Trial-by-trial adaptation of decision making performance - a model-based EEG analysis

van Vugt, M. K.; Simen, P.; Cohen, J. D.

Trial-by-Trial Adaptation of Decision Making Performance - A Model-Based EEG Analysis

van Vugt, M.[1], Simen, P.[2], Cohen, J.D.[2]

[1] University of Groningen, The Netherlands, ² Princeton University, New Jersey, USA

Theoretical Background

Whenever we engage in a task, it is crucial we monitor our performance to make sure that it does not decline. When it does decline, the performance monitoring system takes action to remedy that, e.g., by slowing down responding (Laming, 1979; Rabbitt, 1966). Botvinick, Braver, Barch, Carter, and Cohen (2001) proposed that this slowing of response times (RTs) reflected conflict that stirred a performance monitoring system to action. There exist various decision systems that adapt performance (Daw, Niv, & Dayan, 2005), each of which have been associated with specific neural correlates. Some of the behavioral adjustment is thought to be implemented by the medial frontal cortex (MFC), thought to implement reinforcement learning mechanisms for behavioral adjustment (Cohen & Ranganath, 2007). This is contrasted to the striatum that is thought to implement rule-based behavioral adjustments.

Activity of the MFC in EEG (electroencephalography) is typically associated with two components: the error-related negativity (ERN) in the first 100 ms after a response, and the feedback-related negativity (FRN) about 200–400 ms post-response. Both components are more negative after errors when compared to correct trials, and opinions differ about what cognitive processes they reflect. The ERN covaries with individual differences in personality traits, e.g., with a participants' tendency to learn more from negative than from positive feedback (Frank, D'Lauro, & Curran, 2007). The FRN covaries with valence of the feedback but not so much with its magnitude (Hajcak, Moser, Holroyd, & Simons, 2006), although this has been disputed (Bellebaum & Daum, 2008). The ERN and FRN have both been related to reinforcement learning (RL) models of learning in decision making, and are thought to reflect prediction errors. Although good at describing across-trial dynamics, these RL models do not describe within-trial dynamics.

While Drift Diffusion Models (DDMs; Ratcliff, 1978) of decision making do not describe learning from successive decisions, they do describe within-trial dynamics. We have begun to relate electrophysiological data to DDMs (van Vugt et al, in preparation), and we wondered whether the ERN and/or FRN could have a role in setting the speed-accuracy trade-off (SAT) in a perceptual decision making task. Answering this question could elucidate how the ERN and FRN, which are fairly similar potentials, differ (e.g., Heldmann, Rüsseler, & Münte, 2008), and how these across-trial dynamics relate to the within-trial dynamics described by the DDM.

Drift diffusion models describe decision processes as random walks toward thresholds corresponding to the response alternatives, e.g., "left" and "right". As soon as a threshold is crossed, the participant emits the corresponding response. The drift rate of the random walk reflects the quality of the data on which the decision is based, whereas the decision threshold is under the participant's control and reflects SATs. When a participant increases her decision threshold, she will accumulate more evidence before responding, i.e., act more conservatively. Her RT consists of the time it takes to reach the decision threshold, together with a non-decision time, reflecting fixed perceptual and motor delays.

We predict that behavioral adjustments after errors are reflected in an increased decision threshold. Furthermore, if the ERN/FRN not only reflects error awareness but also commitment to behavioral adjustment, then we predict that the magnitude of this component predicts the magnitude of the threshold adjustment. Finally, the magnitude of the ERN within a subject should predict RTs on the next trial.

Methods

Task: Participants performed a perceptual decision making task in which they judged the direction of motion of a display of randomly moving dots, a subset of which moved coherently to the left or the right. These random dot kinematograms were similar to those used in a series of psychophysical and decision making experiments involving monkeys (e.g., Britten, Shadlen, Newsome, & Movshon, 1992; Gold & Shadlen, 2001; Shadlen & Newsome, 2001). Stimuli consisted of an aperture of approximately 7.6 cm diameter viewed from approximately 100 cm (approximately 4 degrees visual angle) in which white dots (2 x 2 pixels) moved on a black background. A subset of dots moved coherently either to the left or to the right on each trial, whereas

the remainder of dots were distractors that jumped randomly from frame to frame. Motion coherence was defined as the percentage of coherently moving dots. Dot density was 17 dots/square degrees, selected so that individual dots could not easily be tracked. Following the procedure used in Simen, Contreras, Hu, Holmes, and Cohen (2009), stimuli remained visible until participants made a response, at which point the stimulus disappeared and a variable response-to-stimulus interval ensued. Correct responses were rewarded with $0.01, errors were unrewarded. Reward feedback was displayed visually and signaled with a tone after each trial.

We manipulated response bias and response-to-stimulus interval (RSI; see Simen et al. (2009) for a review of the effects of these variables on behavioral performance in two-alternative forced-choice tasks with response-terminated stimuli, and a comparison to the predictions of Bogacz, Brown, Moehlis, Holmes, and Cohen (2006)). Response bias was manipulated by changing the probability that the dots would move in one of the two directions from 0.5 (no bias) to 0.6, 0.75 and 0.9. For the purposes of this analysis, we collapsed across the two RSI levels. The experiment presentation code was written in PsychToolbox (Brainard, 1997). Dot stimuli were presented with PsychToolbox extensions written by J. I. Gold (http://code.google.com/p/dotsx/).

Participants: Twenty-one members of the Princeton Community (15 female, mean age 20.1) participated in our experiment in exchange for payment. The experiment was approved by the Institutional Review Board of Princeton University. Subjects participated in four separate hour-long training sessions in which they became familiar with the task. At the end of these training sessions, we used performance on a psychometric block to determine the coherence at which they performed at approximately 80% correct. This coherence level was used for the two hour-long EEG sessions.

Recording Methods: We recorded EEG data from 128 channels using Neuroscan EEG caps with a Sensorium EPA-6 amplifier. All data were referenced to the average of the mastoids and off-line rereferenced to an average reference after automatic bad-channel removal (Friederici, Wang, Herrmann, Maess, & Oertel, 2000; Hestvik, Maxfield, Schwartz, & Shafer, 2007).

Data Analysis: We fitted the behavioral data to the Drift Diffusion Model using the DMA toolbox (VandeKerckhove & Tuerlinckx, 2007). We computed the error- and feedback-related negativity in electrode Fz, where it

is known to be maximal, on 12-Hz low-pass filtered data (Holroyd, Nieuwenhuis, Yeung, & Cohen, 2003). We baseline corrected each trial to the 100 ms prior to the response. The ERN was taken to be the average of the signal between 0–100 ms post-response. The FRN was taken to be the difference between the positive peak in the window 120–200 ms post-response and the first negative peak thereafter, but before 325 ms (Holroyd et al., 2003). This peak window was determined from visual inspection of the ERN waveforms.

Results

We first asked whether participants adapted their decision threshold height after an error on the previous trial. Figure 1 shows that as predicted, participants adapt their decision thresholds after they commit an error, especially for trials in the highly biased conditions. Post-correct and post-error thresholds differ from each other for the 0.75 and 0.9 bias conditions [$t(20) = 2.41$, $p = 0.026$ and $t(20) = 5.86$, $p < 0.001$, respectively], marginally for the 0.6 bias [$t(20) = 1.77$, $p = 0.092$] but not in the no-bias condition [$t(20) = 1.039$, *n.s.*]).
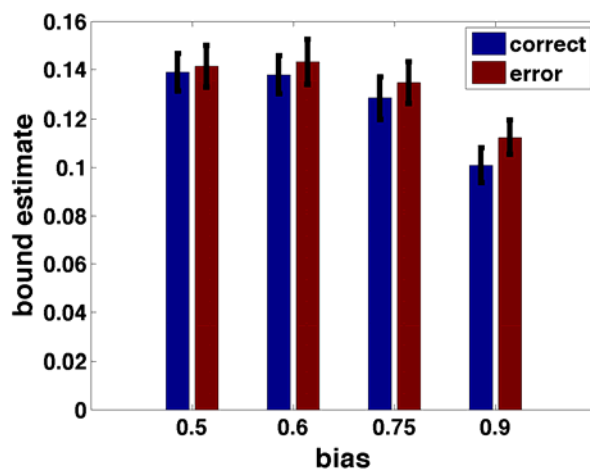

Fig.1   Fitted DDM decision thresholds in the trials after a correct (blue) and incorrect (red) response. The thresholds are shown separately for the different response bias levels. Error bars reflect standard error of the mean.
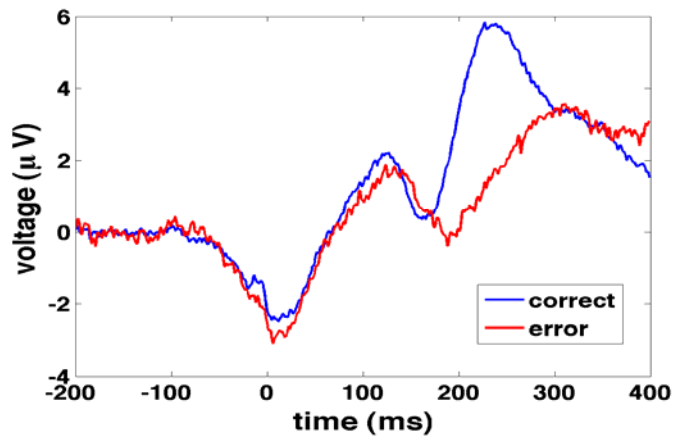
144

Fig.2   Error-related negativity for the 0.75 bias condition (the condition that shows a modulation of the decision threshold).
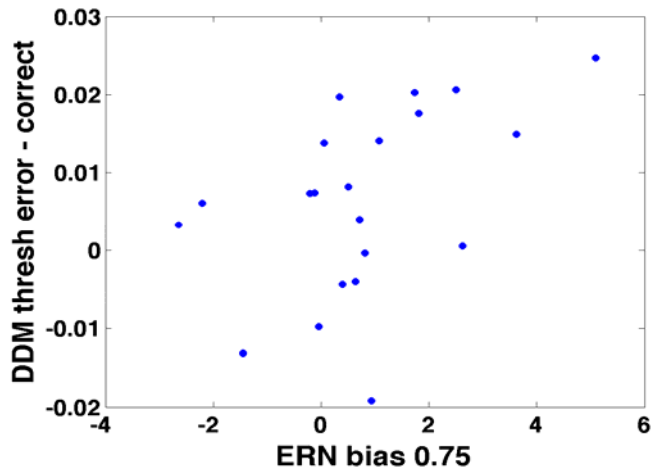

Fig.3   Individual differences in threshold adaptation correlate with the magnitude of the ERN in the 0.75 bias condition [$r(19) = 0.47$, $p = 0.032$].

We then examined the ERN magnitude to replicate the basic ERN result. Figure 2 shows that in fact there is an effect of errors in our data: there is both a difference between correct and error trials across all conditions in the

145

ERN window [0–100 ms post-response; $t(83) = 2.13$, $p = 0.036$] and in the FRN window [approximately 200-400 ms post-response; $t(83) = 3.91$, $p < 0.001$].

Next, we asked whether the magnitude of the ERN correlates with individual differences in adaptation of the decision threshold. Figure 3 shows that for the 0.75 bias condition (which showed a significant difference in decision threshold between post-correct and post-error trials), there is indeed such a relation. For the other bias levels, this relation is suggestive, though not significant. In contrast to the ERN, there was no relation between the later FRN (appearing around 200-400 ms) and decision threshold adaptation [all correlations $p > 0.17$], except for a negative correlation between change in bias level and FRN magnitude at 0.9 [$r(20) = -0.56$, $p = 0.0081$], which we think may be spurious because this correlation is negative, rather than positive, as we would expect. Moreover, in the 0.9 bias condition participants typically resort to non-integrative responding (Simen et al., 2009).
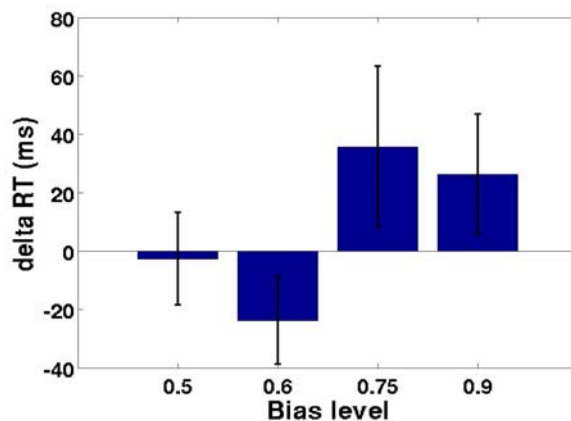

Fig.4   Increases in the ERN are associated with a longer RT in the next trial ($p = 0.0291$; sign-rank test). Error bars reflect standard error of the mean.

If the ERN is related to decision threshold adjustments, then that should have consequences for behavior. We therefore asked whether the magnitude of the ERN in the 0.75 bias condition could predict RT on the next trial. In other words: an increase in the decision threshold should result in a longer RT. To assess this, we split each participant's data in two halves based on ERN magnitude. We then computed the average RT in the small ERN and large

146

ERN trials for every participant. Figure 4 shows that as the ERN increases, the average RT will tend to increase for the 0.75 bias level. Although this is different from predicting RT on a single trial, it is, given the noisy nature of EEG, nevertheless a step in the right direction.

Conclusions & Future Directions

We have shown that after an error, the decision threshold increases. At the intermediate bias level of 0.75 for which the thresholds differed significantly between post-correct and post-error trials, this increase in decision threshold was associated with individual differences in the magnitude of the ERN (but not the later FRN). The ERN may therefore reflect consequences for performance in the next trial through an adjustment of the decision threshold, and this process of performance adjustment occurs very early on after error detection. This further supports findings by Bellebaum, Kobza, Thiele, and Daum (2010) that the ERN is modulated by the degree of learning people do from task performance. It also agrees with Gentsch, Ullsperger, and Ullsperger (2009), who found that the ERN was associated with an adjustment of performance, whereas the FRN was not. Post-error behavioral adjustment effects have previously been described in the N2 component in a flanker task (Forster, Carter, Cohen, & Cho, 2010). It remains to be determined whether N2 modulations also occur in our perceptual decision task, and if so, whether they too relate to decision threshold adaptations.

One may wonder why we only observe decision threshold adjustments and ERN effects at the intermediate bias level (0.75), where more than half of the trials the stimuli move in a biased direction. It may be the case that participants are more focused on learning in these trials, because they are trying to figure out whether there is a bias, and how they should adapt to that. When there is no bias, participants simply focus on the stimuli; when there is a large bias participants resort to pre-emptive responding and do not pay too much attention to the stimulus (Simen et al., 2009).

Our results have potentially important applications in the field of education research. The demonstrated relation between the ERN, one of the most robust electrophysiological indices, and adjustments in cognitive control, could be used to assess the effectiveness of educational interventions. Many of these interventions involve monitoring current behavior and adjusting it as necessary (i.e., executive function), making the ERN a viable target for assessing its effects on the human brain.

147

Additionally, our work forms a bridge between cognitive modeling and cognitive neuroscience, by giving a formal description of the ERN in terms of model components. This link between the DDM and the ERN could be used to make strong predictions about the effect of behavioral manipulations on the ERN. Finally, this work suggests that it may be useful to extend the DDM to account for sequential effects, such that it can explain the post-error threshold adjustment that we observed.

Literature

Bellebaum, C., & Daum, I. (2008). Learning-Related Changes in Reward Expectancy Are Reflected in the Feedback-Related Negativity. *European Journal of Neuroscience*, 27(7), 1823–1835.

Bellebaum, C., Kobza, S., Thiele, S., & Daum, I. (2010). It Was Not MY fault: Event-Related Brain Potentials in Active and Observational Learning from Feedback. *Cerebral Cortex*, 20, 2874–2883.

Bogacz, R., Brown, E., Moehlis, J., Holmes, P., & Cohen, J. D. (2006). The Physics of Optimal Decision Making: A Formal Analysis of Models of Performance in Two-Alternative Forced-Choice Tasks. *Psychological Review*, 113(4), 700–765.

Botvinick, M. M., Braver, T. S., Barch, D. M., Carter, C. S., & Cohen, J. D. (2001). Conflict Monitoring and Cognitive Control. *Psychological Review*, 108(3), 624–652.

Brainard, D. H. (1997). The Psychophysics Toolbox. *Spatial Vision*, 10, 443-446.

Britten, K. H., Shadlen, M. N., Newsome, W. T., & Movshon, J. A. (1992). The Analysis of Visual Motion: A Comparison of Neuronal and Psychophysical Performance. *Journal of Neuroscience*, 12(12), 4745–4765.

Cohen, M. X., & Ranganath, C. (2007). Reinforcement Learning Signals Predict Future Decisions. *Journal of Neuroscience*, 27, 371–378.

Daw, N. D., Niv, Y., & Dayan, P. (2005). Uncertainty-Based Competition between Prefrontal and Dorsolateral Striatal Systems for Behavioral Control. *Nature Neuroscience*, 8, 1704–1711.

Forster, S. E., Carter, C. S., Cohen, J. D., & Cho, R. Y. (2010). Parametric Manipulation of the Conflict Signal and Control-State Adaptation. *Journal of Cognitive Neuroscience*, 23(4), 923–935.

Frank, M. J., D'Lauro, C., & Curran, T. (2007). Cross-Task Individual Differences in Error Processing: Neural, Electrophysiological, and Genetic Components. *Cognitive, Affective, & Behavioral Neuroscience*, 7(4), 297–308.

Friederici, A. D., Wang, Y., Herrmann, C. S., Maess, B., & Oertel, U. (2000). Localization of Early Syntactic Processes in Frontal and Temporal Cortical Areas: A Magnetoencephalographic Study. *Human Brain Mapping*, 11(1), 1–11.

Gentsch, A., Ullsperger, P., & Ullsperger, M. (2009). Dissociable Medial Frontal Negativities from a Common Monitoring System for Self- and Externally Caused Failure of Goal Achievement. *NeuroImage*, 47, 2023–2030.

Gold, J. I., & Shadlen, M. N. (2001). Neural Computations that Underlie Decisions about Sensory Stimuli. *Trends in Cognitive Science*, 5(1), 10–16.

Hajcak, G., Moser, J. S., Holroyd, C. B., & Simons, R. F. (2006). The Feedback-Related Negativity Reflects the Binary Evaluation of Good Versus Bad Outcomes. *Biological Psychology*, 71(2), 148–154.

Heldmann, M., Rüsseler, J., & Münte, T. F. (2008). Internal and External Information in Error Processing. *BMC Neuroscience*, 9, 33.

Hestvik, A., Maxfield, N., Schwartz, R. G., & Shafer, V. (2007). Brain Responses to Filled Gaps. *Brain and Language*, 100(3), 301–316.

Holroyd, C. B., Nieuwenhuis, S., Yeung, N., & Cohen, J. D. (2003). Errors in Reward Prediction Are Reflected in the Event-Related Brain Potential. *NeuroReport*, 14(18), 2481–2484.

Laming, D. (1979). Choice Reaction Performance Following an Error. *Acta Psychologica*, 43, 199-224.

Rabbitt, P. M. A. (1966). Errors and Error Correction in Choice-Response Tasks. *Journal of Experimental Psychology*, 71(2), 264–272.

Ratcliff, R. (1978). A Theory of Memory Retrieval. *Psychological Review*, 85, 59-108.

Shadlen, M. N., & Newsome, W. T. (2001). Neural Basis of a Perceptual Decision in the Parietal Cortex (area LIP) of the Rhesus Monkey. *Journal of Neurophysiology*, 86(4), 1916–1936.

Simen, P., Contreras, C., D.and Buck, Hu, P., Holmes, P., & Cohen, J. D. (2009). Reward-Rate Optimization in Two-Alternative Decision Making: Empirical Tests of Theoretical Predictions. *Journal of Experimental Psychology: Human Perception and Performance*, 35(6), 1865–1897.

VandeKerckhove, J. A., & Tuerlinckx, F. (2007). Fitting the Ratcliff Diffusion Model to Experimental Data. *Psychonomic Bulletin & Review*, 14(6), 1011–1026.

Acknowledgements