generally close to the maximum. (B) Subjective scaling of errors via the exponential parameter was significantly different across experiments. (C) The sensitivity parameter was significantly larger in Experiment 2. (D) Learning rates were significantly larger in Experiment 2, consistent with the hypothesis that learning rates would be higher in the more volatile environment.

Table 1. *Δ*AIC are summed AIC values relative to the winning model. Abbreviations: AIC, Akaike Information Criterion; SD, standard deviation; ESS, error-state-space models; OTB, one-trial-back models.

Table 2. *Δ*AIC are summed AIC values relative to the winning model. Abbreviations: AIC, Akaike Information Criterion; SD, standard deviation; ESS, error-state-space models; OTB, one-trial-back models.

Table T1. *Note.* T-tests are compared to target error model R^2, unless stated otherwise. Abbreviations: df, degrees of freedom, SD, standard deviation.

bioRxiv preprint doi: https://doi.org/10.1101/2023.01.17.524436; this version posted January 20, 2023. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC 4.0 International license.

Figures

Figure 1.

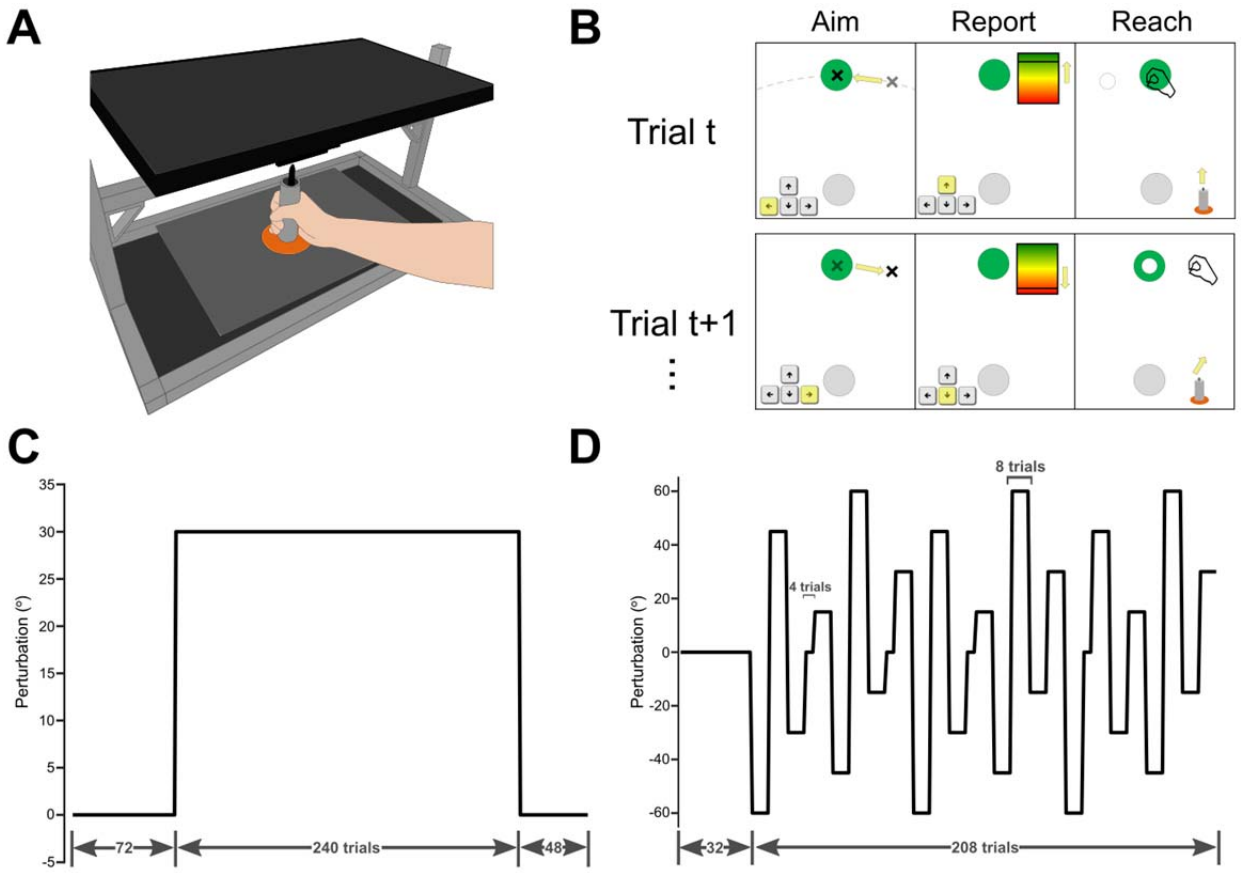

bioRxiv preprint doi: https://doi.org/10.1101/2023.01.17.524436; this version posted January 20, 2023. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC 4.0 International license.

Figure 2.

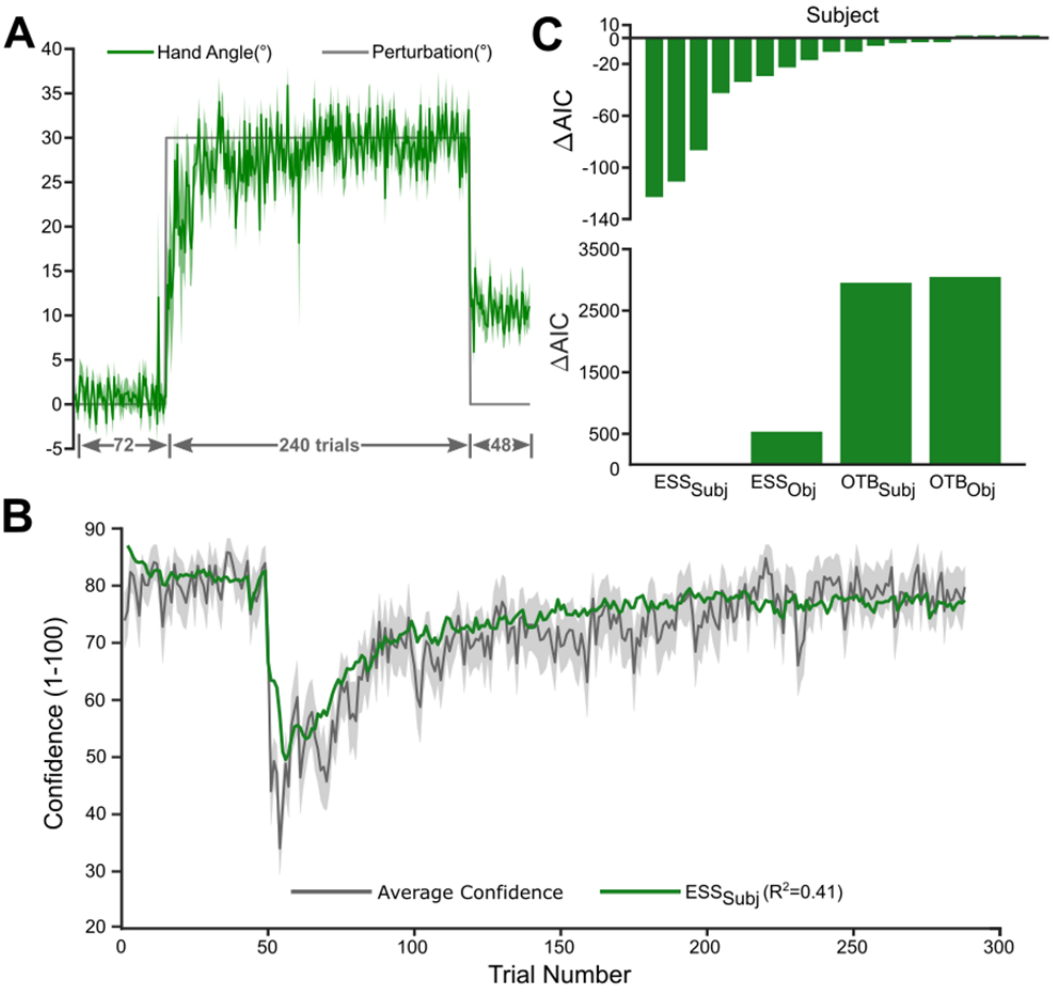

bioRxiv preprint doi: https://doi.org/10.1101/2023.01.17.524436; this version posted January 20, 2023. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC 4.0 International license.

Figure 3.

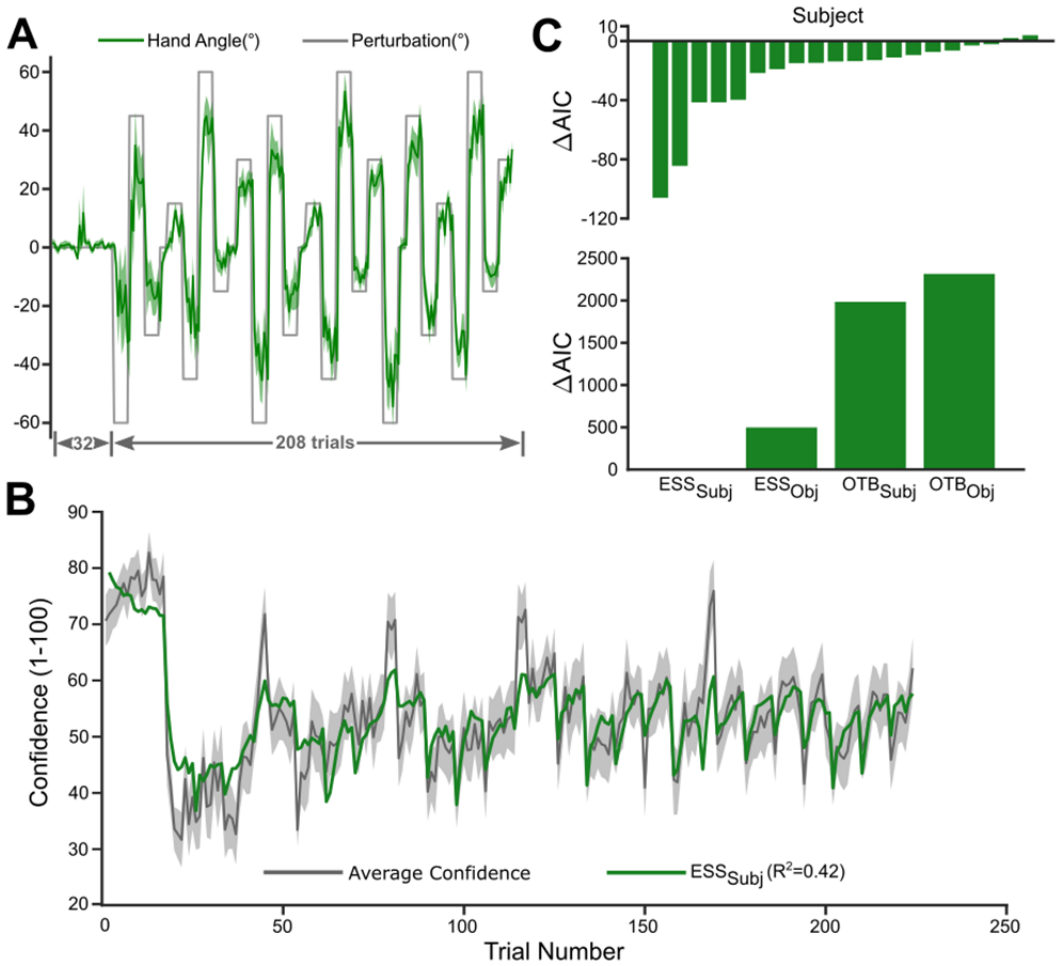

bioRxiv preprint doi: https://doi.org/10.1101/2023.01.17.524436; this version posted January 20, 2023. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC 4.0 International license.

Figure 4.

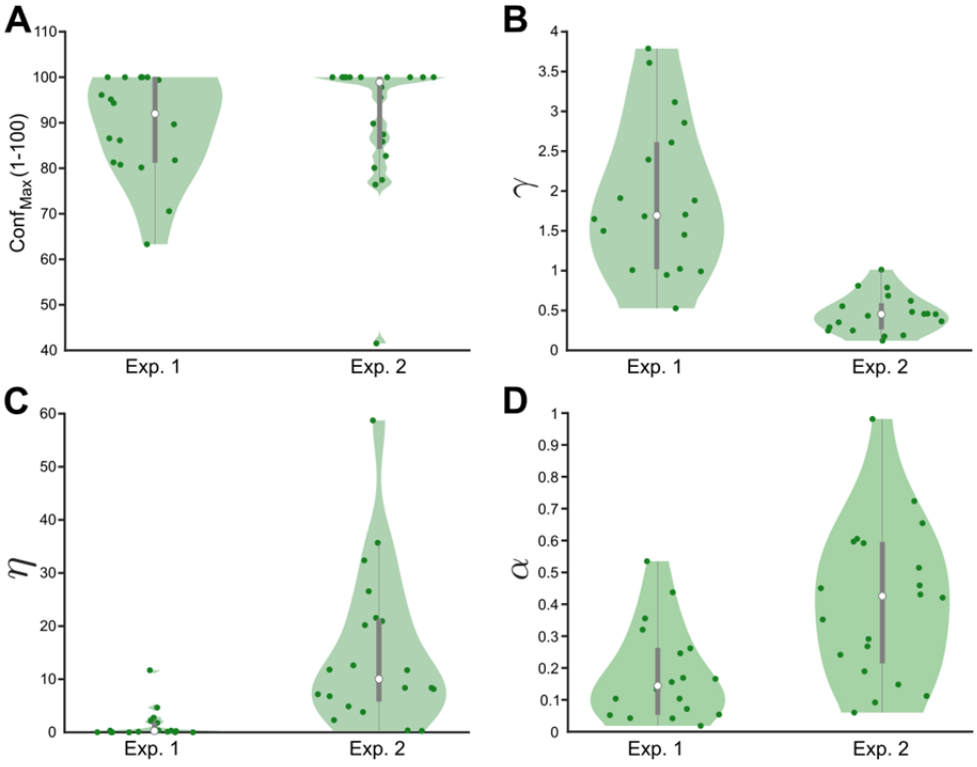

bioRxiv preprint doi: https://doi.org/10.1101/2023.01.17.524436; this version posted January 20, 2023. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC 4.0 International license.

Tables

Table 1. Model Fit R^2 and ΔAIC Values in Experiment 1

Model	R^2 (SD)	ΔAIC
ESS_{subj}	0.4066 (0.2321)	---
ESS_{obj}	0.3587 (0.2272)	525.5
OTB_{subj}	0.1417 (0.0667)	2943
OTB_{obj}	0.1208 (0.0540)	3038

bioRxiv preprint doi: https://doi.org/10.1101/2023.01.17.524436; this version posted January 20, 2023. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC 4.0 International license.

Table 2. Model Fit R^2 and ΔAIC Values in Experiment 2

Model	R^2 (SD)	ΔAIC
ESS_{subj}	0.4207 (0.1974)	---
ESS_{obj}	0.3539 (0.1848)	498.1
OTB_{subj}	0.2373 (0.1009)	1984
OTB_{obj}	0.1741 (0.0830)	2316

bioRxiv preprint doi: https://doi.org/10.1101/2023.01.17.524436; this version posted January 20, 2023. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC 4.0 International license.

Supplemental Table

Table T1. Model Fit R^2 upon Isolating Explicit and Implicit Components

Model	R^2 (SD)	t	df	p
Experiment 1				
Cursor Error (Explicit + Implicit)	0.41 (0.23)	--	--	--
Explicit Component	0.27 (0.26)	-3.59	17	0.002**
Implicit Component	0.15 (0.19)	-5.23	17	6.2×10^{-5}**
Explicit vs Implicit		2.82	17	0.01*
Experiment 2				
Cursor Error (Explicit + Implicit)	0.42 (0.20)	--	--	--
Explicit Component	0.37 (0.21)	-1.61	19	0.12
Implicit Component	0.27 (0.23)	-5.74	19	1.6×10^{-5}**
Explicit vs Implicit		2.34	19	0.03*