MODULATION OF EVIDENCE ACCUMULATION RATE AND DECISION CRITERIA VIA INHIBITION AND EXCITATION OF LEFT DORSOLATERAL PREFRONTAL CORTEX BY TRANSCRANIAL MAGNETIC STIMULATION

by

ESİN TÜRKAKIN

Submitted to the Graduate School of Social Sciences and Humanities

in partial fulfilment of the requirements for

the degree of Master of Arts in Psychology

at Koç University

August, 2019

Koc University

Graduate School of Social Sciences and Humanities

This is to certify that I have examined this copy of a master's thesis by

ESİN TÜRKAKIN

and have found that it is complete and satisfactory in all respects,
and that any and all revisions required by the final
examining committee have been made.

Committee Members:

Prof. Dr. Fuat Balcı

Asst. Prof. Çağlar Akçay

Assoc. Prof. Mine Mısırlısoy

Date: August 09, 2019

STATEMENT OF AUTHORSHIP

This thesis contains no material that has been accepted for any award or any other degree or diploma in any university or other institution. It is affirmed by the candidate that, to the best of her knowledge, the thesis contains no material previously published or written by another person, except where due reference is made in the text of the thesis.

Signed Esin Türkakın

ABSTRACT

Accumulation-to-bound models provide a neurally plausible mechanistic framework for the decision process. Evidence accumulation rate in these models depends both on the stimulus properties and the agent's abilities, but evades cognitive control. Dorsolateral prefrontal cortex (DLPFC) activity is one of the neural correlates of how well an agent integrates the available evidence. Supporting the correlational imaging studies, brain stimulation studies elicit reduced evidence accumulation rates by inhibiting this region via transcranial magnetic stimulation (TMS). This study aimed to (1) partially replicate these findings, (2) investigate whether this latent decision variable can also be improved via excitation of left DLPFC, and (3) explore potential effects of activity modulation in this region on any other decision parameters, particularly decision criteria. We did not replicate the drift-rate suppression effect due to inhibition, but did find a trend for the drift-rate enhancing effect of DLPFC excitation specifically when the signal to noise ratio of the sensory input was low. We also found that inhibition of the left DLPFC results in more cautious responding compared to when the same region was targeted with excitation. The novel findings of this study challenge the existing studies regarding the role of DLPFC in the decision process and adds a new perspective to the literature that investigates evidence accumulation and decision thresholds in isolation.

Keywords: Perceptual Decision Making; Transcranial Magnetic Stimulation; Computational Modelling; Drift Diffusion Model; Dorsolateral Prefrontal Cortex

ÖZET

Eşiğe doğru delil birikimi modelleri, karar verme süreçlerini anlamak için altta yatan sinirsel mekanizmalarla da uyumlu bir bakış açısı sunar. Bu süreçlerdeki delil toplama hızı hem uyaranın özelliklerine hem de karar vericinin becerilerine bağlıdır ve bilişsel olarak kontrolü mümkün değildir. Mevcut delilin etkili olarak bütünleştirilmesinin sinirsel eşleniklerinden biri dorsolateral prefrontal korteks (DLPFK) aktivitesi olarak görülmektedir. Bu yöndeki ilintisel görüntüleme çalışmalarını destekleyen beyin uyarım çalışmaları da bölgedeki aktivitenin azaltılması ile delil bütünleştirme hızının da düştüğünü göstermektedir. Bu çalışma (1) delil toplama hızını düşüren önceki beyin uyarım çalışmalarının kısmen tekrarlanmasını ve teyidini, (2) sol DLPFK aktivitesini artırarak ile delil toplamayı hızlandırmanın mümkün olup olmadığının araştırılmasını ve (3) bu bölgedeki aktivite değişimlerinin karar eşikleri gibi diğer karar parametrelerine etkisinin keşfini amaçlamaktadır. Çalışma sonucunda önceki çalışmalardaki sürüklenme hızı azaltma etkisi tekrarlanamamış, fakat sinyal seviyesi düşük (zor) kararlarda DLPFK aktivitesini artırıcı uyarım sürüklenme hızını artırmaya yönelik bir eğilim göstermiştir. Aynı zamanda DLPFK baskılanmasının, aynı bölgedeki aktivitenin artırıldığı duruma kıyasla daha temkinli karar stratejileri belirlenmesine sebep olduğu bulunmuştur. Bu çalışma mevcut araştırmalara karşıt bulguları ile DLPFK'nin karar verme süreçlerinin farklı aşamalarındaki rolünü tartışmakta ve bu aşamaları birbirinden ayrı değerlendiren alanyazına yeni bir perspektif sunmaktadır.

Anahtar Kelimeler: Algısal Karar Verme; Transkraniyal Manyetik Uyarım; Computational Modelling; Drift Diffusion Model; Dorsolateral Prefrontal Cortex

ACKNOWLEDGEMENTS

I would like to extend my gratitude to my thesis advisor Prof. Dr. Fuat Balcı for the opportunity to pursue this research question alongside several other related studies, none of which would have been possible without his stimulating mentorship and continued support throughout my studies. I would also like to thank my thesis committee members, Asst. Prof. Çağlar Akçay and Assoc. Prof. Mine Mısırlısoy, for their time and valuable feedback on this thesis.

I would like to thank Gamze, Bilgesu, and Ekin for their time and valuable assistance in conducting the study discussed in this thesis. I also thank all the participants for volunteering to take part in the experiments – the ideas put forth in this thesis were tested thanks to their time and diligent participation.

I am also deeply grateful to my labmates in the Timing and Decision-Making Lab and friends beyond for their emotional support through it all. Thanks for all the fun, inspiration, thoughts, challenges, discussions, and food. Cheers!

Finally, I thank mom and dad for their endless love, trust, and support. This thesis is dedicated to them, with hopes to make up for the time spent so close yet so far apart for the past three years.

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

1. INTRODUCTION

1.1. Background

Classifying noisy input from the environment is a task people undertake every day in different forms. Understanding what happens in the brain during these decisions is therefore useful for insight into a wide range of human behavior. Generative computational models of decision making provide valuable insight into mechanistic aspects of decisions that are traditionally investigated through descriptive outcomes, namely response time and accuracy. A widely used family of such models conceptualize decisions as noisy evidence accumulation up to a set boundary (for a review, see Ratcliff, Smith, Brown & McKoon, 2016). Neural correlates of the psychologically-meaningful parameters of this model is a rich research area, with imaging studies exploring potential regions of interest and brain stimulation studies modulating the activity in these regions to establish causal links. The rate of evidence accumulation in this process is an especially intriguing target, as it evades explicit cognitive control. There are imaging studies that show neural correlates of this parameter (Mulder, van Maanen & Forstman, 2014), as well as neuromodulation studies that cause a reduction in this parameter (Philiastides, Auksztulewicz, Heekeren, & Blankenburg, 2011, Georgiev, Rocchi, Tocco, Speekenbrink, Rothwell, & Jahanshahi, 2016), but none that boost it. In this study, we set out to replicate the previous brain stimulation findings, further investigate whether this latent decision variable can be boosted through excitation of the same area, and explore the effects of neuromodulation in this region to other model parameters, guided by a conflicting literature.

1.2 Drift Diffusion Model

The Drift Diffusion Model (DDM) integrates the two primary and interdependent behavioral outcomes of two alternative forced choice (2AFC) tasks (response time and accuracy) in a generative framework that underlies the observed decision behaviors (Ratcliff, 1978; Ratcliff, 2016; Forstmann 2015). In this model, noisy information accumulates over time at a certain rate (i.e. drift rate) towards one of two decision boundaries until either threshold is reached, at which point the agent executes the corresponding decision. The drift rate parameter indicates the efficiency with which evidence can be perceived and integrated, and the separation of the two thresholds corresponds to the decision strategy (i.e. level of caution). Investigating 2AFC task performance in terms of these parameters instead of observed behavioral data (latency and accuracy of decisions) provides a way to integrate these two interdependent measures to allow for easier and more intuitive interpretation.

Another advantage of using the drift diffusion model to as the generative process behind perceptual decision-making is its compatibility with neural dynamics in evidence accumulation in both humans and non-human animals. For example, neural recordings in the lateral intraparietal area (LIP) and frontal eye field (FEF) of monkeys show the neural populations in these regions increasing their firing rate until they hit a plateau as they observe a stimulus before choosing between two alternative responses, in a way one would expect from a decision particle of the drift diffusion model (e.g. Kim & Shadlen, 1999; Gold & Shadlen, 2007, Ding & Gold, 2013). The neural correlates of the drift diffusion model are also being investigated in humans. Drift rate, related to both the signal-to-noise ratio of the stimulus (i.e. task difficulty) and the ability of the individual to integrate evidence, seems to correlate with blood-

oxygen-level-dependent (BOLD) activity in various regions in the fronto-parietal network: BOLD activity in DLPFC is found to be positively correlated with the drift rate (Heekeren 2004, Rolls 2010, Philiastides 2007), whereas the right insular cortex, the frontal eye field (FEF) and intraparietal sulcus (IPS) are negatively correlated (Ho et al., 2009; Basten et al., 2010, Liu & Pleskac, 2011). Threshold setting, related to caution and therefore inhibition or disinhibition of responses under a response policy, is correlated with BOLD activity in regions of the frontostriatal network, particularly pre-supplementary motor area (pre-SMA), anterior cingulate cortex (ACC), striatum, and the sub-thalamic nucleus (STN) (Bogacz, 2010; Forstmann, 2008). Compatibility with neural dynamics makes DDM an especially useful explanatory tool for decision-making and the generative processes behind it.

1.3. Neuromodulation Studies

While correlational functional magnetic resonance imaging (fMRI) studies are useful as a rough guide, they do not provide a causal link between the neural activity and the observed behaviour, and the discrepancy of the time scale between the task and the measurements is a fundamental shortcoming in understanding what exactly is going on in the brain during decision-making and how it leads to behavior. Studies looking to modulate the activity of brain regions associated with DDM parameters to observe any corresponding changes in the parameters therefore provide an invaluable causal link to these associations. Threshold setting, being associated with a relatively well-studied network, has been the target of such studies within the drift diffusion model framework, albeit with inconsistent results. deHollander and colleagues have found no effect of inhibition of pre-SMA using transcranial direct current stimulation (tDCS) on decision thresholds, while Georgiev et al. (2016) have found that inhibition of pre-

SMA with a continuous theta burst stimulation (cTBS) protocol lowered the threshold, going against what the literature would predict. Recently, two studies (Tosun, Berkay, Sack, Cakmak & Balci, 2017; Berkay, Eser, Sack, Cakmak & Balci, 2018) have found that inhibition of the preSMA increases the threshold and results in more cautious responding and that excitation of the same area leads to the opposite effect, a finding that is in line with what the imaging literature predicts.

The relationship between the drift rate and the cognitive processes it relates to is still under investigation. The interdependence of task difficulty, motivation, and attention make it difficult to pinpoint the exact mapping between these cognitive processes, their neural correlates, and drift rate (Mulder et al., 2014). Despite these issues, drift rate is also emerging as a target parameter for neuromodulation. Philiastides et al. (2011) have demonstrated a causal effect of inhibiting the left DLPFC on reduction in drift rate in healthy individuals using a face-versus-car discrimination task. Georgiev et al. (2016) tested whether this was also possible by inhibition of the region on the right hemisphere and with a lower-level perceptual task, and found that inhibiting the right DLPFC with cTBS also resulted in a reduction in drift rate, although the effect was observed only in easy tasks (with high signal to noise ratio) and not harder ones (with low signal to noise ratio). These two studies, in combination, present a convincing, albeit unidirectional, case regarding the causal role of DLPFC in the rate of evidence accumulation in either hemisphere and across various visuo-perceptual tasks, at least at certain task difficulties.

In a more comprehensive study investigating the hierarchical progression of various aspects of decisions in the cortex, Rahnev et al. (2016) disrupted activity in the frontal

eye field (FEF), DLPFC and anterior prefrontal cortex (aPFC) (all in the right hemisphere) using cTBS in separate sessions. They argue that their findings are consistent with activity changes in rDLPFC affecting decision criteria (i.e. threshold), whereas FEF was associated with changes in drift rate and aPFC with changes in metacognition (self-evaluation of responses). It is worth noting that the findings of this latter study are based on RT differences and are not directly modeled, and that the conclusions on model parameters stem from the researchers' simulation of comparable data using a diffusion model with two accumulators (which introduces a metacognitive aspect). However, while recognizing that the model and the analysis may not be directly comparable to the previous studies and the one we have conducted, there are correlational studies that suggest that DLPFC plays a role in decision criterion setting that lend support to Rahnev et al.'s findings (e.g. Ivanoff et al., 2008; Van Veen et al., 2008, Vallesi et al., 2012, Weigard et al., 2018). These studies provide a good reason to look out for and explore any effects of DLPFC stimulation and inhibition on threshold settings as well as drift rate.

The neuromodulation studies targeting drift rate have so far aimed exclusively to disrupt activity in areas positively correlated with the parameter in order to dampen it. Considering that the cognitive processes thought to underlie drift rate are attention, motivation and perceptual evidence integration efficiency in general, an increase in drift rate could potentially entail a valuable form of cognitive enhancement. Cognitive enhancement using noninvasive stimulation techniques such as tDCS and TMS is an increasingly active research area with popular interest in the outcomes (Luber 2013, Demeter 2016). As an example to a study suggesting that a boost to a cognitive process like attention is possible, He et al. (2013) have induced an improvement in alerting

and executive attention in an Attention Network Test (ANT) by stimulating the right dorsolateral prefrontal cortex (rDLPFC) of healthy subjects using an iTBS protocol. Since this region is also implicated by drift rate, and the inhibition of which has resulted in a reduction of drift rate with a cTBS protocol (Georgiev et al., 2016), this study provides an additional indication that the DLPFC is a suitable target for studying drift rate enhancement with TMS.

1.4. Aims and Hypotheses

In light of these previous findings, we designed this study to fill a missing link in the existing literature by modulating drift rate in healthy individuals by means of stimulating the left dorsolateral prefrontal cortex (DLPFC) with offline transcranial magnetic stimulation (TMS) before a perceptual decision task. The primary aim of the proposed study was to boost cognitive evidence accumulation efficiency for healthy individuals during a low-level perceptual task. A secondary aim was to replicate the previous findings demonstrating that drift rate can be reduced by inhibiting the DLPFC, thereby establishing a functional role of DLPFC in bidirectional modulation of drift rate in a single study. A third and final aim was to explore the effects of DLPFC activity on other DDM parameters, particularly decision threshold, as suggested by the literature.

We hypothesized that, in line with previous findings, cTBS targeting DLPFC would result in a reduction of drift rate especially in the easy version of the perceptual task where signal-to-noise ratio was high, providing room for reduction. Similarly, we also hypothesized that the iTBS protocol on the same area would increase the drift rate especially in the hard version of the task with a low enough signal-to-noise ratio that

allows for a boost to make a difference. The changes to the threshold settings were exploratory, but with a cautionary note that the changes in the two parameters may affect the outcomes of each other that we predicted in isolation, we hypothesized that the inhibition of DLPFC would lead to a higher threshold, and the excitation of the area would decrease the decision threshold. We did not have a prediction for how the threshold settings would interact with difficulty levels.

2. METHODS

2.1. Participants

The experiment was announced through a university-wide newsletter in Koc University and those who were interested were sent anonymous health forms. Health forms contained questions about the exclusion criteria; people with any personal and family history of neurological and psychiatric disorders, psychoactive drug use in the past month, or metal implants were excluded. We recruited only right-handed people between the ages of 18 and 30 in order to minimize potential variance in brain structure and function. 40 healthy adult volunteers took part in the study for at least one session, of which 35 completed two sessions (21 Females, aged 20.03 ± 1.24) and 25 completed all four sessions (16 Females, aged 19.96 ± 1.40). All participants gave written consent for participation in the study and received 20 TL per stimulation session and 10 TL per baseline session without stimulation as compensation for their time. The participants were asked to attend all sessions well-rested and to avoid ingesting any stimulants such as caffeine or nicotine 4 hours prior to testing.

2.2. Design

The study took place over four sessions on separate days. The first session only had behavioral testing (Dot Motion Discrimination Task, detailed below) with no stimulation, and the remaining three sessions consisted of a transcranial magnetic stimulation protocol, followed by behavioral testing. The randomization of the order of stimulation protocols (inhibition, excitation, and sham) for these three sessions was conducted in two steps: (1) the condition for the second session followed a pseudo-random pattern, with each condition appearing once in every three-participant sequence, and (2) the third and fourth sessions alternated in an orderly fashion between the two remaining alternative stimulation conditions, if the participant decided to continue after the second session.

The data from the first behavioral session was used in the single session group-comparison analyses as a comparison point for the effects of stimulation in the following session. One participant was removed from this analysis due to having an extremely slow median response time in the baseline session across difficulty levels (3.67s, $z = 7.6$). The baseline session was not used for the within-subjects analyses, eliminating the initial learning phase as a confound for the analysis. The decision to plan and randomize for two separate analyses was in anticipation of a high drop-out rate after the second session due to either scheduling issues or potential intimidation by the stimulation procedure, and aimed to make use of the data from participants with only two sessions in a between-subjects analysis. We present both analyses as planned, but note that the single session analyses that use between-subjects comparisons are highly under-powered due to high inter-individual variability in baseline performance, learning and TMS responsivity.

The a priori power analysis were conducted based on the within-subject effects found in Georgiev et al. (2015), and indicated that twenty-four participants taking part in all four sessions would be sufficient for %80 power, assuming that stimulation order has no effect in our design. We note that our design and analyses differ from the study that constituted the basis of the expected effects in the power analysis, but also note that that same study used only fifteen participants and still found significant effects.

Table 1. Number of participants in each stimulation order condition for both analyses.

	Sham		**Inhibition**		**Excitation**		**Total**
Single Session	11		11		12		34
Ses. 3 Cond.	**Inh.**	**Exc.**	**Shm.**	**Exc.**	**Shm.**	**Inh.**	
Three Sessions	4	3	5	5	4	4	25

2.3. Task

Dot motion discrimination task consisted of white dots (2 px) moving on a black background within a circular field with a diameter of 10 visual degrees. Most of the dots were displaced in random directions at each frame, while a certain portion of the dots moved coherently in one direction, either to the left or to the right. The task of the participant was to indicate the direction of the coherent dots as fast and as accurately as possible in a free-response paradigm. Correct answers were followed by an auditory positive feedback, whereas errors received no feedback. The response-to-stimulus intervals (RSI) were sampled from a uniform distribution with a range of 0.5-1.5s. The task was written and presented in MATLAB (2015b) using Psychtoolbox (v3.0.14,

Brainard, 1997) on an iMac computer (late 2012, 21.7 inch, OSX 10.9.5), and responses were collected via a mechanical keyboard (Zalman ZM-K500). The participants were seated approximately 85 cm away from the monitor and were instructed to use both hands when responding (any finger of the left hand on the A key to respond left, and any finger of the right hand on the L key to respond right).

The percentage of coherent dots determines the objective difficulty of the task. We used three different coherence levels: high (35%), moderate (15%), and low (5%), where each block of 50 trials had a constant coherence level throughout. The coherence levels were chosen to test the interaction of DLPFC-related drift rate changes and difficulty previously found in the Georgiev et al. (2016) study, where inhibition caused a drop in drift rate in high (35% and 50%) but not in lower coherence levels. We presented the coherences in a block design (instead of in interleaved trials) in order to observe any potential changes in decision thresholds differentially for the coherence levels. The participants were not cued about the coherence levels at the beginning of each block, but were informed in the initial instructions that there were three difficulty levels that would come up in random order and remain constant within blocks.

In all sessions, the participants completed one practice block of each coherence level in fixed order (high-moderate-low). The data from practice blocks was not used for the analyses. In the baseline session, the practice part was immediately followed by the experimental part of twelve blocks with a pseudo-random order of coherences, where all three coherences appeared in random order in four groups of three blocks. In the other sessions, the practice part was followed by the stimulation procedure

(detailed below), after which the participant was asked to complete the twelve experimental blocks of the computerized task. In each session, participants completed a total of 750 trials, with 200 non-practice trials per difficulty level.

2.4. Transcranial Magnetic Stimulation

2.4.1. Equipment

We used Magstim Super Rapid^2 magnetic simulator with a 70-mm figure-of-eight coil for all TMS sessions. Region localization was made using the 10-20 system (g.GAMMAcap, G.Tec Medical Engineering GMBH, Austria). Despite being a region spanning a large area containing three Brodmann's Areas (8, 9, 46), the dorsolateral prefrontal cortex (DLPFC), is often targeted using the F3 (left) and F4 (right) electrode sites in tDCS and TMS studies (e.g. Georgiev et al. 2016, also see Herwig, Satrapi & Schonfeldt-Lecuona, 2003), and we stimulated over F3 following this convention.

2.4.2. Motor threshold

In the second session (the first session with stimulation), we determined the active motor threshold of the participants using single pulse stimulation over the left motor cortex (C3), with the coil angled at 45 degrees to the midline and the coil handle pointing backwards. Initially, the region was stimulated several times at 65% of the maximum stimulator intensity with slight changes to the coil location and/or increased intensity (up to a maximum of 80%) until the motor hotspot was detected and a movement was observed on the right hand. This initial search was conducted at a higher-than-average intensity in order to ensure correct coil localization, and participants where no movement could be elicited at 80% intensity were excluded from the study. Once a movement was elicited on the right hand, the coil location was

stabilized, and stimulation intensity was dropped by fifteen percentage points to start the staircase protocol. C3 was then stimulated at each intensity, increasing by five percentage points until a movement was observed on the right hand for more than five out of ten pulses, and dropped by one percentage point until a movement was observed for five or fewer out of ten pulses. The final motor thresholds had a range of 44-74 and a mean of 60.82. In the third and fourth sessions, the motor cortex was stimulated at the motor threshold for ten pulses in order to reduce the perceived difference between sessions for the participants, but no changes were made to the motor threshold depending on the observed motor output in these sessions.

2.4.3. Targets and Stimulation Protocol

We chose theta burst stimulation protocols over varying frequencies of repetitive TMS, as TBS protocols allow applying many pulses within a much shorter time frame, with longer lasting effects (Huang et al. 2005). These advantages make TBS a safer, more practical, more effective and less variable option compared to rTMS. For the inhibition sessions, we applied continuous theta burst stimulation (cTBS) over the F3 electrode site, with three 50Hz pulses repeated every 200 milliseconds continuously for 40 seconds (Huang et al., 2005). For the excitation sessions, we applied intermittent theta burst stimulation (iTBS) over the F3 electrode site, with 2-second trains of three 50Hz pulses every 10 seconds for 190 seconds (Huang et al., 2005). For the sham sessions, we applied intermediate theta burst stimulation (imTBS) to the vertex (Cz), in order to avoid any task-relevant brain region while preserving most of the peripheral sensation of the stimulation on the scalp. The imTBS protocol is shown to not have any inhibitory or excitatory effects, and was chosen to both avoid the risk of modulating the activity of the nearby sensory and motor areas, as well as to provide a

different auditory output than the other two protocols and avoid creating misleading expectations.

For all stimulation conditions, a total of 600 pulses were applied. All stimulations were done at 80% of the participant's active motor threshold. For those with a stimulation intensity higher than 50%, the theta burst frequency was automatically readjusted from 50 Hz (1 Hz drop every 2 percentage point increases in intensity, down to a minimum of 45 Hz in our sample) due to stimulator constraints. Before starting the full stimulation protocol, we applied the theta burst trains for less than a second over the stimulation area in order to acclimate the participants and to check for any discomfort. If any discomfort was reported, we changed the coil angle (from 45 degrees to the anterior posterior midline) until the participant reported no severe discomfort. The cTBS and iTBS protocols were expected to have effects lasting up to 60 and 20 minutes respectively, and the experimental part of the behavioral task was designed to take an average 20 minutes immediately following stimulation in order to be completed before the neuromodulatory effects tapered off.

2.5. Data Analysis

2.5.1. Behavioral Outcomes

We analyzed the two descriptive behavioral outcomes with linear mixed effects models, separately for the three difficulty levels. All models had trial-by-trial data as input and random intercepts for each participant. The main three-session analyses compared the response time (RT) and accuracy in the second, third and fourth sessions, in which the participants received the inhibition, excitation, and sham TMS protocols in counterbalanced order. The effect of interest was the effect of TMS protocol

separately in each difficulty level, but session order of each condition was also included in the models as a continuous and centered within-subject fixed variable to control for any learning over the sessions. The single-session analyses compared the RT and accuracy changes from the first (baseline) session to the second session for three groups of participants that received one of three different TMS protocols in the second session. The linear mixed effects models were conducted in jamovi (The jamovi project, 2019), using the gamlj plugin (Galluci, 2019). Holm-Bonferoni corrections were used for multiple comparisons where appropriate.

2.5.2. Model Parameter Estimation

For DDM parameter estimation and comparison, we used hierarchical drift diffusion model (HDDM), a tool that conducts a hierarchical Bayesian estimation of the parameters with Markov Chain Monte Carlo (MCMC) sampling to approximate the posterior distributions of the parameters (Wiecki, Sofer & Frank, 2013). We chose HDDM over other DDM parameter estimation tools as it increases power (Wiecki et al., 2013, Ratcliff & Childers, 2015), uses priors informed by published decision making studies in the literature compiled by Matzke & Wagenmakers (2009), and provides the whole (approximated) posterior distribution of the parameters instead of a single estimate. Models that included different combinations of the drift rate (v), threshold (a), non-decision time (t) and inter-trial drift-rate variability (sv) (Table 2) varying with the stimulation condition (sham/inhibition/excitation), as well as a null model, were fit separately to data from three difficulty levels. The best model was chosen using deviance information criterion (DIC) scores of the models for each difficulty, with a reduction of 10 indicating a justified increase in model complexity. Once the best model was chosen, posterior distributions of the parameters were

compared to determine whether inhibition and excitation had an effect on the drift rate or threshold compared to the sham and no-stimulation (baseline) conditions at each coherence level. Similarly to the behavioral analyses, two separate analyses were conducted, one to see the effect of a single TMS protocol over the baseline session as a between-subjects comparison, and another to see the effects of three different TMS protocols over three sessions as a within-subjects comparison.

Table 2. Parameters and DIC scores of compared models.

Model Name	Included Parameters	Dependent Parameters	DIC (Easy)	DIC (Moderate)	DIC (Hard)
M0	a, v, t	-	-14502	-2717	21015
M1	a, v, t	v	-14527	-2731	21012
M2	a, v, t	v, a	-14525	-2735	20994
M3	a, v, t, sv	v, a	-14946	-3357	20132
M4	a, v, t, sv	v, a, t	-14961	-3355	20130

For all HDDM fits, trials with RTs faster than 100 ms were excluded as these responses were too fast to have come from an actual decision process and thus were not informative for the parameter estimation. We also excluded the first five trials of each block, as the participants were not informed about the difficulty level of the upcoming block and thus would have needed a few trials to gauge the coherence level adjust their decision strategies accordingly. The MCMC sampling for all models used 5000 samples and discarded the first 20 as burn-in. Parameter convergence was evaluated using the sampling traces, autocorrelations and distributions of the MCMC chains. The posterior distributions of the parameters were compared on a sample-by-sample basis, with the end result of each comparison indicating the percentage of the samples where the relevant hypothesis held true (denoted with a capital P). We used a significance

cutoff of 0.95 by convention, but it is important to note that these values denote a fundamentally different and intuitive probability value than the frequentist p value and the cutoffs should be interpreted as such.

3. RESULTS

3.1. Behavioral Outcomes

3.1.1. Three-Sessions, Within-Subjects

3.1.1.1. Accuracy

For the three-session linear mixed model analyses of accuracy, there was no significant interaction between stimulation condition and session progress at any of the three difficulty levels. For the moderate and hard levels, but not for the easy level, there was a positive main effect of session progress on accuracy (mod: β_{ses}=0.0124, SE=0.0031, $p<0.001$; hard: β_{ses}=0.0121, SE=0.0047, $p<0.05$) pointing to perceptual learning, with a 1.2 percentage point increase in accuracy in consecutive sessions regardless of stimulation condition. The main effect of stimulation condition on accuracy was not significant at any difficulty level.

3.1.1.2. Response Time

The interaction of session progress and stimulation condition on response times was also non-significant for moderate and hard levels but was significant for the easy level (Figure 1). We observed a negative main effect of session progress on response time for all three difficulty levels (easy: β_{ses}=-0.0248 SE=0.0029, $p<0.001$; mod: β_{ses}=-0.051, SE=0.0039, $p<0.001$; hard: β_{ses}=-0.119, SE=0.0074, $p<0.001$), pointing to increasing response speed in consecutive sessions due to perceptual learning. We also observed a significant main effect of stimulation condition on RTs in all three

difficulty levels. Given the non-interaction of the learning effect with that of stimulation condition in moderate and hard levels, we conducted post-hoc z-tests for the condition effect in these levels. In the moderate level, regardless of session progress, the response time estimate after both inhibition and excitation sessions were faster than that after sham stimulation (Exc-Shm: β=-0.0318, SE=0.0077, p_{Holm}<0.001; Inh-Shm: β=-0.0435, SE=0.0077, p_{Holm}<0.001), but inhibition and excitation RT estimates were not significantly different than each other (Exc-Inh: β=-0.0117, SE=0.0077, p_{Holm}=0.129). In the hard condition, controlling for perceptual learning over sessions, the RT estimate after excitation was significantly faster than that after sham (β=-0.0409, SE=0.0148, p_{Holm}=0.018) and inhibition (β=-0.037, SE=0.0149, p_{Holm}=0.026), while the difference between inhibition and sham sessions was not significant (β=0.0038, SE=0.0149, p_{Holm}=0.798).

In the easy task, where we did observe an interaction between session progress and stimulation condition, we investigated the simple effect of TMS condition at each session. In the first stimulation session, RTs after both excitation and inhibition were faster than those after sham (Exc-Shm: β=-0.0714, SE=0.0122, p<0.001; Inh-Shm: β=-0.0509, SE=0.0116, p<0.001), but the 95% confidence intervals (CI) of the excitation and inhibition effects (Exc: [-0.0954 -0.0475], Inh: [-0.0736 -0.0283]) did not exclude the estimates of each other, indicating no differential effect of the two stimulation conditions. In the second stimulation session, both excitation and inhibition resulted in faster RTs compared to the sham condition (Exc-Shm: β=-0.0370, SE=0.0059, p<0.001; Inh-Shm: β=-0.0238, SE=0.0059, p<0.001), and the 95% CI of the two effects (Exc: [-0.0486 -0.0255], Inh: [-0.0254 -0.0122]) excluded the estimates of each

other, pointing to faster RTs after excitation compared to inhibition. The third stimulation session showed no effect of stimulation condition. Note that while both variables of this interaction were within-subject, this simple effect ends up being a between-subjects comparison, as each participant received only one of the stimulation conditions at a given session.

3.1.2. Single-Session – Between-Subjects

3.1.2.1. Accuracy

In the single-session linear mixed effects model analyses (Figure 2, right panel), the interaction of session and stimulation group effects on accuracy was significant for the moderate ($F(2,12206)=5.485$ $p=0.004$) and hard ($F(2,12196)=6.24$, $p=0.002$) conditions, but not for the easy condition ($F(2,12206)=2.112$, $p=0.121$). In the easy level, there was a significant positive main effect of session on accuracy ($\beta=0.0087$ $SE=0.0039$, $p=0.025$), pointing to perceptual learning regardless of the stimulation condition. For the moderate and hard levels, where we observed an interaction between the two variables, we investigated the simple effect of session (between baseline and stimulation sessions) for each stimulation group. In the moderate level, while neither the inhibition (Post-Pre: $\beta=0.0157$, $SE=0.0095$, $p=0.101$) nor the excitation group (Post-Pre: $\beta=0.0157$, $SE=0.0091$, $p=0.085$) had a significant change in accuracy in the stimulation session over the baseline, the sham group showed a significant reduction in accuracy (Post-Pre: $\beta=-0.0227$ $SE=0.0095$, $p=0.017$). In the hard level, while both the inhibition (Post-Pre: $\beta=0.0747$, $SE=0.0148$, $p<0.001$) and the excitation (Post-Pre: $\beta=0.0793$, $SE=0.0142$, $p<0.001$) group showed a similar significant improvement in accuracy in the stimulation session, the sham group did not show a significant change

(Post-Pre: β=0.0135, SE=0.0148, p=0.362). The change in inhibition and excitation groups were not significantly different from each other.

3.1.2.2. Response Time

The linear mixed effects models on response time (Figure 2, left panel) show a significant interaction of session and stimulation group in all three difficulty conditions (easy: $F(2,12206)=36.906$ $p<0.001$; mod: $F(2,12206)=3.2$ $p<0.041$; hard: $F(2,12296)=40.585$ $p<0.001$). The simple effects investigating this interaction in the easy level indicate that both the inhibition (Post-Pre: β=-0.0736, SE=0.0109, p<0.001) and the excitation group (Post-Pre: β=-0.0981, SE=0.0104, p<0.001) responded faster in the stimulation session compared to their baseline performance, while the sham group performed significantly slower in the stimulation session (Post-Pre: β=0.0252, SE=0.0109, p=0.02). The improvement in RT for the excitation group (95%CI: [-0.1184 -0.0776]) was significantly higher than that for the inhibition group (95%CI: [-0.0949 -0.0523]). In the moderate difficulty level, all stimulation groups responded significantly faster in the stimulation session compared to the baseline (Post-Pre, Inh: β=-0.244, SE=0.02, p<0.001; Exc: β=-0.191, SE=0.019, p<0.001; Shm: β=-0.176, SE=0.02, p<0.001), but the inhibition group showed significantly higher speeding than the other two groups (95%CI: [-0.282 -0.205]). In the hard level, all three stimulation groups had significantly faster RTs in the stimulation session compared to the baseline (Post-Pre, Inh: β=-0.62, SE=0.032, p<0.001; Exc: β=-0.446, SE=0.030, p<0.001; Shm: β=-0.218, SE=0.032, p<0.001), with the inhibition group having the most improvement (95%CI: [-0.682 -0.558]) and the sham group having the least (95%CI: [-0.280 -0.156]).

3.2. DDM parameters

All parameters in the models showed acceptable convergence. The best model, used for all difficulty levels, had drift rate (v), threshold (a), and non-decision time (t) varying with stimulation condition and included a drift rate variability (s) parameter fitted as a constant for all conditions and participants (see Table 2 for model comparisons).

3.2.1. Three-Sessions – Within-Subjects

Please consult Figure 3 for the posterior distributions that are reported in this section. The actual values of the parameters for the Sham condition (the intercept for the model) are presented in Table 3.

Table 3. Means of posterior distributions for parameters at the intercept (Sham).

	a	v	t	sv
Easy	1.675	3.308	0.261	0.867
Moderate	1.756	2.443	0.261	0.997
Hard	1.920	1.015	0.242	1.505

3.2.1.1. Drift rate

For participants who completed all three stimulation sessions, the drift rate in both inhibition and excitation sessions was equally higher than that in sham sessions in easy (Inh>Shm: P=0.995; Exc>Shm: P=0.989; Exc>Inh: P=0.374) and moderate (Inh>Shm: P=0.995; Exc>Shm: P=0.953; Exc>Inh: P=0.186) levels. In the hard level, the drift rate in the excitation session showed a trend to be higher than the drift rate in

both the sham and inhibition sessions (Exc>Shm: P=0.937; Exc>Inh: P=0.928; Inh>Shm: P=0.441).

3.2.1.2. Threshold

In all three difficulty levels, the threshold in the inhibition sessions was higher compared to the excitation sessions (Inh>Exc, easy: P>0.999; mod: P=0.996; hard: P>0.999). In the easy and hard levels, the threshold in inhibition sessions was higher than that in sham sessions (Inh>Shm, easy: P>0.999; hard: P>0.999), and in the moderate level, the excitation resulted in a lower threshold compared to the sham sessions (Exc<Shm, mod: 0.996).

3.2.1.3. Non-Decision Time

In all three difficulty levels, the non-decision times after inhibition and excitation sessions tended to be lower than those in the sham sessions. The inhibition session resulted in notably lower (faster) non-decision times compared to sham and excitation sessions in the easy and hard levels (Inh<Shm, easy: P>0.999; hard: P=0.995; Inh<Exc, easy: P=0.999 ; hard: P=0.953). The magnitude of differences in non-decision times between sessions were negligible (the highest mean difference, between inhibition and sham in the easy level, was 10.1 ms), when taken in context with the average frame duration of the monitor (60 Hz refresh rate, 16.67ms frame duration).

3.2.2. Single-Session – Between-Subjects

For participants that completed the first two sessions, the comparisons between the two sessions and different stimulation groups did not reveal any differences for easy and moderate levels (Figure 4). In the hard level, there was a significant increase in

drift rate and a reduction in threshold setting from the baseline to the stimulation sessions for the inhibition and excitation groups (Drift: Post>Pre, Inh: P=0.98; Exc: P=0.955; Threshold: Post<Pre, Inh: P=0.986; Exc: P=0.939), but not for the sham group (Drift, Post>Pre: P=0.75; Threshold: Post<Pre: P=0.739). However, the difference between the sessions were not different between the groups (Drift: Exc_{diff}>Shm_{diff}: P=0.815, Inh_{diff}>Shm_{diff}: P=0.781; Threshold: Exc_{diff}>Shm_{diff}: P=0.732, Inh_{diff}>Shm_{diff}: P=0.877).

4. DISCUSSION

This study demonstrated effects of excitatory and inhibitory TMS protocols in both behavioral and model outcomes of a perceptual decision-making task. While brain stimulation did not result in systematic differential changes in accuracy, we observed diverging response times under different stimulation conditions, particularly when the signal-to-noise ratio was low. The inhibition protocol resulted in the most speeding over the baseline after a single stimulation session, while the excitation protocol resulted in the fastest responses when only the stimulation sessions were compared, controlling for any learning-related improvement over the sessions. The drift diffusion modelling of the data establishes differential effects of TMS on particularly drift rate and threshold parameters at different difficulty levels. While we observe improved evidence integration after both excitation and inhibition compared to the sham condition in easy and moderate levels, the hard task seems to differentiate between excitation and inhibition, with a trend for the drift rate after excitation to be higher than both inhibition and sham conditions. We also show the threshold parameter to be higher after inhibition compared to excitation in all difficulty levels, and higher than the sham condition in easy and hard levels. Overall, the changes to decision strategies in response to inhibition and excitation of the left DLPFC seem to be the most

consistent across difficulty levels, but this study cannot dissociate this effect from changes to evidence accumulation rate due to the same stimulation.

The results of this study present some discrepancies with the previous literature. One of the key differences is the lack of a reduction of drift rate due to inhibition of DLPFC, which was a finding presented in two independent studies (Philiastides et al., 2011, Georgiev et al., 2016). On the contrary, we have observed a significant *increase* in drift rate after inhibition, even in the easiest condition in which we initially expected a ceiling effect that might prevent any enhancing effects of excitation, due to the task being too easy to leave any room for improvement. We did, however, find evidence for the novel hypothesis that stemmed from those non-replicated findings, namely that excitation of the same region would lead to an enhancement of the evidence accumulation rate in a task where the signal to noise ratio was low. This study did not aim to be an exact replication of any previous study (it used a different task than Philiastides et al. (2011) and stimulated the opposite hemisphere of Georgiev et al. (2016)), but the complete contradiction of the previously demonstrated effect of DLPFC inhibition was striking. We do not consider differences in task and hemisphere choice to be the cause of this discrepancy, as the previous literature points to the association between drift rate and left DLPFC to be generalizable across tasks (Heekeren, Marrett, Ruff, Bandettini & Ungerleider, 2006), and particularly to the task that we have used. One main reason for this discrepancy may be the approach to modeling – neither of the previous neuromodulation studies tested models where decision threshold also depended on stimulation condition alongside drift rate. While the literature shows a clear connection of DLPFC activity and drift rate, some studies also point to the possibility of an association of this region with top-down control of

the entire decision process, and thus threshold setting. We do see consistent changes to decision thresholds as a result of stimulation in our, and the omission of this parameter would compromise the validity of the other parameter estimates. If DLPFC plays a role in both integrating evidence and setting decision boundaries, their combined effect may well be expected to differ from their isolated effects.

The most consistent result we have observed was increased cautious responding after inhibition of the left DLPFC, compared to the decision strategies after excitation, regardless of task difficulty. While there were studies that hinted at a connection between DLPFC activity and decision strategy setting, previous brain modulation studies either targeted a different brain region (Tosun et al., 2017, Berkay et al., 2018), a different parameter (Philiastides et al., 2011, Georgiev et al., 2016), or different outcome measures (Rahnev et al., 2016). Studies by Berkay et al. (2018) and Tosun et al. (2017) clearly show that preSMA inhibition and excitation respectively increase and decrease decision thresholds, similar to the patterns we observe with DLPFC stimulation. These two regions have previously been suggested as the source of cortical signals in two alternative theories of speed-accuracy trade-off (SAT) modulation. PreSMA is the cortical center that provides input to the striatum which in turn modulates the level of control on actions by the basal ganglia in the *striatal theory*. DLPFC is considered to be the main source of input to cortical integrator neuron populations in the *cortical theory*, where decision strategies are controlled by modulating the baseline activity of integrator neurons and thus the distance to threshold (for a review of SAT modulation theories, see Bogacz, Wagenmakers, Forstmann, 2010). These two theories make similar predictions with regards to decision outputs in response to modulation of the relevant targets, therefore seeing

congruent results with preSMA modulation studies is not surprising. As an alternative to considering two alternative theories, a recent study by Weigard et al. (2018) pointed to a directional connection between preSMA and DLPFC, suggesting that DLPFC has top-down control over preSMA which in turn acts as a hub for threshold setting. This is also in line with the posterior-to-anterior hierarchical structure suggested by Rahnev (2017), building on the findings of Rahnev et al. (2016) that pins down DLPFC as a region responsible for criterion setting and having top-down control over more posterior regions. These recent findings may link and unify the SAT modulation theories (also see Standage, Blohm, Dorris, 2014), and our study informs these by providing a causal link to the relationship between threshold setting and DLPFC activity. However, it is important to note that this study does not dissociate the effects of DLPFC modulation on evidence accumulation rate and decision criteria, and therefore cannot suggest that this region is solely responsible for one or the other aspect of the decision process.

The effect of DLPFC activity modulation on both drift rate and threshold parameters is a novel finding, however there are both behavioral (Rae, Heathcote, Donkin, Averall & Brown, 2014) and neuromodulatory (Tosun et al., 2017, Berkay et al., 2018) studies pointing to verbal task instruction (e.g. accuracy emphasis) and brain stimulation (e.g. preSMA inhibition) that lead to higher thresholds also result in an increase in drift rate. However, the changes reported in previous studies are not entirely consistent with our data. While we do partially see such coordinated effects in easy and moderate levels in our data (increased drift rate and threshold after inhibition of DLPFC), the hard level (with 5% coherence, closest to the 8% used in Tosun et al., 2017 and Berkay et al. 2018) has the two parameters and stimulation conditions dissociated: the excitation of

DLPFC leads to an increase in drift rate but no change in threshold setting, and inhibition leads to an increase in threshold but no change in drift rate. Our original predictions for the DLPFC effect on drift rate, guided by previous studies and partially supported only in the low signal-to-noise ratio scenarios, were also in the opposite direction to the preSMA modulation and behavioral findings (which would predict a reduction in drift rate following DLPFC excitation that also lowered decision thresholds). This dissociation warrants further study and potentially points to differential roles of DLPFC and preSMA in the decision process rather than a direct top-down connection.

This study has several limitations that warrant caution when interpreting its results. One key issue is that the counterbalancing scheme did not result in an equal weight in all stimulation order conditions due to unbalanced drop-outs. This was likely not a result of the nature of the stimulations, as the participants mostly reported dropping out due to scheduling issues on their part and not a problem that they had with the particular stimulation protocol they were exposed to. The uneven distribution of participants to order conditions presents an issue particularly when combined with the presence of a learning effect over sessions, since more subjects that were exposed to a condition in the first as opposed to the last session could skew the outcomes of that condition towards slower response times and lower accuracies. The results of these analyses will be more robust once the order conditions are balanced out, assuming that the stimulation conditions do not interfere with the learning process differentially.

Another limitation of the study is the craniometric localization of the target region, as opposed to MRI-, or better yet, fMRI-guided stimulation. Using the 10-20 system for

localization is a cheap and practical option but results in a trade-off with precision compared to methods that take into account the structural and functional idiosyncrasies of the participants' brains. However, the DLPFC is one of the most often targeted regions in brain stimulation studies, and localization studies indicate that the F3 electrode site is a decent approximation (Herwig et al. 2003) for the part of the DLPFC that is of most interest to us (i.e. Broadmann Areas 8 and 9) based on fMRI studies (Heekeren et al., 2008). A related limitation, shared by the previous studies as well, is the inter-individual variance in responses to TMS stimulation – while cTBS and iTBS are shown to inhibit and excite the areas that they target in general (Huang et al., 2005), some individuals simply do not respond or even respond in the opposite direction to these protocols (Hamada, Murase, Hasan, Balaratnam & Rothwell, 2013). This study does not measure the changes to neural activity after the stimulations and simply assumes that the intended modulation does take place, and therefore can only realistically claim to show a relationship between model parameters and stimulation type, not directly neural activity.

5. CONCLUSION

The current study links several aspects of the current computational modeling and neuromodulation literature together and challenges some of the previous findings. We show that, contrary to what previous studies found, inhibition of the left DLPFC does not lead to a decrease but an increase in drift rate in easy and moderate difficulty tasks compared to sham stimulation. Excitation of the area behaves identically to inhibition with regards to increasing evidence integration efficiency in low and medium signal-to-noise ratio scenarios, but also shows a trend for leading to higher drift rates compared to both sham and the inhibition conditions specifically in the hard task, providing evidence towards our main hypothesis. We also show a bi-directional effect

of DLPFC stimulation on decision thresholds: inhibiting the DLPFC leads to more cautious decision-making compared to the excitation condition in all difficulty levels, mimicking the effects of preSMA stimulation on decision strategies. These results provide partial evidence for the proposed roles of DLPFC in both evidence integration and top-down control over the decisions and hopes to guide further studies in dissociating these roles and better elucidating the neural correlates of decision-making.

6. REFERENCES

Basten U, Biele G, Heekeren HR, Fiebach CJ (2010) How the brain integrates costs and benefits during decision making. Proc Natl Acad Sci U S A 107:21767–21772.

Berkay, D., Eser, H. Y., Sack, A. T., Çakmak, Y. Ö., & Balcı, F. (2018). The modulatory role of pre-SMA in speed-accuracy tradeoff: a bi-directional TMS study. *Neuropsychologia*, *109*, 255-261.

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., & Vision, S. (1997). The psychophysics toolbox. *Spatial vision*, *10*, 433-436.

Demeter, E. (2016). Enhancing Cognition with Theta Burst Stimulation. *Current Behavioral Neuroscience Reports*, *3*(2), 87-94.

Ding, L., & Gold, J. I. (2013). The basal ganglia's contributions to perceptual decision making. *Neuron*, *79*(4), 640-649.

Domenech, P., & Dreher, J. C. (2010). Decision threshold modulation in the human brain. Journal of Neuroscience, 30(43), 14305-14317.

Erhan, C., Bulut, G. Ç., Gökçe, S., Ozbas, D., Turkakin, E., Dursun, O. B., ... & Balcı, F. (2017). Disrupted latent decision processes in medication-free pediatric OCD patients. *Journal of affective disorders*, *207*, 32-37.

Forstmann, B.U., Dutilh, G., Brown, S., Neumann, J., Von Cramon, D.Y., Ridderinkhof, K.R., & Wagenmakers, E.J. (2008). Striatum and pre-SMA facilitate decision-making under time pressure. Proceedings of the National Academy of Sciences of the United States of America, 105, 17538-42.

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. New York: Springer. p. 139-156.

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

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*, *9*(4), 601-608.

Gold, J. I., & Shadlen, M. N. (2007). The neural basis of decision making. Annu. Rev. Neurosci., 30, 535-574.

Hamada, M., Murase, N., Hasan, A., Balaratnam, M., & Rothwell, J. C. (2012). The role of interneuron networks in driving human motor cortical plasticity. *Cerebral cortex*, *23*(7), 1593-1605.

He, X., Lan, Y., Xu, G., Mao, Y., Chen, Z., Huang, D., & Pei, Z. (2013). Frontoparietal regions may become hypoactive after intermittent theta burst stimulation over the contralateral homologous cortex in humans. Journal of neurophysiology, 110(12), 2849-2856.

Heekeren HR, Marrett S, Bandettini PA, Ungerleider LG (2004) A general mechanism for perceptual decision-making in the human brain. Nature 431:859–862.

Heekeren, H. R., Marrett, S., Ruff, D. A., Bandettini, P. A., & Ungerleider, L. G. (2006). Involvement of human left dorsolateral prefrontal cortex in perceptual

decision making is independent of response modality. *Proceedings of the National Academy of Sciences*, *103*(26), 10023-10028.

Heekeren, H. R., Marrett, S., & Ungerleider, L. G. (2008). The neural systems that mediate human perceptual decision making. Nature reviews neuroscience, 9(6), 467-479.

Herwig, U., Satrapi, P., & Schönfeldt-Lecuona, C. (2003). Using the international 10-20 EEG system for positioning of transcranial magnetic stimulation. Brain topography, 16(2), 95-99.

Ho TC, Brown S, Serences JT (2009) Domain general mechanisms of perceptual decision making in human cortex. J Neurosci 29:8675–8687.

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(2), 201-206.

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

Kim, J. N. & Shadlen, M. N. Neural correlates of a decision in the dorsolateral prefrontal cortex of the macaque. Nature Neurosci. 2, 176–185 (1999).

Liu T, Pleskac TJ (2011) Neural correlates of evidence accumulation in a perceptual decision task. J Neurophysiol 106:2383–2398.

Luber, B., & Lisanby, S. H. (2014). Enhancement of human cognitive performance using transcranial magnetic stimulation (TMS). *Neuroimage*, *85*, 961-970.

Matzke, D., & Wagenmakers, E. J. (2009). Psychological interpretation of the ex-Gaussian and shifted Wald parameters: A diffusion model analysis. Psychonomic bulletin & review, 16(5), 798-817.

Mulder, M. J., Van Maanen, L., & Forstmann, B. U. (2014). Perceptual decision neurosciences–a model-based review. Neuroscience, 277, 872-884.

Philiastides MG, Sajda P (2007) EEG-informed fMRI reveals spatiotemporal characteristics of perceptual decision making. J Neurosci 27:13082–13091.

Philiastides, M. G., Auksztulewicz, R., Heekeren, H. R., & Blankenburg, F. (2011). Causal role of dorsolateral prefrontal cortex in human perceptual decision making. *Current Biology*, *21*(11), 980-983.

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.

Rahnev, D. (2017). Top-Down Control of Perceptual Decision Making by the Prefrontal Cortex. *Current Directions in Psychological Science*, 096372141770980. https://doi.org/10.1177/0963721417709807

Rahnev, D., Nee, D. E., Riddle, J., Larson, A. S., & D'Esposito, M. (2016). Causal evidence for frontal cortex organization for perceptual decision making. *Proceedings of the National Academy of Sciences*, *113*(21), 6059-6064.

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

Ratcliff, R., & Childers, R. (2015). Individual differences and fitting methods for the two-choice diffusion model of decision making. *Decision*, *2*(4), 237.

Ratcliff, R. & McKoon, G. The diffusion decision model: theory and data for two-choice decision tasks. Neural Comput. 20, 873–922 (2007).

Ratcliff, R., Smith, P. L., Brown, S. D., & McKoon, G. (2016). Diffusion decision model: current issues and history. *Trends in cognitive sciences*, *20*(4), 260-281.

Rolls ET, Grabenhorst F, Deco G (2010) Decision-making, errors, and confidence in the brain. J Neurophysiol 104:2359–2374.

The jamovi project (2019). *jamovi*. (Version 1.0) [Computer Software]. Retrieved from https://www.jamovi.org.

Tosun, T., Berkay, D., Sack, A. T., Çakmak, Y. Ö., & Balcı, F. (2017). Inhibition of pre–supplementary motor area by continuous theta burst stimulation leads to more cautious decision-making and more efficient sensory evidence integration. *Journal of cognitive neuroscience*, *29*(8), 1433-1444.

Vallesi A, McIntosh AR, Crescentini C, Stuss DT. fMRI investigation of speed-accuracy strategy switching. Hum Brain Mapp 2012 33:1677-1688

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*(May 2008), 1952–1965. https://doi.org/10.1162/jocn.2008.20146

Weigard, A., Beltz, A., Reddy, S. N., & Wilson, S. J. (2019). Characterizing the role of the pre-SMA in the control of speed/accuracy trade-off with directed functional connectivity mapping and multiple solution reduction. *Human Brain Mapping*, *40*(6), 1829–1843. https://doi.org/10.1002/hbm.24493

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.

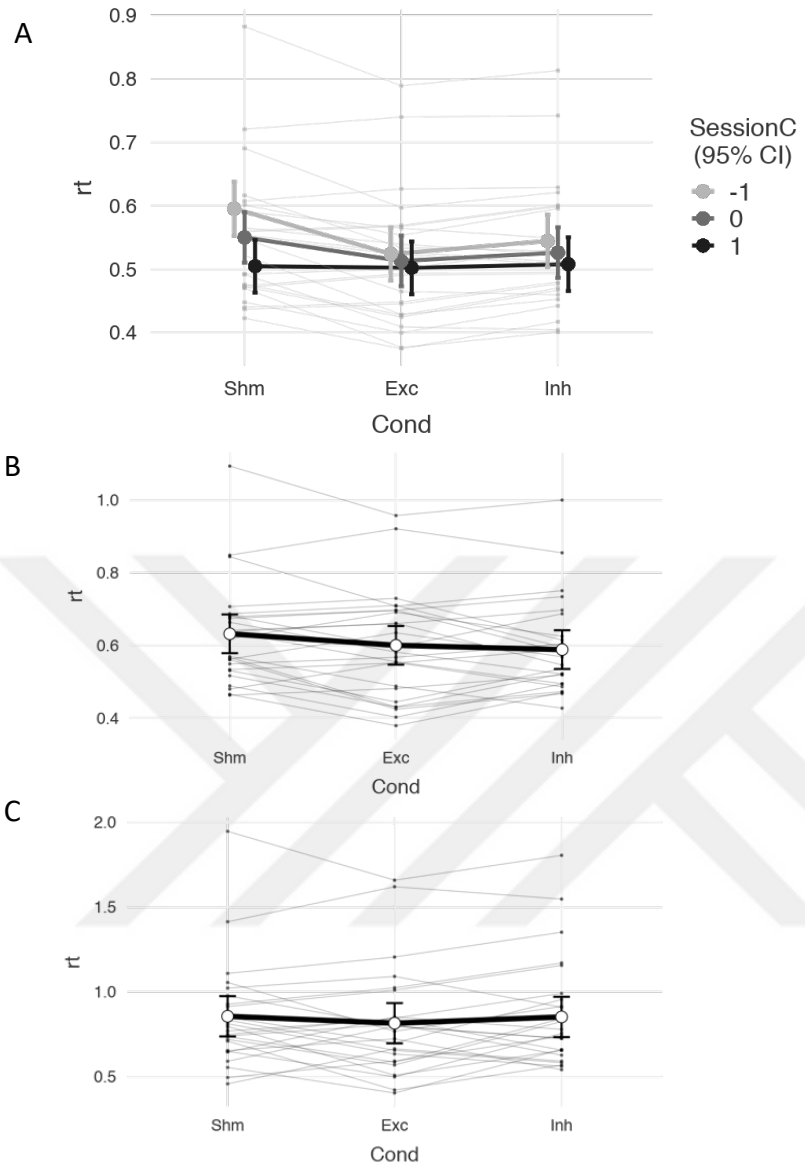

Figure 1. Response time (s) estimates from the linear mixed effects models of the three-session analyses. (A) is the session*stimulation interaction in the easy level, where the lightest gray is the first stimulation session and the darkest gray is the final session. (B) is the main effect of stimulation condition in the moderate level, and (C) is the main effect of stimulation condition in the hard level. Bars stand for 95% confidence intervals of the estimated marginal means (not the effects), thin lines stand for random effects (participants).

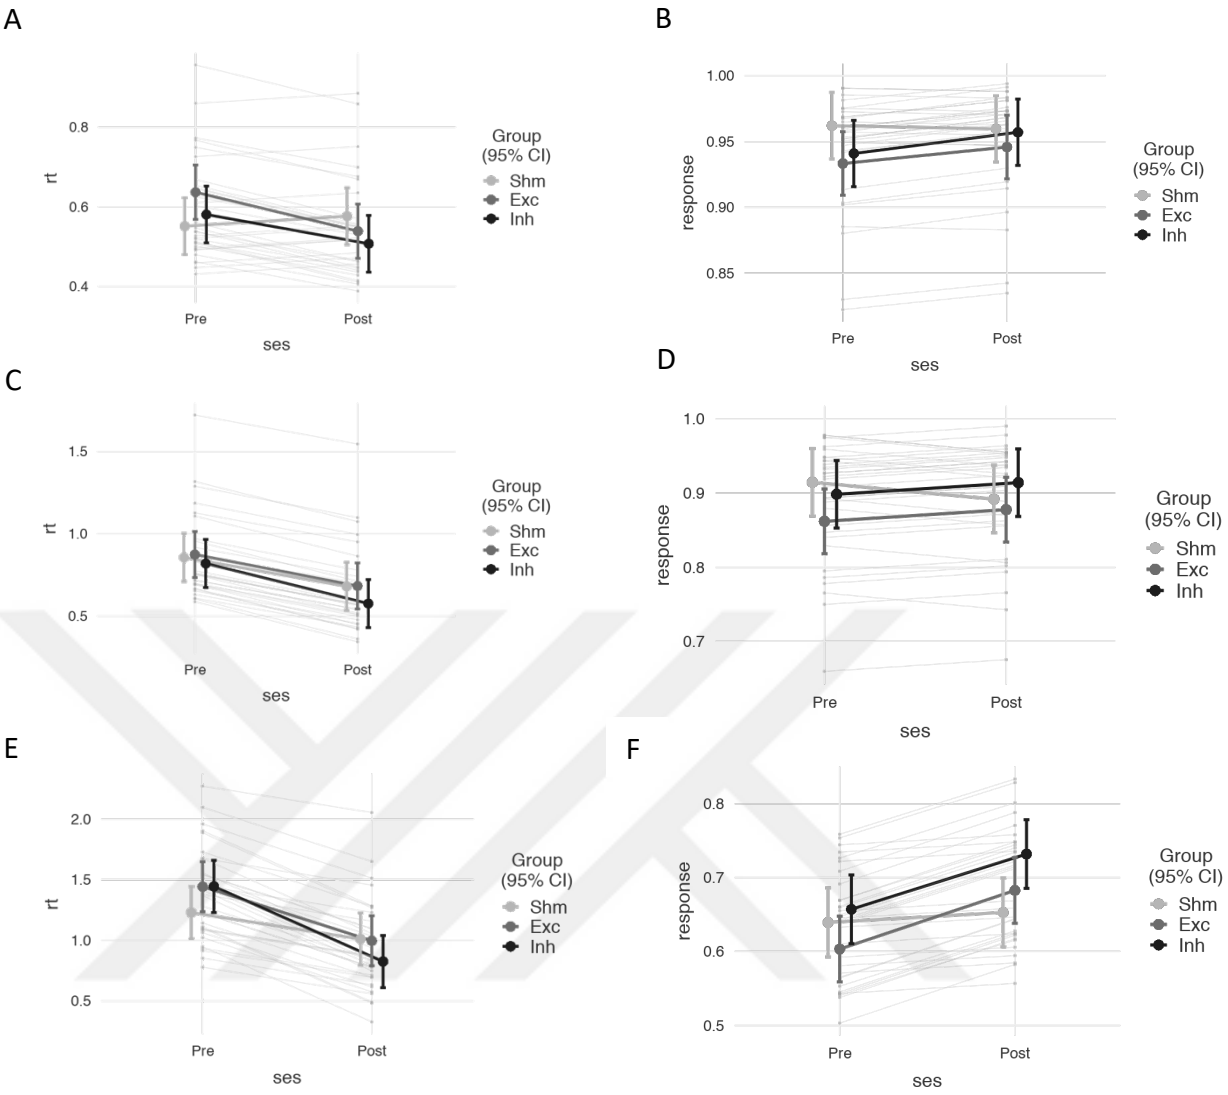

Figure 2. Response time (s) (left panel) and accuracy (right panel) estimates from linear mixed effects models of the single-session analyses. Top row (A and B) are for the easy level (note that the group*session interaction was not significant for this level, but the interaction plot is presented for consistency with the other levels). Middle row (C and D) are for the moderate level, and bottom row (E and F) are for the hard level. Light gray lines are for Sham group estimates, moderate gray lines are for Excitation group estimates and dark gray lines are for Inhibition group estimates. Bars stand for 95% confidence intervals of the estimated marginal means (not the effects), thin lines depict random effects (participants).

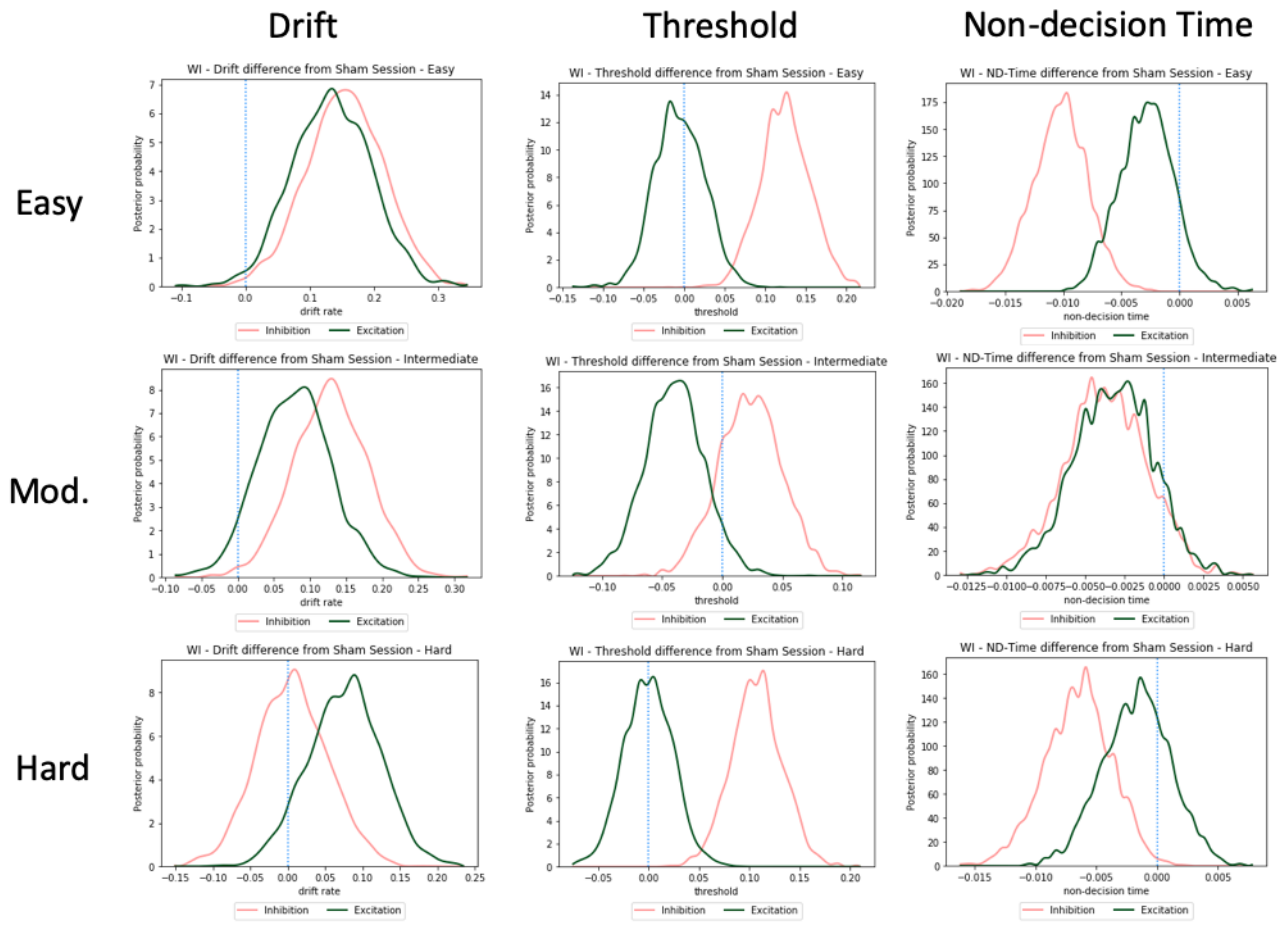

Figure 3. Posterior distributions of drift rate, threshold and non-decision times after inhibition and excitation, compared to the sham session. Three-session comparisons for participants who completed all four sessions. In the grid, rows are for difficulty levels, and columns are for model parameters. The dashed lines at zero stand for the sham session references (intercepts). The pink (light gray) distributions are the difference of the inhibition session from the sham session, gray (dark gray) distributions are the difference of excitation sessions from sham.

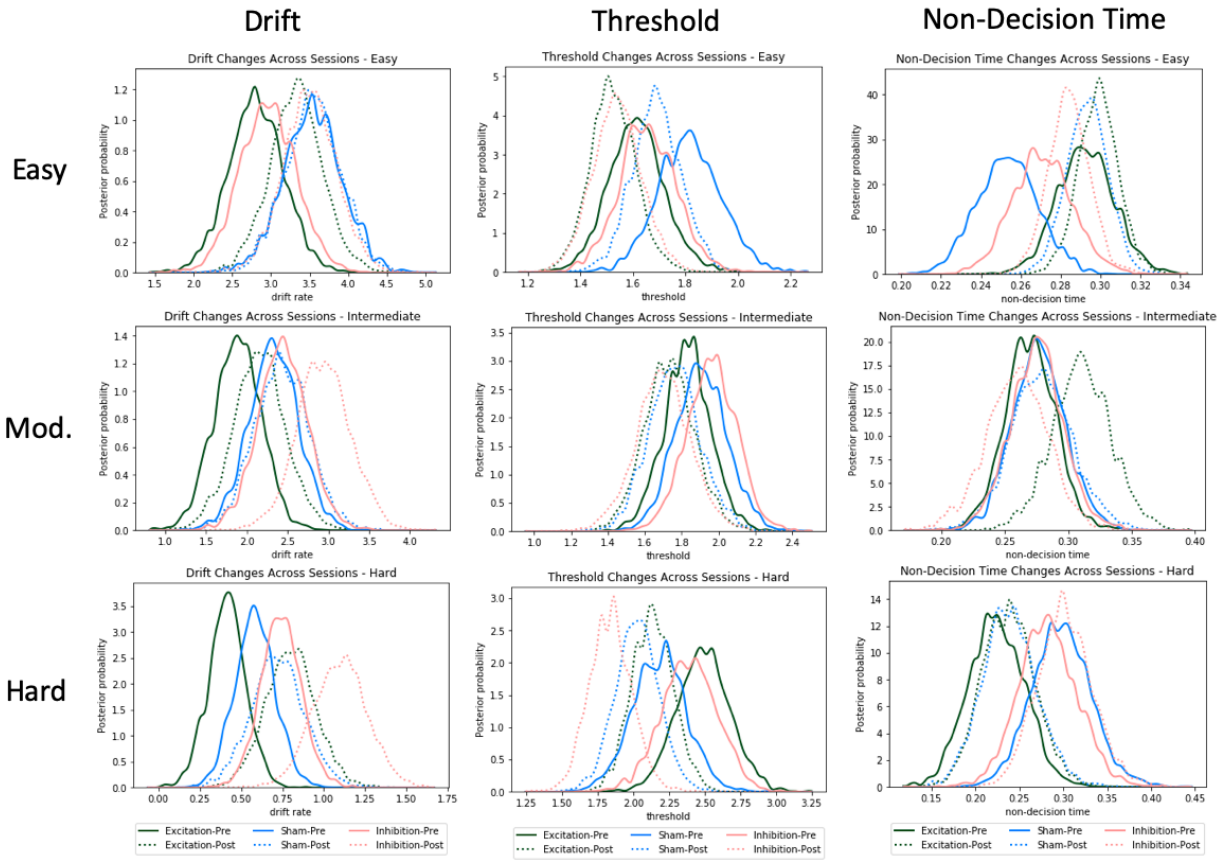

Figure 4. Posterior distributions of drift rate, threshold and non-decision times after in baseline and first stimulation sessions, for inhibition, excitation, and sham groups. Single-session comparisons for all participants who completed the first two sessions. In the grid, rows are for difficulty levels, and columns are for model parameters. Solid lines are for the baseline sessions and dotted lines are for the stimulation sessions. Green (dark gray) distributions are for the excitation group, blue (mod. gray) distributions are for the sham group, and pink (light gray) distributions are for the inhibition group. The posteriors represent the actual parameter estimation values and not difference scores (no intercepts were used).