Inhibition of Pre–Supplementary Motor Area by Continuous Theta Burst Stimulation Leads to More Cautious Decision-making and More Efficient Sensory Evidence Integration

Tuğçe Tosun[1], Dilara Berkay[2], Alexander T. Sack[3], Yusuf Ö. Çakmak[4], and Fuat Balcı[1]

Abstract

■ Decisions are made based on the integration of available evidence. The noise in evidence accumulation leads to a particular speed–accuracy tradeoff in decision-making, which can be modulated and optimized by adaptive decision threshold setting. Given the effect of pre-SMA activity on striatal excitability, we hypothesized that the inhibition of pre-SMA would lead to higher decision thresholds and an increased accuracy bias. We used offline continuous theta burst stimulation to assess the effect of transient inhibition of the right pre-SMA on the decision processes in a free-response two-alternative forced-choice task within the drift diffusion model framework. Participants became more cautious and set higher decision thresholds following right pre-SMA inhibition compared with inhibition of the control site (vertex). Increased decision thresholds were accompanied by an accuracy bias with no effects on post-error choice behavior. Participants also exhibited higher drift rates as a result of pre-SMA inhibition compared with the vertex inhibition. These results, in line with the striatal theory of speed–accuracy tradeoff, provide evidence for the functional role of pre-SMA activity in decision threshold modulation. Our results also suggest that pre-SMA might be a part of the brain network associated with the sensory evidence integration. ■

INTRODUCTION

We routinely make various simple perceptual and economic decisions such as choosing between different meal options in the cafeteria or deciding to either pass a slower vehicle in front of us now or rather wait for a safer moment. The outputs of such simple decisions have been successfully accounted for in a unified fashion by sequential sampling models of decision-making (Ratcliff, Smith, Brown, & McKoon, 2016). For instance, neurophysiological and psychological data suggest that, during two-alternative forced-choice (2AFC) perceptual decision-making, the brain integrates sensory evidence supporting one alternative over the other over time before making a choice (Shadlen & Newsome, 2001; Ratcliff, 1978; Laming, 1968). This integration process is required for accurate decisions due to the limited reliability of sensory evidence, stemming from the noise in the sensory input and/or its transduction/processing.

The decision outputs gathered from 2AFC tasks can be successfully modeled with computational models such as the drift diffusion model (DDM; e.g., Bogacz, 2007; Gold & Shadlen, 2007; Smith & Ratcliff, 2004). The DDM implements the optimum decision procedure for the 2AFC behavior (Laming, 1968) and decomposes it to estimate various decision parameters, which represent different latent psychological processes/variables. It assumes that (a) the difference between the evidence from the noisy sources of sensory information supporting each of the two alternatives is integrated over time and (b) when the accumulated information reaches one of the decision thresholds representing the two response alternatives, either above or below the initial belief state (starting point), the corresponding option is chosen. The RT of choice behavior is determined by the first threshold crossing time (Ratcliff & McKoon, 2008; Ratcliff & Rouder, 1998; Ratcliff, 1978).

This model explains the dilemma between faster but more error-prone versus more accurate but slower decisions, namely the speed–accuracy tradeoff (SAT; e.g., Bogacz, Brown, Moehlis, Holmes, & Cohen, 2006). Specifically, when the decision boundary is set narrow, the decision process takes less time to hit a threshold but with a heightened risk of hitting the incorrect threshold due to noise. On the other hand, although the risk of hitting the incorrect threshold can be reduced by setting a wider decision boundary, this strategy leads to longer delays to first threshold crossing and thus longer RTs. Therefore, reward rate maximization requires finding a balance between how fast and accurately an agent aims to respond, which

[1]Koç University, Istanbul, Turkey, [2]University of Pennsylvania, [3]Maastricht University, [4]University of Otago, Dunedin, New Zealand

© 2017 Massachusetts Institute of Technology

Journal of Cognitive Neuroscience 29:8, pp. 1433–1444
doi:10.1162/jocn_a_01134

can be accomplished through adaptive decision threshold modulation (Bogacz, Wagenmakers, Forstmann, & Nieuwenhuis, 2010; Bogacz et al., 2006).

Although the SAT and related latent decision processes have been successfully studied within a computational framework, the neural mechanisms of this adaptive function have not been investigated sufficiently (Bogacz, Wagenmakers, et al., 2010). Providing theoretical guidance to empirical undertakings, neurocomputational studies implicate a role of the cortico-basal ganglia circuitry for decision threshold modulation (Bogacz & Gurney, 2007; Frank, 2006; Lo & Wang, 2006; Gurney, Prescott, Wickens, & Redgrave, 2004). BG is the main constituent of this circuitry, which comprise a distributed set of brain structures and primarily involve the control of voluntary actions. At resting state, the globus pallidus interna (GPi), one of the output nuclei of the BG, inhibits the thalamus and consequently the cortical areas, so that no premature responses are executed (e.g., DeLong & Wichmann, 2007). Striatum, one of the input nuclei of BG, is activated when it receives consistent information supporting a particular action from cortical regions. This activation exerts a selective suppression on the GPi. Through suppression of the inhibitory effect of GPi on thalamus, associated cortical regions are released from inhibition leading to action execution. On the basis of this functional architecture, BG are proposed to implement an action selection mechanism that disinhibits the desirable actions while inhibiting others (Forstmann et al., 2008).

Within this framework, when speed instructions are given, cortical neurons (e.g., within pre-SMA) send excitatory signals to the striatum. Increased striatal activity reduces the inhibitory effect of the BG, which in turn allows the execution of faster but often premature responses (Bogacz, Wagenmakers, et al., 2010; Forstmann et al., 2008). Consistent with this framework, a number of previous studies have shown that the pre-SMA and striatum show increased activity under speed-stressed conditions in perceptual decision-making tasks (Mansfield, Karayanidis, Jamadar, Heathcote, & Forstmann, 2011; van Maanen et al., 2011; Wenzlaff, Bauer, Maess, & Heekeren, 2011; Ding & Gold, 2010; Forstmann et al., 2008, 2010; Ivanoff, Branning, & Marois, 2008; van Veen, Krug, & Carter, 2008). Additionally, the increased activity in pre-SMA and striatum was correlated with the modulation of decision threshold; as the activity increased, the decision threshold decreased (Mansfield et al., 2011; Forstmann et al., 2008). Lending further support for the role of corticostriatal circuitry in modulating SAT, Forstmann et al. (2010) showed in a structural MRI study that participants who demonstrated larger adjustments of the decision boundary (i.e., who changed their decision thresholds considering the task demands in a 2AFC task) had stronger connectivity between pre-SMA and striatum.

Although imaging studies point out that pre-SMA and striatum play a role in SAT, as these studies provide only correlational information, they do not enable us to draw causal links. Recently, some brain stimulation studies were conducted to causally investigate the neural mechanisms underlying SAT. However, in addition to being quite limited in number, these studies did not converge on consistent conclusions; one of them (de Hollander et al., 2016) reported a null effect for the role of pre-SMA in threshold modulation in SAT, which could have possibly been attributed to the stimulation technique used (see Discussion), whereas the other (Georgiev et al., 2016) found that decreased activity in pre-SMA led to lower threshold setting under accuracy-stressed condition, which runs counter to these earlier findings and approaches.

Because the causal role of pre-SMA in SAT modulation is understudied and debatable, we investigated this issue by combining noninvasive brain stimulation with model-based unified analysis of decision outputs. We experimentally manipulated the activity of the right pre-SMA by applying continuous theta burst stimulation (cTBS) before having participants perform a 2AFC task. The RT and accuracy were modeled within the framework of the DDM to elucidate the effects of the inhibition of pre-SMA on decision threshold setting. We specifically hypothesized that the cTBS-induced inhibition of right pre-SMA would lead to more cautious decisions by leading to wider decision boundaries and an accuracy bias with respect to optimality.

METHODS

Participants

Twenty-four right-handed healthy volunteers (13 women), aged 19–24 years ($M = 20.9$, $SD = 1.7$) participated in the study after giving written consent. Participants were recruited through an announcement on a publicly available newsletter and received monetary reward based on their task performance. A preexperimental health form was used to screen for contraindications of TMS. Any participant who did not meet the eligibility criteria (Rossi, Hallett, Rossini, Pascual-Leone, & Safety of TMS Consensus Group, 2009) was excluded from the experiment. The study was approved by the institutional review board at Koç University.

Design

All participants were tested in three fixed-duration (approximately 45 min) sessions, all of which were held in different days. In the first (familiarization) session, participants were tested without brain stimulation. In the second and the third sessions, participants completed the same task after either pre-SMA or vertex inhibition in a counterbalanced order. Vertex was chosen as the control site as it plays no active role in the neural processes being investigated.

Stimuli and Apparatus

The stimulus was a circular field of randomly moving white dots (3 × 3 pixels) that appeared in a 3-in. diameter

kinematogram in the center of the computer screen with a black background (see Gold & Shadlen, 2001; Shadlen & Newsome, 2001). On each trial, a particular portion of the dots moved either toward left or right (with equal probability) with a fixed speed, whereas the rest of the dots were randomly repositioned over time. The stimulus was generated in MATLAB (The MathWorks Natick, MA) and presented on a 21.5-in. MAC monitor via Psychophysics Toolbox extension (Brainard, 1997; Pelli, 1997). Participants were seated approximately 60 cm from the monitor and responses were collected via computer keyboard.

For brain stimulation, a Magstim Super Rapid^2 magnetic stimulator (70-mm figure-of-eight coil) was used. Intensity and locations for repetitive TMS (rTMS) application were determined for each participant before behavioral testing. The international 10–20 system for EEG electrode placement was utilized for the localization of the target brain sites. EEG caps with 74 positions designed according to the 10–20 system (The g.GAMMAcap, G.Tec Medical Engineering GMBH, Austria) were used to define the 10–20 positions.

Procedure

Free-response Dot Motion Discrimination

Each session comprised nine 4-min test blocks of free-response (FR) dot motion discrimination task with 8% motion coherence and two 2-min signal detection blocks to determine non-decision times (to be used in optimality analysis presented below). Participants could take a break of up to 4 min at the end of the FR part. For FR trials, participants were instructed to respond as quickly and accurately as possible by pressing the "M" key (for rightward motion) with their right index finger or "Z" key (for leftward motion) with their left index finger. Each response terminated the stimulus presentation. The response-to-stimulus interval was sampled from a truncated exponential distribution (mean = 2 sec). Correct responses were signaled by a short audio tone whereas no feedback was provided for incorrect responses. Participants collected one point, corresponding to 0.04 TRY, for each correct response. Responses emitted during response-to-stimulus interval or within the first 100 msec of stimulus presentation were considered premature responses and were followed by a 4-sec timeout period, which started after a buzzing sound. The cumulative scores were displayed on the screen after every 10 trials. For signal detection trials, participants were instructed to press the "M" key and "Z" key in different blocks as soon as they saw the stimulus on the screen without considering the coherent motion.

rTMS Protocol

An offline cTBS was applied over the target brain site (either right pre-SMA or vertex in a counterbalanced order) at the beginning of the second and the third sessions, before participants performed the task. This protocol consists of three pulses of stimulation given at 50 Hz, repeated every 200 msec for 40 sec in an uninterrupted fashion, resulting in 600 pulses (Huang, Edwards, Rounis, Bhatia, & Rothwell, 2005). The cTBS protocol was preferred over the traditional rTMS protocol, as it enables having a longer-lasting inhibitory effect (as long as 60 min) with a considerably shorter stimulation duration (Huang et al., 2005), which makes the stimulation more comfortable for the participants. Vertex is defined as the intersection of the midpoint between nasion and inion with the midpoint between the left and right intertragal notches (Cz site in the international 10–20 EEG system). For the localization of the right pre-SMA, the center of the magnetic coil was placed over the Fz site and moved 1 cm lateral to the right from the mid-sagittal line (Cavazzana, Penolazzi, Begliomini, & Bisiacchi, 2015; Hsu et al., 2011). To set the specific intensity of stimulation for each participant, single-pulse TMS was applied to motor cortex at increasing intensities, and the active motor threshold of each individual was determined according to the criterion that a given intensity evokes a muscle twitch in the contralateral hand (Huang et al., 2005). The stimulation intensity was determined individually as 80% of the active motor threshold.

Data Analysis

The units of analysis were the accuracy and RT data obtained from the test sessions. Premature responses (RTs < 100 msec, 0.07% of all trials) were excluded from the analysis.

Behavioral Analysis

We used paired samples *t* tests to compare the RT and accuracy between different stimulation conditions. To treat RT and accuracy in a unified fashion, we also calculated three different measures commonly used for integrating these two variables as an alternative to their isolated analyses (for a review, see Vandierendonck, 2017). To calculate inverse efficiency score (IES), we divided the average correct RT by the proportion of correct responses. The rate of correct score (RCS) was calculated by dividing the number of correct responses by the sum of all RTs in a given condition. Lastly, linear integrated speed–accuracy score (LISAS) was calculated by adding up the mean RT in a given condition and the proportion of errors weighted by the ratio of RT and error proportion standard deviations. To complement the conventional *t* tests, Bayesian *t* tests were conducted where appropriate using JASP software (JASP Team, 2016). The default Cauchy distribution with width 0.707 was used as the prior distribution (for details, see Rouder, Speckman, Sun, Morey, & Iverson, 2009).

Hierarchical Bayesian DDM

Data were fit by the DDM using the Hierarchical Bayesian estimation of DDM parameters (HDDM) in Python (Wiecki, Sofer, & Frank, 2013). To estimate the within-subject effects, we used HDDMRegressor class. In the Bayesian estimation procedure used in HDDM, quantification of parameter estimates is performed in the form of the posterior distributions, which are approximated by Markov chain Monte Carlo sampling methods. Prior distributions for each parameter were informed by a collection of 23 studies reporting best-fitting DDM parameters recovered on a range of decision-making tasks (Wiecki et al., 2013; Matzke & Wagenmakers, 2009). To obtain smooth parameter estimates, 10,000 samples were drawn from the posterior distribution, and the first 1000 were discarded as burn-in.

We fit four different models with varying degrees of complexity. In Model 1, we allowed the decision threshold to vary between conditions to test the most theoretically constrained model. For completeness, we fit three additional models with increasing number of parameters allowed to vary between conditions (see Table 1 for model specifications).

To evaluate the performance of our four alternative models with varying complexity, we used each model's deviance information criterion (DIC). The DIC values for Models 1, 2, 3, and 4 were 63,484, 63,478, 63,477, and 58,721, respectively, favoring Model 4 with intertrial variability parameters as the model that fit the data best (Burnham & Anderson, 2003). Because the differences between the DIC values of the reduced models (Models 1–3) were not higher than 10, none of these models performs significantly better than the other. Note that other models with different intertrial variability parameters were also fit to the data, but their DIC scores were much higher than the full model (62,255 for the model with only drift rate variability, 60,020 for the model with only st, 63,478 for the model with only starting point variability). These models led to identical conclusions regarding decision threshold and drift rate effects based on the same results with the four models presented in the study. For brevity, these other models are not presented in the article.

Reward Rate Maximization

The expected reward rate (RR) in FR 2AFC tasks is calculated as in Equation 1 (Gold & Shadlen, 2002),

$$\mathrm{RR} = \frac{1 - \mathrm{ER}}{\mathrm{DT} + T_0 + \mathrm{RSI}} \tag{1}$$

where ER denotes error rate, DT represents decision time, T_0 is time required for all non-decision-related processes such as sensory processing and motor execution, and RSI is the response-to-stimulus interval. DT is calculated by subtracting the non-decision-related times from the RTs. Reward maximizing mean-normalized decision times are calculated as follows (Bogacz et al., 2006),

$$\frac{\mathrm{DT}}{D_{\mathrm{tot}}} = \frac{1}{\frac{1}{\mathrm{ER}\log\left(\frac{1-\mathrm{ER}}{\mathrm{ER}}\right)} + \frac{1}{1-2\mathrm{ER}}} \tag{2}$$

where $D_{\mathrm{tot}} = T_0 + \mathrm{RSI}$. This optimal performance curve was derived from the reduced form of the DDM (Bogacz et al., 2006) with no trial-to-trial variability parameters (e.g., Models 1–3 in our study).

In previous studies utilizing 2AFC tasks, the majority of the participants were shown to set their decision thresholds higher than the reward maximizing thresholds early in training (e.g., Balcı et al., 2011; Bogacz, Hu, Holmes, & Cohen, 2010; Maddox & Bohil, 1998). This accuracy bias was formulated by Bogacz et al. (2006) as below,

$$\mathrm{RR}(q) = \frac{(1 - \mathrm{ER}) - q\mathrm{ER}}{\mathrm{DT} + D_{\mathrm{tot}}} \tag{3}$$

where q is the additional parameter representing a penalty for errors in the form of the weight assigned to accuracy relative to reward rate. Considering different error rates, the related normalized optimal decision times (given that

Table 1. Model Specifications

	Model Descriptions			
Parameter	*Model 1*	*Model 2*	*Model 3*	*Model 4*
Threshold	Varied	Varied	Varied	Varied
Drift rate	Fixed	Varied	Varied	Varied
Non-decision time	Fixed	Fixed	Varied	Varied
Starting point	No bias	No bias	No bias	No bias
Drift rate variability (sv)	0	0	0	Fixed
Non-decision time variability (st)	0	0	0	Fixed
Starting point variability (sz)	0	0	0	Fixed

q stands for an actual penalty) are calculated using the formula below (Bogacz et al., 2006):

$$\frac{DT}{D_{tot}} = (1+q)\frac{1}{\frac{\frac{1}{ER}-\frac{q}{1-ER}}{\log\left(\frac{1-ER}{ER}\right)}+\frac{1-q}{1-2ER}} \quad (4)$$

When q is equal to zero, Equation 4 prescribes the optimal performance curve for a task with no penalty for errors (as in the case of the current task; Equation 4 reduces to Equation 2), and thus, the best fit q value indicates the degree of accuracy bias assuming that the participant optimizes this alternative function with subjective penalty for errors. Thus, to find this subjective cost for each participant, we obtained their average error rates and normalized decision times in each condition. We then used these inputs to estimate individual q parameters based on the simplex search algorithm that starts from a predefined starting point and aiming to find a local minimum for the parameter of interest (Lagarias, Reeds, Wright, & Wright, 1998).

Post-error Slowing

Post-error slowing (PES) scores were quantified based on two different methods. In the first method, PES was determined as the difference in the RTs between the post-error and post-correct trials. We also quantified PES as described by Dutilh et al. (2012) to avoid the potential effect of global fluctuations in performance.

RESULTS

RT and Accuracy Comparisons

We first examined whether there was a difference in the RT and the accuracy between the pre-SMA and vertex inhibition conditions. Although the change in both RTs and accuracy levels were in the predicted directions (Table 2), these were not statistically significant ($t(23) = 1.20, p = .24$ and $t(23) = 1.15, p = .26$, respectively). To further investigate the strength of evidence in favor of these null findings, we conducted Bayesian t tests (Rouder et al., 2009). As the results revealed, the odds were 2.47:1 (weak evidence; Raftery, 1995) and 2.57:1 (weak evidence) in favor of the null hypothesis that there was no significant difference between conditions for the RTs and the accuracy levels, respectively.

To further examine the behavioral differences between the two stimulation conditions, three different methods of combining decision outputs were utilized as alternative to analyzing them in isolation. The comparison of the obtained scores between pre-SMA and vertex inhibition conditions revealed no significant difference for any of these measures (IES: $t(23) = 1.29, p = .21$; RCS: $t(23) = .97, p = .34$; LISAS: $t(23) = .46, p = .65$). Bayesian paired samples t tests also provided evidence in favor of the null hypothesis that there is no difference between conditions in terms of IES ($BF_{01} = 2.24$), RCS ($BF_{01} = 3.06$), and LISAS ($BF_{01} = 4.24$).

Table 2. Means and *SD*s of RTs and Accuracy Levels in Pre-SMA and Vertex Inhibition Conditions

	RT (msec)		*Accuracy Levels*	
Condition	*M*	*SD*	*M*	*SD*
Pre-SMA inhibition	998.85	248.89	0.78	0.13
Vertex inhibition	957.96	267.42	0.76	0.11

Effects on the Latent Decision Process

To examine whether the latent decision processes associated with the control of the SAT has been affected by our experimental manipulation, we fit four different DDMs to the data. The first three models were reduced DDMs, whereas the last model was an extended DDM with intertrial variability parameters (see Table 1 for model specifications).

Using the hierarchical DDM, the posterior distribution for the difference between the pre-SMA and vertex inhibition conditions was generated. The degree of overlap between the posterior distribution of this difference and the value of 0 was used as the comparison metric for the two stimulation conditions, where at least 95% of nonoverlap indicates a significant difference between the conditions. To assess whether the Markov Chain Monte Carlo chains successfully converged from the starting point to the posterior distribution, we calculated the (Gelman-Rubin) statistic from five separate runs (each containing 10,000 samples) for each model. All of the values for all model parameters were lower than 1.1, indicating successful convergence. To assess whether our models reproduced the key features in our data, we conducted posterior predictive checks as described in the HDDM tutorial (ski.clps.brown.edu/hddm_docs/tutorial_post_pred.html). For each model, we simulated data sets from the posterior samples of the corresponding fitted model and obtained summary statistics from these simulations to compare them to the summary statistics describing our actual data. For all models, the observed RTs were within 95% credible interval of the predicted data (see Figure 1 for observed and predicted RTs for different quantiles for both correct and error trials).

Model 1 is the most theoretically compliant model with our hypothesis since only the decision threshold parameter was allowed to vary between the conditions. The posterior distribution of the threshold parameter did not overlap with zero, indicating a significant effect of the condition on this parameter estimate (100% of posterior > 0). Because the within-subjects effect of the pre-SMA inhibition condition on decision threshold parameter with

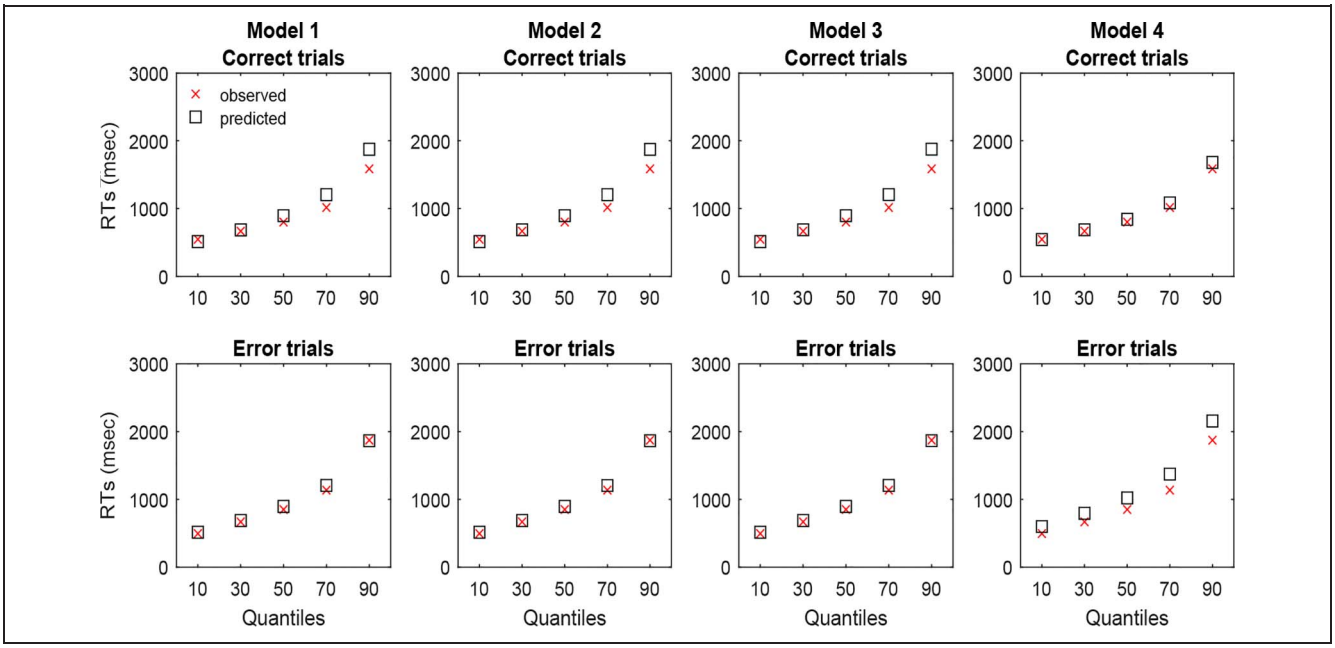


Figure 1. Observed (red crosses) and predicted (black open squares) RTs for 0.1, 0.3, 0.5, 0.7, and 0.9 quantiles separately for correct responses (top row) and incorrect RTs (bottom row) and four different models (columns).

respect to the control condition is higher than 0 (see Figure 2A), decision threshold estimated for the pre-SMA condition is significantly higher than threshold in the control condition. This result indicates that when the excitability of the right pre-SMA was reduced, participants set higher decision thresholds, exhibiting a more cautious decision strategy.

In Model 2, both decision threshold and drift rate parameters were allowed to vary between the conditions. The results of fits of this model (see Figure 2B) showed that threshold in the pre-SMA inhibition condition were higher than the threshold in the vertex inhibition condition (100% of the posterior > 0). Thus, consistent with the findings gathered based on the simpler model, inhibition of right pre-SMA has a significant effect on the response cautiousness levels of the participants, making them to set higher decision thresholds. The within-subject effect of the pre-SMA inhibition on the drift rate parameter was also significant (99% of posterior > 0). This result indicated that participants had higher evidence accumulation rates when their right pre-SMA was inhibited compared with the condition in which their vertex was inhibited.

In Model 3, along with the threshold and drift rate parameters, non-decision time parameter was also allowed to vary between the conditions. Consistent with the results

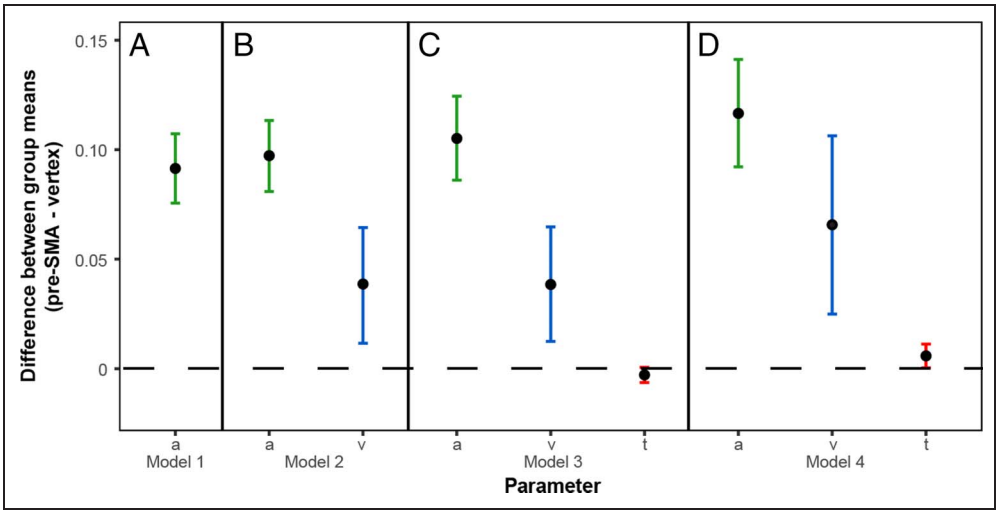


Figure 2. Mean difference in the threshold (a, green), drift rate (v, blue), and non-decision time (t, red) parameters between pre-SMA and vertex inhibition conditions with 95% credible intervals separately for Models 1 (A), 2 (B), 3 (C), and 4 (D). Note that positive differences indicate higher values in the pre-SMA inhibition compared with vertex inhibition condition. Also note that only the parameters allowed to vary between these two conditions are displayed for each model.

of the first and the second model fits, threshold (100% of posterior > 0) and drift rate (99% of posterior > 0) were significantly higher in the pre-SMA inhibition condition compared with the vertex inhibition condition, whereas the effect of right pre-SMA inhibition on the non-decision time parameter was not significant (94% of posterior < 0; see Figure 2C).

In Model 4, the intertrial variability parameters for drift rate, non-decision time, and starting point were also included in the model. Consistent with the other models, threshold (99% of posterior > 0) and drift rate estimates (100% of posterior > 0) were significantly higher in the pre-SMA inhibition condition compared with the vertex inhibition condition (see Figure 2D). Additionally, non-decision time estimates for the pre-SMA inhibition condition were also higher than those for the vertex inhibition condition (98% of posterior higher than 0).

Figure 2A–D show the means and 95% credible intervals of the difference between the pre-SMA and vertex inhibition conditions for the parameters relevant for their comparison in Models 1, 2, 3, and 4, respectively. This difference reflects the within-subject effect of pre-SMA inhibition condition with respect to the control condition. In all models (Figure 2A–D), mean difference between the conditions and 95% credible intervals for threshold and drift rate parameters were above 0, indicating a significant increase in these parameters in pre-SMA inhibition condition. In Model 3 (Figure 2C), 95% credible intervals for the difference in non-decision time parameter between the conditions included 0, indicating that this parameter was not affected by the stimulation, whereas in Model 4 (Figure 2D) mean difference in this parameter between the conditions and 95% credible intervals was above 0.

Speed–Accuracy Tradeoff and Reward Rate Maximization

Because in free response fixed session duration tasks, reward rate maximization requires participants to optimize the balance between the speed and accuracy of their decisions, we also assessed whether inhibition of pre-SMA had an effect on this reward rate maximization process. As the previous 2AFC studies where the majority of the participants were shown to set higher decision thresholds than the reward maximizing thresholds predicted by the DDM (e.g., Balcı et al., 2011; Bogacz, Hu, et al., 2010; Maddox & Bohil, 1998), we calculated the subjective cost participants attributed to errors (i.e., accuracy bias) in different conditions. To assess whether the performance of the participants differed based on their accuracy bias in different conditions, we compared the values of the parameter q corresponding to the error rates of each participant in both experimental and control sessions (see Equation 4). The q values calculated for both the experimental (M = .47, SD = .40) and the control conditions (M = .35, SD = .26) were significantly greater than 0, $t(23) = 4.13, p < .001$, and $t(23) = 4.26$, $p < .001$, respectively. This result indicates that, regardless of the condition, participants had an accuracy bias. Comparison of the accuracy bias in different conditions showed that, in line with our predictions, the weight assigned to accuracy relative to reward rate was significantly higher in the pre-SMA compared with the vertex inhibition condition, $t(23) = 2.35, p < .05$. Thus, consistent with our findings regarding decision threshold differences, participants displayed a more cautious performance under the pre-SMA inhibition condition.

Post-error Slowing

To explore whether the increased cautiousness levels in the pre-SMA inhibition condition was related to any change in the tendency to slow down after erroneous responses, we evaluated the PES scores. According to the results of the standard method, PES score was significantly higher than 0 in both experimental (M = 114.96, SD = 180.55; $t(23) = 3.05$, $p < .01$) and control conditions (M = 98.62, SD = 223.67; $t(23) = 2.11$, $p < .05$), pointing at a tendency to slow down after erroneous responses in both conditions. However, the comparison of these two conditions indicated no significant effect of stimulation site on the level of PES, $t(23) = .78$, $p = .44$. An estimated Bayes factor revealed that the odds were 3.53:1 in favor of the null hypothesis, providing substantial evidence for that there was no difference between groups in terms of PES. Similar results were obtained from the set of analyses conducted based on the method described in Dutilh et al. (2012), indicating that the difference between the post-error RTs and the pre-error RTs were significantly higher than the value of 0 both in the pre-SMA (M = 124.23, SD = 214.62; $t(23) = 2.54, p = .02$) and the vertex inhibition conditions (M = 83.26, SD = 134.57; $t(23) = 3.16, p$ = .004). There was no significant difference between the two conditions, $t(23) = 1.30$, $p = .21$. We also conducted a Bayesian t test, which showed that the odds were 1.64:1 in favor of the null hypothesis (weak evidence), indicating that there was no difference between pre-SMA and vertex inhibition conditions in terms of PES.

We also investigated whether the accuracy of decisions increased after errors by comparing post-correct and post-error accuracy rates. The results indicated that the difference between the post-correct and post-error accuracy rates were not significantly higher than the value of 0 in either pre-SMA (M = 0.01, SD = 0.04; $t(23) = 0.93$, $p = .36$) or the control (M = 0.01, SD = 0.05; $t(23)$ = 1.35, $p = .19$) condition. Also, the stimulation condition did not have a differential effect on the difference in the accuracy rates between post-error and post-correct trials, $t(23) = -0.36, p = .72$. Bayesian t tests revealed that the odds were 3.15:1 and 2.09:1 in favor of the null hypothesis that the difference between post-correct and post-error accuracy rates was not different from the value of 0 in pre-SMA and control conditions, respectively. Additionally,

the odds were 4.39:1 in favor of the null hypothesis that the stimulation site did not have a differential effect on the difference between the post-correct and post-error accuracy rates.

To examine whether there was any difference in threshold setting between post-correct and post-error trials, and if so, whether this difference varied between the pre-SMA and vertex stimulation conditions, we fit a DDM to the data allowing only the decision threshold parameter to vary. The results of this analysis revealed a higher threshold setting for the post-error trials compared with the post-correct trials regardless of the stimulation site (100% of posterior > 0). However, the stimulation site did not have a differential effect on the difference between post-error and post-correct trials in terms of threshold setting (64% of posterior > 0). We repeated this set of modeling analyses also using the alternative PES coding method (Dutilh et al., 2012). These results, too, revealed that threshold settings for the post-error trials were higher than the threshold settings for post-correct (pre-error) trials (100% of posterior > 0), whereas the experimental manipulation did not have a differential effect on the difference in threshold settings between the post-error and post-correct trials (85% of posterior > 0).

DISCUSSION

In the current study, we investigated the functional relationship between the activity within the right pre-SMA and the SAT in human perceptual decision-making by combining noninvasive brain stimulation with model-based approaches to choice behavior. Our results revealed that humans indeed exhibit a more cautious choice behavior by setting higher decision thresholds as a result of right pre-SMA inhibition compared with the inhibition of a control region, vertex.

Importantly, this difference was present despite the null effects revealed when accuracy levels and RTs were compared in isolation or even in a unified fashion (i.e., IES, RCS, LISAS), reflecting the importance of model-based approaches in cognitive neuroscience research (Ly et al., in press; Erhan & Balcı, 2017; Georgiev et al., 2016; Forstmann & Wagenmakers, 2015; Voss, Nagler, & Lerche, 2013; Voss, Rothermund, & Voss 2004). Along with the increase in the decision thresholds, we also found that the weight assigned to accuracy relative to reward rate (i.e., accuracy bias; Balcı et al., 2011; Bogacz et al., 2006; Maddox & Bohil, 1998) was significantly higher in the right pre-SMA as compared with the vertex inhibition condition. This result reflects the effect of right pre-SMA inhibition at the level of behavioral output when it is evaluated with respect to the optimality benchmark.

On the basis of the neurocomputational models of decision-making (Bogacz & Gurney, 2007; Frank, 2006), we propose that the inhibition of right pre-SMA exerts its downstream effect on the related cortico-BG pathways. To this end, there are four major neurocomputational theories that have been proposed in relation to the control of speed–accuracy tradeoff (for a review, see Bogacz, Hu, et al., 2010). The cortical theory asserts that the baseline activity of cortical integrators would be increased by additional excitatory input, which would be effectively equivalent to decreased threshold (e.g., Ivanoff et al., 2008; van Veen et al., 2008). To this end, van Veen et al. (2008) showed that dorsolateral pFC is the source of these excitatory signals that modulate SAT. The striatal theory, on the other hand, asserts that the excitability of striatum modulated by cortical inputs (specifically pre-SMA) effectively results in changing decision boundaries by modulating the inhibitory effect of BG over its efferents (Forstmann et al., 2008, 2010).

More specifically, according to the striatal theory, as a result of increased right pre-SMA activity, humans execute faster but often premature responses because an increase in the activity of cortical nonintegrator neurons excites striatum, which in turn decreases the inhibitory effect of the output nuclei of BG (GPi) over the cortical areas related to motor execution via thalamus (Bogacz, Wagenmakers, et al., 2010; Forstmann et al., 2008). In the same line of reasoning, inhibition of right pre-SMA decreases the activity of the striatum, which causes the activity of GPi to increase. This in turn leads to a decreased activity in thalamus, resulting in a more cautious decision strategy (i.e., effectively corresponding to increased threshold setting).

Subthalamic nucleus (STN) theory of SAT asserts that cortical areas can effectively modulate decision thresholds through their projections to STN (e.g., pre-SMA, inferior frontal gyrus, ACC). Under this view, lower activity of the cortical afferents of STN would decrease the thresholds (Aron, Behrens, Smith, Frank, & Poldrack, 2007; Aron & Poldrack, 2006). Finally, the synaptic theory of SAT asserts that SAT can be modulated by changing the weights of the corticostriatal synapses through processes such as long-term potentiation induced by reinforcement learning (Lo & Wang, 2006). According to this view, strengthening of these synapses would effectively decrease whereas weakening (e.g., through long-term depression) of these synapses would effectively increase the decision thresholds.

Under these approaches to the neural basis of SAT modulation, our findings are more in favor of the striatal theory of SAT compared with the cortical theory given the site of stimulation (assuming that pre-SMA does not contain integrator units) and its functional connectivity to striatum (Bogacz, Wagenmakers, et al., 2010). If, on the other hand, pre-SMA is contained within the integrator network, our findings cannot distinguish between the striatal and cortical theories of SAT. Our findings do not support the STN theory given the fact that, under this view, pre-SMA inhibition should have led to a decrease, not an increase, in decision thresholds. Although the directionality of the predictions is similar for striatal and synaptic theories (but also cortical theory), the synaptic

theory is a less plausible account of our findings given the speed with which such synaptic changes would take place in the brain (see also Furman & Wang, 2008).

The primary assertion of the striatal theory of SAT is further supported by the findings of a recent rTMS-fMRI study in the context of stop-signal response inhibition task (Watanabe et al., 2015). Watanabe et al. (2015) showed that rTMS of pre-SMA modulated the connectivity between pre-SMA and striatum as well as the connectivity between striatum and globus pallidus during a response inhibition task. Watanabe et al. (2015) also observed that the changes in response inhibition performance were associated with the degree of change in the resting state functional connectivity between these areas. These findings point to a causal relationship between pre-SMA and globus pallidus through striatum during response inhibition, which further suggests a functional interaction between these regions in the context of decision-making.

In addition to these functional connectivity findings, the study of Forstmann et al. (2010) provided evidence on the relationship between the individual differences in behavior and in the structural features of particular brain regions. Their results revealed that the participants who demonstrated larger adjustments of the decision boundary based on task demands had stronger connectivity between the pre-SMA and the striatum. Despite these consistent findings regarding the role of pre-SMA in decision threshold setting, a recent study by de Hollander et al. (2016) did not observe an effect of transcranial direct current stimulation (tDCS) of pre-SMA on this parameter. However, lack of effects in that study can be attributed to multiple factors such as the inefficacy of tDCS in decision-making domain (e.g., Horvath, Forte, & Carter, 2015; Tremblay et al., 2014) and/or the localization of the stimulation site.

Our results regarding decision thresholds differ from the findings of Georgiev et al. (2016). This study had two task conditions; in one condition (SAT condition) speed and accuracy was differentially emphasized between trials in a pseudorandom fashion (50% coherence was used), whereas in the other condition (Coherence condition) participants were asked to be as accurate and fast as possible (identical instructions to the current study) and a range of different coherence levels was used (5–50%). Georgiev et al. (2016) found a reduction in decision thresholds with inhibition of right pre-SMA when accuracy was emphasized whereas no difference was observed when speed was emphasized or in the Coherence condition. In contrast to Georgiev et al. (2016), we found an increase in decision thresholds as a result of inhibition of right pre-SMA. There are a number of factors that might have contributed to the results revealed in our study as compared with the findings reported by Georgiev et al. (2016).

One of the main differences is that our task was more difficult (8% coherence) compared with the SAT condition in Georgiev et al.'s (2016) study (50%). This difference in signal quality was also clearly reflected in the observed accuracy (77% vs. 95%) and RT levels (~980 msec vs. ~400 msec). The same also applies to the speed condition, where no effect on decision threshold was found. Different difficulty levels are known to be differentially sensitive to factors of interest. For instance, Banca et al. (2015) found medium level of coherence to be most sensitive to differences between obsessive compulsive disorder and control groups in terms of decision thresholds and drift rates, with low level of coherence being specifically sensitive to differences in decision thresholds and high level of coherence being specifically sensitive to difference in drift rates. However, please note that no effect of pre-SMA inhibition was observed in the coherence condition in which a range of difficulty levels was used. Furthermore, Georgiev et al. (2016) fit their data using fast-dm (Voss & Voss, 2007), which has lower statistical power compared with HDDM (Wiecki et al., 2013). Overall, results of the current study are more in line with the results of neuroimaging studies conducted in this field (Bogacz, Wagenmakers, et al., 2010).

In addition to the increase in threshold setting, our findings also indicated an increase in drift rate, which represents the amount of evidence processed in unit time, under pre-SMA inhibition condition. Although there has been a selective emphasis on the decision threshold modulation to explain the SAT (Bogacz, Wagenmakers, et al., 2010; Forstmann et al., 2008, 2010; Bogacz et al., 2006), some studies have indicated that SAT is better captured with changes in both threshold setting and drift rate parameters (Rae, Heathcote, Donkin, Averell, & Brown, 2014; Heathcote & Love, 2012; Vandekerckhove & Tuerlinckx, 2007). For instance, in a study by Rae et al. (2014), it has been shown in three different tasks that, under accuracy instructions, which, in terms of expected effects, correspond to the pre-SMA inhibition condition in our study, there was an increase not only in the amount of evidence required to commit to a decision, but also in the quality of the evidence accumulated during this process. However, to our knowledge, there is no study indicating a difference in drift rates along with the activity change in pre-SMA. Besides, drift rates have been typically associated with dorsolateral pFC function (e.g., Bogacz, Wagenmakers, et al., 2010; Domenech & Dreher, 2010; Heekeren, Marrett, Bandettini, & Ungerleider, 2004), and the recent cTBS and tDCS studies did not report any effect of pre-SMA inhibition on drift rates in either speed or in accuracy conditions or in conditions where differential speed–accuracy instructions were not provided (de Hollander et al., 2016; Georgiev et al., 2016). Thus, this finding might be investigated in future imaging studies. In addition to consistent effect of pre-SMA inhibition on decision threshold and drift rates, we found that inhibition of pre-SMA also led to longer non-decision times in Model 4 but not in Model 3. Because this was not a consistent finding as in the case of decision thresholds and drift rates and we did not expect this finding based on any theoretical approach, future studies are needed to provide further

insight regarding the credibility of the effect of pre-SMA inhibition on non-decision times.

In our further exploratory analyses, we also investigated the possible relationship between PES and SAT, because the underlying neural mechanisms of PES also include similar cortical and subcortical structures of the right hemisphere (Danielmeier & Ullsperger, 2011). In the review by Danielmeier and Ullsperger (2011), a network consisting of pre-SMA, lateral inferior frontal areas and the STN are suggested to be crucial for PES. In the hyperdirect pathway, STN receives direct cortical input and projects directly to the GPi to act as a global brake on the striatal output (Cavanagh, Sanguinetti, Allen, Sherman, & Frank, 2014). Thus, increased STN activity following erroneous responses enables acting more cautious in the next trial by increasing the response threshold (Cavanagh et al., 2014). In our study, participants in both conditions responded slower by setting higher thresholds after they made an error; however, this increase in the RTs or thresholds did not differ between pre-SMA and vertex inhibition conditions. Thus, participants did not exhibit differential post-error behavior in the pre-SMA inhibition condition compared with the control condition. In light of the findings of Watanabe et al. (2015), which indicated no effect of pre-SMA rTMS on the activity of STN, one can speculate that pre-SMA inhibition in our study did not modulate the activity of STN and therefore did not specifically lead to the modulation of PES. Thus, overall condition-based threshold setting (for which the signal-to-noise ratio and average reward rate are relevant) and error-based threshold modulation (for which the instantaneous consequences/outcomes of previous decisions are relevant) might rely on partially dissociable networks. Future studies may also investigate these relationships using functional imaging.

Acknowledgments

This study was supported by TÜBA (Turkish Academy of Sciences)-GEBİP 2015 award to F. B.

Reprint requests should be sent to Fuat Balcı, Koç University, Rumelifeneri Yolu, Sarıyer, 34450, Istanbul, Turkey, or via e-mail: fbalci@ku.edu.tr, web: mysite.ku.edu.tr/fbalci/.

REFERENCES

Aron, A. R., Behrens, T. E., Smith, S., Frank, M. J., & Poldrack, R. A. (2007). Triangulating a cognitive control network using diffusion-weighted magnetic resonance imaging (MRI) and functional MRI. *Journal of Neuroscience, 27,* 3743–3752.

Aron, A. R., & Poldrack, R. A. (2006). Cortical and subcortical contributions to stop signal response inhibition: Role of the subthalamic nucleus. *Journal of Neuroscience, 26,* 2424–2433.

Balcı, F., Freestone, D., Simen, P., deSouza, L., Cohen, J. D., & Holmes, P. (2011). Optimal temporal risk assessment. *Frontiers in Integrative Neuroscience, 5,* 1–15.

Banca, P., Vestergaard, M. D., Rankov, V., Baek, K., Mitchell, S., Lapa, T., et al. (2015). Evidence accumulation in obsessive-compulsive disorder: The role of uncertainty and monetary reward on perceptual decision-making thresholds. *Neuropsychopharmacology, 40,* 1192–1202.

Bogacz, R. (2007). Optimal decision-making theories: Linking neurobiology with behaviour. *Trends in Cognitive Sciences, 11,* 118–125.

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,* 700–765.

Bogacz, R., & Gurney, K. (2007). The basal ganglia and cortex implement optimal decision making between alternative actions. *Neural Computation, 19,* 442–477.

Bogacz, R., Hu, P. T., Holmes, P. J., & Cohen, J. D. (2010). Do humans produce the speed–accuracy trade-off that maximizes reward rate? *Quarterly Journal of Experimental Psychology, 63,* 863–891.

Bogacz, R., Wagenmakers, E. J., Forstmann, B. U., & Nieuwenhuis, S. (2010). The neural basis of the speed–accuracy tradeoff. *Trends in Neuroscience, 33,* 10–16.

Brainard, D. H. (1997). The psychophysics toolbox. *Spatial Vision, 10,* 433–436.

Burnham, K. P., & Anderson, D. R. (2003). *Model selection and multimodel inference: A practical information-theoretic approach*. New York: Springer.

Cavanagh, J. F., Sanguinetti, J. L., Allen, J. J., Sherman, S. J., & Frank, M. J. (2014). The subthalamic nucleus contributes to post-error slowing. *Journal of Cognitive Neuroscience, 26,* 2637–2644.

Cavazzana, A., Penolazzi, B., Begliomini, C., & Bisiacchi, P. S. (2015). Neural underpinnings of the "agent brain": New evidence from transcranial direct current stimulation. *European Journal of Neuroscience, 42,* 1889–1894.

Danielmeier, C., & Ullsperger, M. (2011). Post-error adjustments. *Frontiers in Psychology, 2,* 233.

de Hollander, G., Labruna, L., Sellaro, R., Trutti, A., Colzato, L., Ratcliff, R., et al. (2016). Transcranial direct current stimulation does not influence the speed–accuracy tradeoff in perceptual decision making: Evidence from three independent replication studies. *Journal of Cognitive Neuroscience, 28,* 1283–1294.

DeLong, M. R., & Wichmann, T. (2007). Circuits and circuit disorders of the basal ganglia. *Archives of Neurology, 64,* 20–24.

Ding, L., & Gold, J. I. (2010). Caudate encodes multiple computations for perceptual decisions. *Journal of Neuroscience, 30,* 15747–15759.

Domenech, P., & Dreher, J. C. (2010). Decision threshold modulation in the human brain. *Journal of Neuroscience, 30,* 14305–14317.

Dutilh, G., van Ravenzwaaij, D., Nieuwenhuis, S., van der Maas, H. L., Forstmann, B. U., & Wagenmakers, E. J. (2012). How to measure post-error slowing: A confound and a simple solution. *Journal of Mathematical Psychology, 56,* 208–216.

Erhan, C., & Balcı, F. (2017). Obsessive compulsive features predict cautious decision strategies. *Quarterly Journal of Experimental Psychology, 70,* 179–190.

Forstmann, B. U., Anwander, A., Schäfer, A., Neumann, J., Brown, S., Wagenmakers, E. J., et al. (2010). Cortico-striatal connections predict control over speed and accuracy in perceptual decision making. *Proceedings of the National Academy of Sciences, U.S.A., 107,* 15916–15920.

Forstmann, B. U., Dutilh, G., Brown, S., Neumann, J., Von Cramon, D. Y., Ridderinkhof, K. R., et al. (2008). Striatum and pre-SMA facilitate decision-making under time pressure. *Proceedings of the National Academy of Sciences, U.S.A., 105,* 17538–17542.

Forstmann, B. U., & Wagenmakers, E. J. (2015). Model-based cognitive neuroscience: A conceptual introduction. In B. U. Forstman & E. J. Wagenmakers (Eds.), *An introduction to model-based cognitive neuroscience* (pp. 139–156). New York: Springer.

Frank, M. J. (2006). Hold your horses: A dynamic computational role for the subthalamic nucleus in decision making. *Neural Networks, 19,* 1120–1136.

Furman, M., & Wang, X. J. (2008). Similarity effect and optimal control of multiple-choice decision making. *Neuron, 60,* 1153–1168.

Georgiev, D., Rocchi, L., Tocco, P., Speekenbrink, M., Rothwell, J. C., & Jahanshahi, M. (2016). Continuous theta burst stimulation over the dorsolateral prefrontal cortex and the pre-SMA alter drift rate and response thresholds respectively during perceptual decision-making. *Brain Stimulation, 4,* 601–608.

Gold, J. I., & Shadlen, M. N. (2001). Neural computations that underlie decisions about sensory stimuli. *Trends in Cognitive Sciences, 5,* 10–16.

Gold, J. I., & Shadlen, M. N. (2002). Banburismus and the brain: Decoding the relationship between sensory stimuli, decisions, and reward. *Neuron, 36,* 299–308.

Gold, J. I., & Shadlen, M. N. (2007). The neural basis of decision making. *Annual Review of Neuroscience, 30,* 535–574.

Gurney, K., Prescott, T. J., Wickens, J. R., & Redgrave, P. (2004). Computational models of the basal ganglia: From robots to membranes. *Trends in Neurosciences, 27,* 453–459.

Heathcote, A., & Love, J. (2012). Linear deterministic accumulator models of simple choice. *Frontiers in Psychology, 3,* 292.

Heekeren, H. R., Marrett, S., Bandettini, P. A., & Ungerleider, L. G. (2004). A general mechanism for perceptual decision-making in the human brain. *Nature, 431,* 859–862.

Horvath, J. C., Forte, J. D., & Carter, O. (2015). Evidence that transcranial direct current stimulation (tDCS) generates little-to-no reliable neurophysiologic effect beyond MEP amplitude modulation in healthy human subjects: A systematic review. *Neuropsychologia, 66,* 213–236.

Hsu, T. U., Tseng, L. Y., Yu, J. X., Kuo, W. J., Hung, D. L., Tzeng, O. J. L., et al. (2011). Modulating inhibitory control with direct current stimulation of the superior medial frontal cortex. *Neuroimage, 56,* 2249–2257.

Huang, Y. Z., Edwards, M. J., Rounis, E., Bhatia, K. P., & Rothwell, J. C. (2005). Theta burst stimulation of the human motor cortex. *Neuron, 45,* 201–206.

Ivanoff, J., Branning, P., & Marois, R. (2008). fMRI evidence for a dual process account of the speed–accuracy tradeoff in decision-making. *PLoS One, 3,* e2635.

JASP Team. (2016). JASP (Version 0.8. 0.0). [Computer software]. Retrieved from https://jasp.stats.org/.

Lagarias, J. C., Reeds, J. A., Wright, M. H., & Wright, P. E. (1998). Convergence properties of the Nelder–Mead simplex method in low dimensions. *SIAM Journal on Optimization, 9,* 112–147.

Laming, D. R. J. (1968). *Information theory of choice-reaction times*. Oxford: Acamedic Press.

Lo, C. C., & Wang, X. J. (2006). Cortico-basal ganglia circuit mechanism for a decision threshold in reaction time tasks. *Nature Neuroscience, 9,* 956–963.

Ly, A., Boehm, U., Heathcote, A., Turner, B. M., Forstmann, B., Marsman, M., et al. (in press). A flexible and efficient hierarchical Bayesian approach to the exploration of individual differences in cognitive-model-based neuroscience. In *Computational models of brain and behavior*. Wiley.

Maddox, W. T., & Bohil, C. J. (1998). Overestimation of base-rate differences in complex perceptual categories. *Attention Perception & Psychophysics, 60,* 575–592.

Mansfield, E. L., Karayanidis, F., Jamadar, S., Heathcote, A., & Forstmann, B. U. (2011). Adjustments of response threshold during task switching: A model-based functional magnetic resonance imaging study. *Journal of Neuroscience, 31,* 14688–14692.

Matzke, D., & Wagenmakers, E. J. (2009). Psychological interpretation of ex–Gaussian and shifted Wald parameters: A diffusion model analysis. *Psychonomic Bulletin & Review, 16,* 798–817.

Pelli, D. G. (1997). The VideoToolbox software for visual psychophysics: Transforming numbers into movies. *Spatial Vision, 10,* 437–442.

Rae, B., Heathcote, A., Donkin, C., Averell, L., & Brown, S. (2014). The hare and the tortoise: Emphasizing speed can change the evidence used to make decisions. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 40,* 1226–1243.

Raftery, A. E. (1995). Bayesian model selection in social research. *Sociological Methodology, 25,* 111–164.

Ratcliff, R. (1978). A theory of memory retrieval. *Psychological Review, 85,* 59–108.

Ratcliff, R., & McKoon, G. (2008). The diffusion decision model: Theory and data for two-choice decision tasks. *Neural Computation, 20,* 873–922.

Ratcliff, R., & Rouder, J. N. (1998). Modeling response times for two-choice decisions. *Psychological Science, 9,* 347–356.

Ratcliff, R., Smith, P. L., Brown, S. D., & McKoon, G. (2016). Diffusion decision model: Current issues and history. *Trends in Cognitive Science, 20,* 260–281.

Rossi, S., Hallett, M., Rossini, P. M., Pascual-Leone, A., & Safety of TMS Consensus Group. (2009). Safety, ethical considerations, and application guidelines for the use of transcranial magnetic stimulation in clinical practice and research. *Clinical Neurophysiology, 120,* 2008–2039.

Rouder, J. N., Speckman, P. L., Sun, D., Morey, R. D., & Iverson, G. (2009). Bayesian *t* tests for accepting and rejecting the null hypothesis. *Psychonomic Bulletin & Review, 16,* 225–237.

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,* 1916–1936.

Smith, P. L., & Ratcliff, R. (2004). Psychology and neurobiology of simple decisions. *Trends in Neurosciences, 27,* 161–168.

Tremblay, S., Lepage, J. F., Latulipe-Loiselle, A., Fregni, F., Pascual-Leone, A., & Théoret, H. (2014). The uncertain outcome of prefrontal tDCS. *Brain Stimulation, 7,* 773–783.

van Maanen, L., Brown, S. D., Eichele, T., Wagenmakers, E. J., Ho, T., Serences, J., et al. (2011). Neural correlates of trial-to-trial fluctuations in response caution. *Journal of Neuroscience, 31,* 17488–17495.

van Veen, V., Krug, M. K., & Carter, C. S. (2008). The neural and computational basis of controlled speed–accuracy tradeoff during task performance. *Journal of Cognitive Neuroscience, 20,* 1952–1965.

Vandekerckhove, J., & Tuerlinckx, F. (2007). Fitting the Ratcliff diffusion model to experimental data. *Psychonomic Bulletin & Review, 14,* 1011–1026.

Vandierendonck, A. (2017). A comparison of methods to combine speed and accuracy measures of performance: A rejoinder on the binning procedure. *Behavior Research Methods, 49,* 653–673.

Voss, A., Nagler, M., & Lerche, V. (2013). Diffusion models in experimental psychology. *Experimental Psychology, 60,* 385–402.

Voss, A., Rothermund, K., & Voss, J. (2004). Interpreting the parameters of the diffusion model: An empirical validation. *Memory & Cognition, 32,* 1206–1220.
Voss, A., & Voss, J. (2007). Fast-dm: A free program for efficient diffusion model analysis. *Behavior Research Methods, 39,* 767–775.
Watanabe, T., Hanajima, R., Shirota, Y., Tsutsumi, R., Shimizu, T., Hayashi, T., et al. (2015). Effects of rTMS over pre-supplementary motor area on fronto-basal-ganglia network activity during stop-signal task. *Journal of Neuroscience, 35,* 4813–4823.
Wenzlaff, H., Bauer, M., Maess, B., & Heekeren, H. R. (2011). Neural characterization of the speed–accuracy tradeoff in a perceptual decision-making task. *Journal of Neuroscience, 31,* 1254–1266.
Wiecki, T. V., Sofer, I., & Frank, M. J. (2013). HDDM: Hierarchical Bayesian estimation of the drift-diffusion model in Python. *Frontiers in Neuroinformatics, 7,* 14.