nature human behaviour

SUPPLEMENTARY INFORMATION
https://doi.org/10.1038/s41562-019-0725-0

In the format provided by the authors and unedited.

Time-dependent competition between goal-directed and habitual response preparation

Robert M. Hardwick[1,2,4*], Alexander D. Forrence[1], John W. Krakauer[1,3] and Adrian M. Haith[1]

[1]Department of Neurology, Johns Hopkins University, Baltimore, MD, USA. [2]Movement Control & Neuroplasticity Research Group, KU Leuven, Leuven, Belgium. [3]Department of Neuroscience, Johns Hopkins University, Baltimore, MD, USA. [4]Present address: Institute of Neuroscience, UC Louvain, Brussels, Belgium. *e-mail: robert.hardwick@uclouvain.be

Supplementary Figures

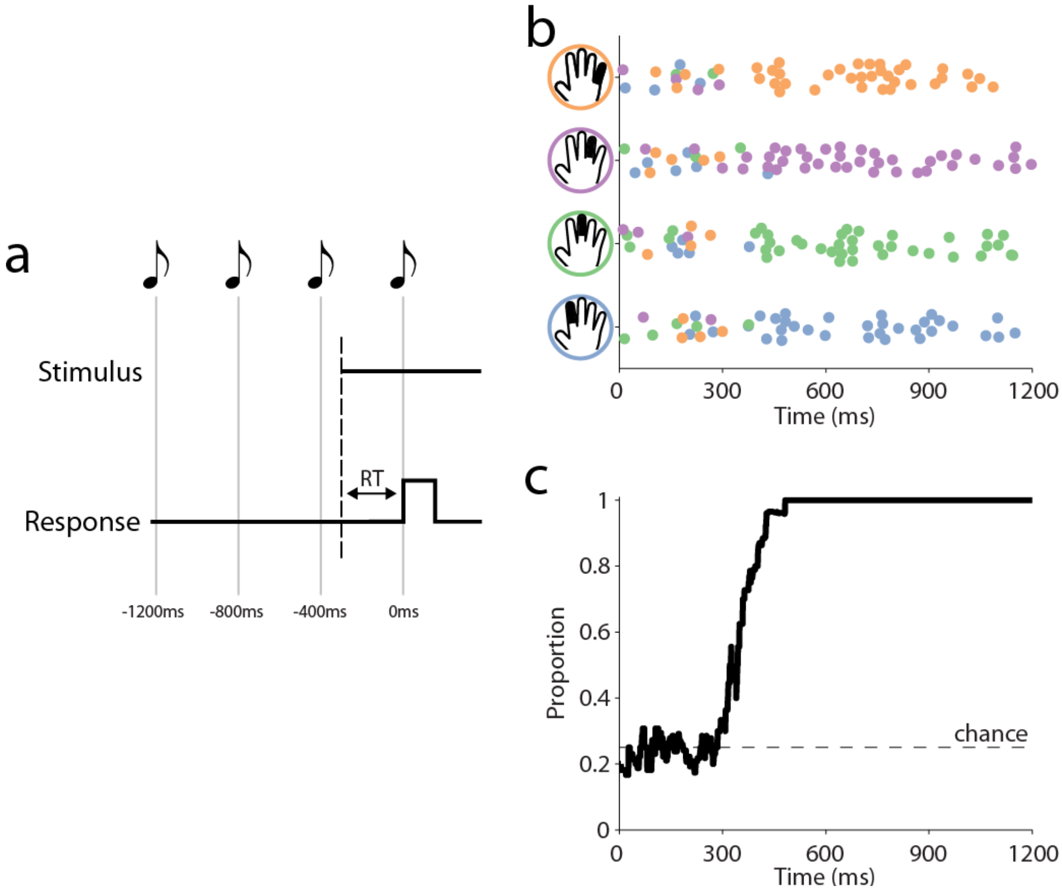

Supplementary Figure 1. Forced response paradigm, example data, and speed-accuracy trade-off. a) Participants heard a sequence of four metronome tones, and were instructed to respond synchronously with the fourth tone. Varying the onset of the stimulus relative to this deadline allowed us to effectively control participant reaction times. b) Data from an example participant on the familiarization task. Participants responded to the appearance of a picture of a hand with a shaded index, middle, ring, or little finger by pressing with the corresponding digit. Illustrations of these stimuli are presented on the y axis (with superimposed colored circles). Responses to each stimulus are shown on the horizontal axis (jitter on the y axis is shown to allow illustration of responses that occurred at the same latency). Circle colors indicate the finger the participant used to respond (e.g. blue circles present responses made with the index finger). c) Data from panel B quantified using a speed-accuracy trade-off. A sliding window (running average across a 100ms window) determined the accuracy for a given time. Note that initial performance is essentially at chance (i.e. at times <300ms participants do not have time to process the stimulus, but must respond synchronously with the fourth tone; therefore, to meet the impending deadline, they select a response at random), after which accuracy increased with time.

Supplementary Figure 2. Experimental Stimuli
As Experiment 1 involved a crossover design, two sets of distinct stimuli (each row presents a different set) were used in a counterbalanced order. The finger to which each stimulus initially corresponded was also randomized in a counterbalanced fashion.

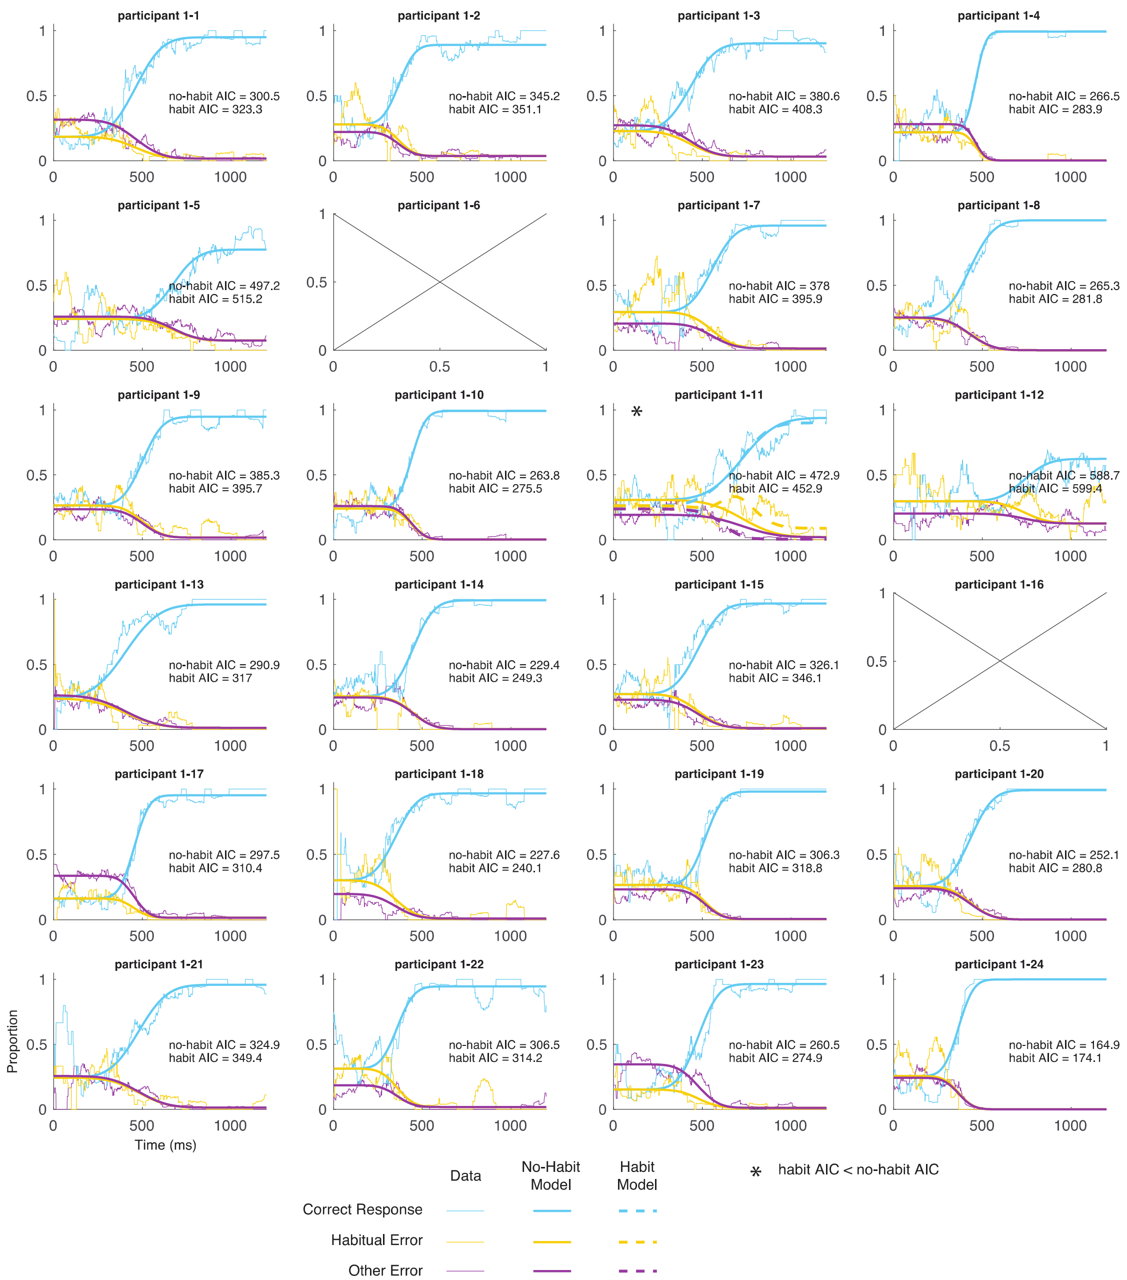

Supplementary Figure 3. Model fits to individual participant data in the Minimal Practice condition Thin lines indicate proportion of different response types, calculated in a 100 ms sliding window (blue = correct response, yellow = habitual error, purple = other error). Text in each panel indicates AIC for each of the three fitted models (no-habit model, and habit model). Thick lines indicate best fits for the no-habit model. For one participant (11, indicated with a *), the no-habit model was outperformed by the habit model. For this participant, the best fit for the habit model is shown as a thick dashed line. X denotes participants that withdrew from the study.

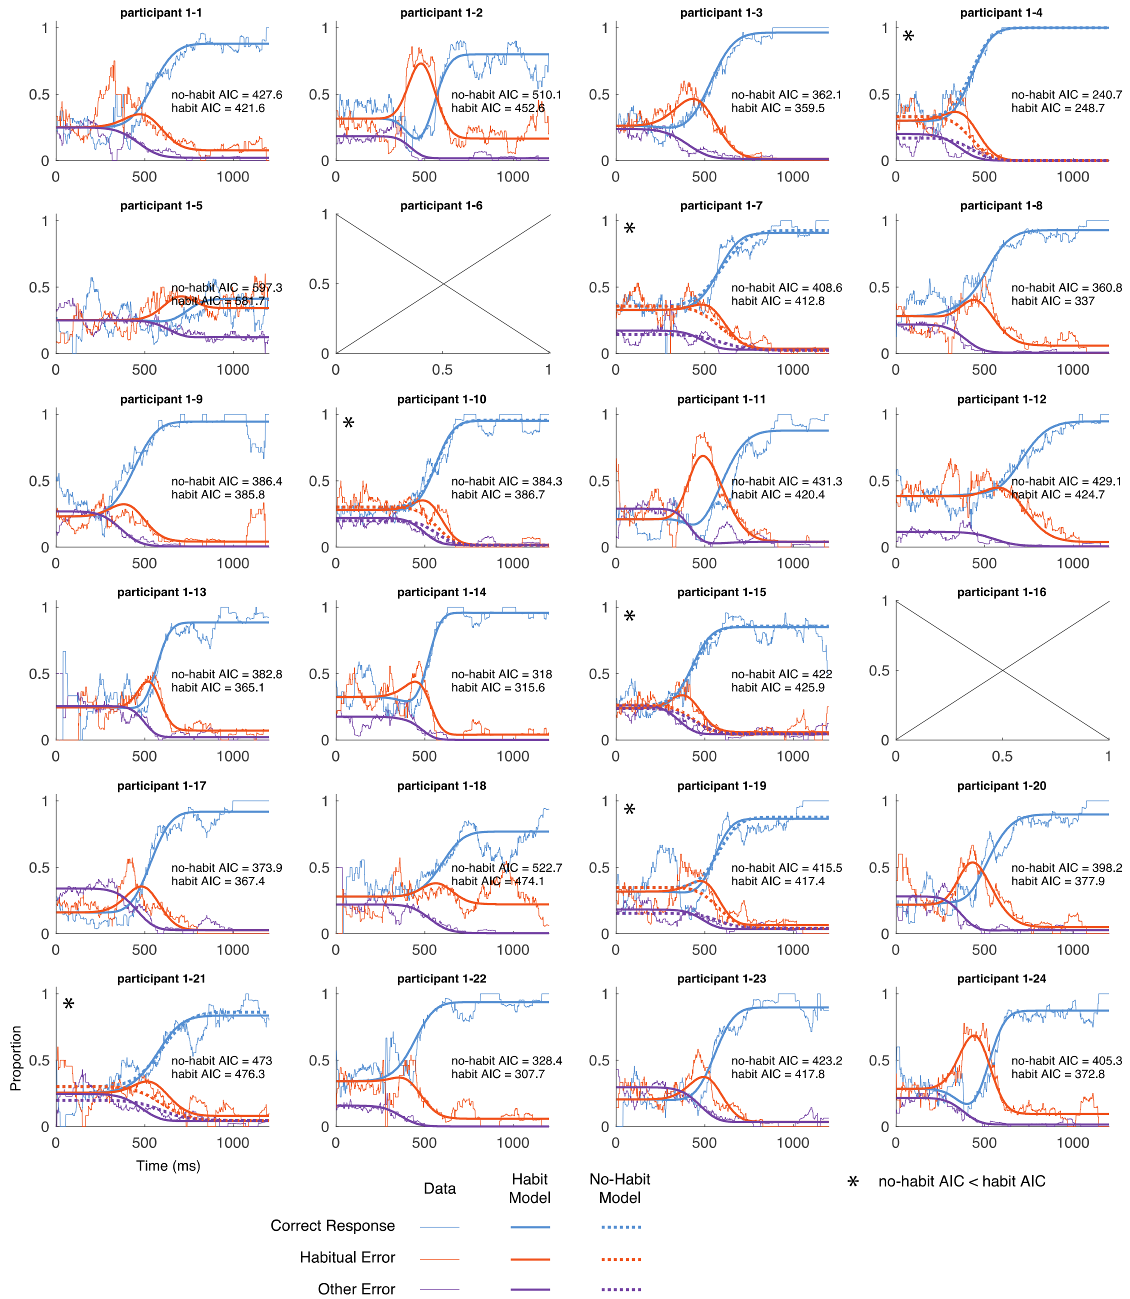

Supplementary Figure 4. Model fits to individual participant data in the 4-Day Practice condition Thin lines indicate proportion of different response types, calculated in a 100 ms sliding window (blue = correct response, orange = habitual error, purple = other error). Text in each panel indicates the AIC for each fitted model. Thick lines indicate best fits for the habit model. * indicates participants for which the no-habit model outperformed the habit model. For these participants, the best fit for the no-habit model is shown as a dotted line. X denotes participants that withdrew from the study.

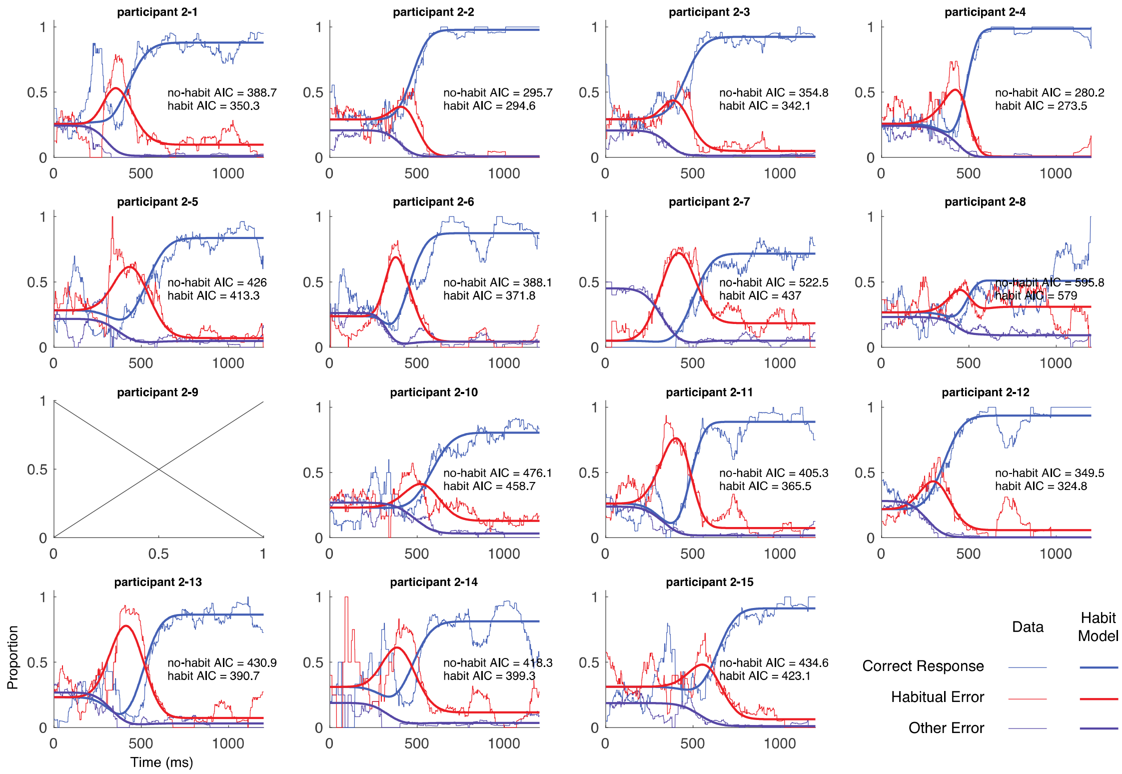

Supplementary Figure 5. Model fits to individual participant data in the 20-Day Practice condition. Thin lines indicate proportion of different response types, calculated in a 100 ms sliding window (blue = correct response, orange = habitual error, purple = other error). Text in each panel indicates AIC for each of the two fitted models. Thick lines indicate best fits for the **habit** model. The habit model outperformed the no-habit model for all participants in the 20-Day Practice condition. X denotes participants that withdrew from the study.

Our data demonstrate that, between the 4-Day practice condition and the 20-Day practice condition, there was a significant increase in the peak probability of expressing a habitual error (Supplementary Figure 6a). To assess whether the increase in selection speed of the originally practiced mapping could account for the increase in the probability of a habitual error, we attempted to predict behavior in the 20-day practice condition based on observed behavior after 4 days of practice combined with observed skill level after 20 days. For our prediction of behavior in the 20-day condition we thus set parameters μ_A and σ_A based on fits to behavior on the unchanged stimuli (Supplementary Figure 6b, purple lines), and set all other parameters ($\mu_B, \sigma_B, q_I, q_B, \rho_B$) based on fits to the 20-day condition. The resulting predictions closely matched the observed behavior, if anything overestimating the increase in the peak probability of a habitual response. This suggested that improved speed of selecting the habitual response accounted for the increased probability of a habitual response.

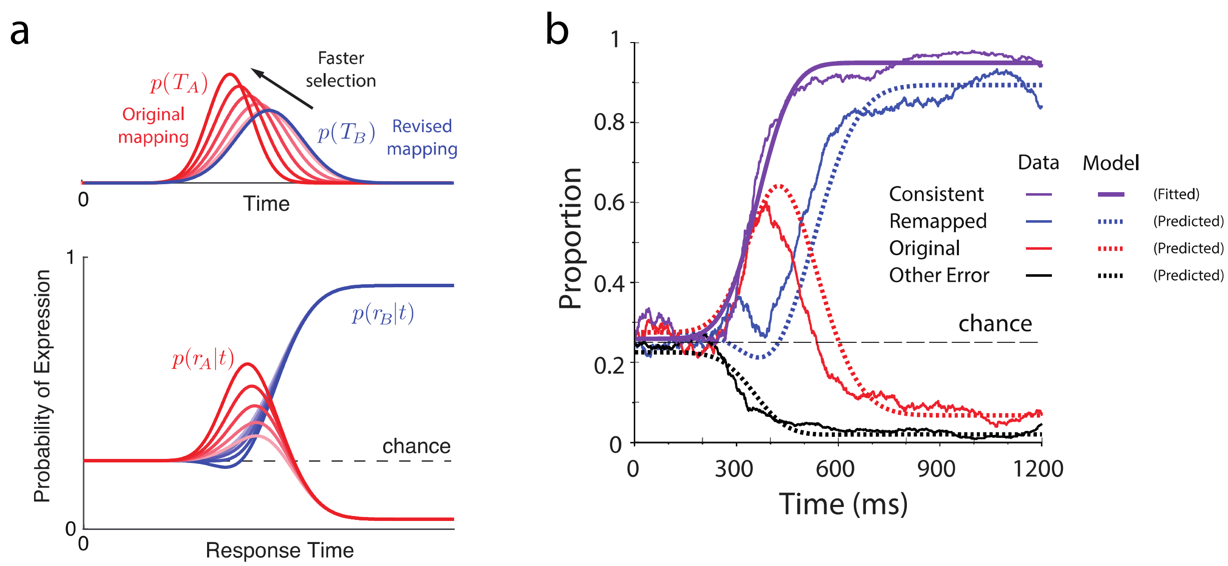

Supplementary Figure 6. Effect of response speed on habit expression: Skill improvement accounts for increased proportion of habitual errors

a) Under the habitual response selection model, more rapid selection of the originally practiced mapping (top panel) leads to a greater peak probability of generating a habitual response (bottom panel). Here darker shades of red indicate earlier response times (lower mean and variance of T_A). b) Predicted behavior in the 20-day Practice condition based on observed improvements in response speed. Thin curves indicate proportion of different responses for consistently mapped stimuli (purple = correct responses) and remapped stimuli (blue = correct, red = habitual error, black = other error). Solid purple line shows model fit to behavior for the consistently mapped stimuli. Dotted lines show predictions for remapped stimuli based on model fits to the 4-Day Practice condition, combined with selection speed after 20 days estimated from consistently mapped stimuli (purple lines).

In order to validate our ability to fit each model to data and distinguish between different models, we conducted a parameter recovery and model recovery analysis (Wilson & Collins, 2019). We first established the extent to which we could recover known model parameters from a synthetic dataset. We fit the no-habit model to data for each participant in the Minimal Practice condition, and then used each fitted model to generate 10 synthetic datasets, maintaining the same set of response times generated by each participant, but sampling each response from the conditional response distribution determined by the model. We then fit the model to this synthetic data via maximum likelihood estimation. In almost all cases, the recovered parameters were well correlated with the parameters used to generate the synthetic data (Supplementary Figure 7a). The only parameter that was not consistently recovered were the variance associated with timing of response preparation, i.e. σ_A. The exact value of this parameter had little qualitative effect on the behavior. We repeated this procedure for the habit model and the flex-habit model, in this case using models fitted to behavior from the 4-Day practice condition and the 20-Day practice condition to generate the synthetic datasets. As for the no-habit model, parameter recovery was strong (Supplementary Figure 7b).

More critical to our conclusions is whether or not we can reliably discriminate which model generated a particular dataset, i.e. model recovery. We fit the habit and no-habit models to synthetic datasets generated by the no-habit model (seeded by model fits to data from the Minimal Practice Condition), and compared the AIC for each model to select the best model. We did this 100 times for each condition, for each participant and recorded the proportion of datasets for which the wrong model was selected to generate a confusion matrix[1]. The no-habit model was reliably recovered for data simulated based on behavior in all three conditions (Supplementary Figure 7c). We did likewise with datasets generated using the habit model. For synthetic datasets seeded from the Minimal Practice condition, recovery of the habit model typically failed, since the two models predicted similar behavior and model comparison therefore favored the more parsimonious no-habit model. The habit model was, however, reliably recovered for synthetic datasets seeded from behavior in the 4-Day and 20-Day practice conditions. The only participants for which the habit model was not reliably recovered were those for which the initial model comparison favored the no-habit model.

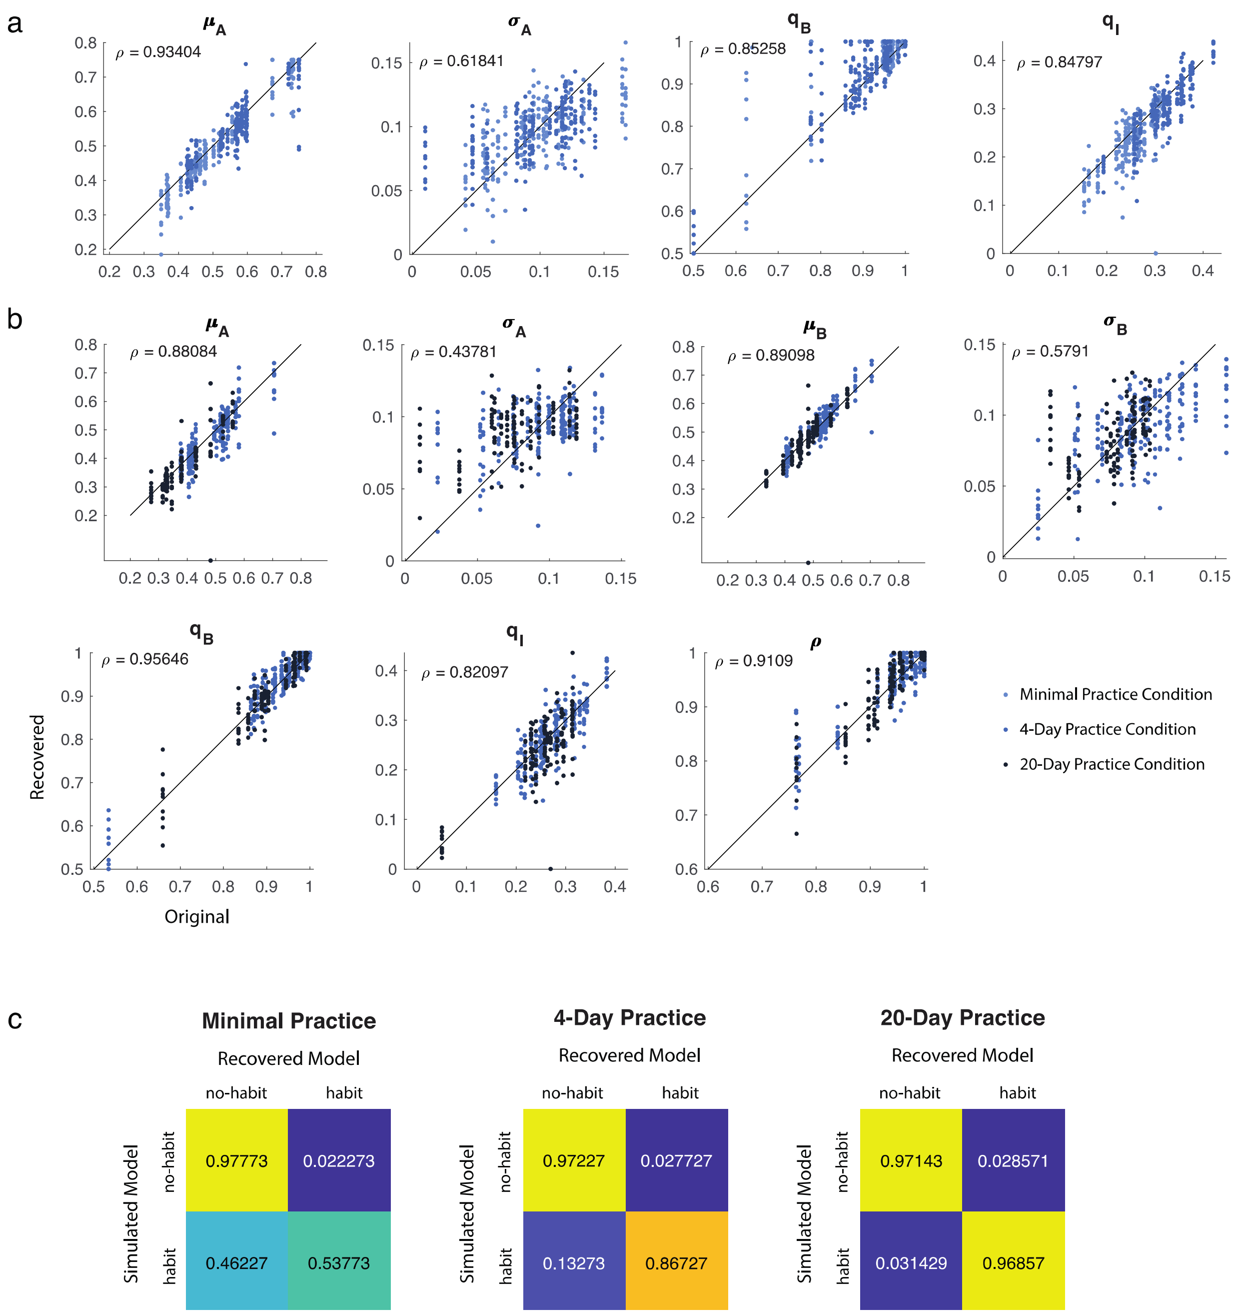

Supplementary Figure 7. Model recovery analysis

a) Recovery of parameters from data simulated using the no-habit model, seeded from model fits to the Minimal Practice and 4-Day Practice conditions. Each panel plots parameter estimates from simulated data (y-axis) against the actual parameter value used to generate the data (x-axis). Inset shows the correlation between original and recovered parameters. b) Analogous parameter recovery for the habit model, in this case seeded from model fits to the 4-Day Practice and 20-Day Practice conditions. c) Confusion matrices for model selection. Each element of the confusion matrix shows how often a particular model had the lowest AIC, for data generated using each model. A diagonal matrix indicates reliable model recovery. Model recovery was conducted separately for datasets seeded from each condition.

Supplementary References

1. Wilson, R. C. & Collins, A. *Ten simple rules for the computational modeling of behavioral data*. (PsyArXiv, 2019). doi:10.31234/osf.io/46mbn