Neural Mechanisms of Action-Selective and Stimulus-Selective Stopping

Ruben van den Bosch[1,2]
Supervisors: Bram Zandbelt[1,2] , Roshan Cools[1,2]

[1]Radboud University Nijmegen, Donders Institute for Brain, Cognition and Behaviour, The Netherlands
[2]Radboud University Medical Centre Nijmegen, Donders Institute for Brain, Cognition and Behaviour, The Netherlands

The past decade has seen a surge of interest in selective stopping. Researchers studying selective stopping have relied on the independent race model of simple stopping. Further- more, they have investigated selective stopping with a heterogeneous set of tasks, including action-selective and stimulus-selective stop tasks. Action-selective stop tasks probe control of specific actions and stimulus-selective stop tasks examine control triggered by specific stimuli. However, it remains unclear whether the independent race model extends to selective stopping and whether selective stopping is a homogeneous or heterogeneous construct. Here, we addressed these important gaps. We tested whether selective stopping performance is in agreement with predictions of the independent race model, using a Bayesian hypothesis testing approach based on the Bayes factor. We performed these tests at the group- and individual-level. We then compared action- and stimulus-selective stopping in terms of performance and brain activation, using functional magnetic resonance imaging.

We found violations of the predictions of the independent race model in 91% of the individuals in action-selective stopping and 74% of the individuals in stimulus-selective stop- ping. These individual violations were almost completely masked by the group performance. Furthermore, performance did not differ between the two selective stopping types and there appeared to be no differences in inhibition-related brain activity.

These results suggest that the independent race model does not generally extend to selective stopping and that action-selective and stimulus-selective stopping form a homogeneous construct.

Keywords: cognitive control, race model, response inhibition, selective stopping, functional MRI

Corresponding author: Ruben van den Bosch; E-mail: r.vandenbosch@donders.ru.nl

The ability to inhibit intended actions in response to environmental changes is an essential act of control in daily life and central to adaptive human behaviour. A drastic form of cognitive control is complete inhibition.

For decades, this has been studied with the simple stop-signal paradigm (Verbruggen & Logan, 2008). In the simple stop-signal task subjects make quick responses to go-stimuli but try to cancel the response when an infrequent stop-signal occurs after a variable delay (stop-signal delay; t_d). The trials on which such a stop-signal occurs can be divided based on whether stopping succeeds (stop-inhibit trials) or fails (stop-respond trials). Stopping performance on this task is characterised by three main findings: (1) the probability of responding (P_r) given a stop-signal increases with t_d ; (2) stop-respond response time (RT) is shorter than no-signal RT; (3) stop-respond RT increases with t_d .

The independent race model provides a theoretical framework from which these findings can be understood (Logan & Cowan, 1984; Logan, Van Zandt, Verbruggen, & Wagenmakers, 2014). This model explains stopping performance as the outcome of a race between a GO process that executes the response and a STOP process that cancels it. If the GO process finishes before the STOP process, the response is executed; if the STOP process finishes before the GO process, the response is canceled. Under these assumptions the model predicts exactly the pattern of findings observed in the standard stop-signal task. Besides a theory of simple stopping, the independent race model provides methods to estimate the latency of the covert STOP process, known as the stop-signal reaction time (SSRT). The SSRT can be estimated from the proportion of stop-respond trials and the distribution of no-signal RTs.

The independent race model has stimulated extensive use of the stop-signal task in various fields of research. This has greatly furthered our understanding of, for example, the lifespan development of control (Coxon, Impe, Wenderoth, & Swinnen, 2012; Van de Laar, Van den Wildenberg, Van Boxtel, & Van der Molen, 2011), response inhibition itself (Logan, 1994; Verbruggen & Logan, 2008) and clinical and neurological disorders, including attention deficit hyperactivity disorder (ADHD; Dimoska, Johnstone, Barry & Clarke, 2003; Janssen, Heslenfeld, Van Mourik, Logan, & Oosterlaan, 2015), Schizophrenia (Thakkar, Schall, Boucher, Logan, & Park, 2011; Zandbelt, Buuren, Kahn, & Vink, 2011) and Parkinson's Disease (Gauggel, Rieger, & Feghoff, 2004; Van de Wildenberg et al., 2006).

Neuroscience studies have demonstrated that simple stopping manifests in the motor system and also involves areas in the frontal lobe and basal ganglia. Neurophysiology studies in animals have shown that eye movement-related activity of neurons in the frontal eye field (Hanes, Patterson, & Schall, 1998) and superior colliculus (Paré & Hanes, 2003) as well as limb movement-related activity of neurons in the dorsal premotor cortex (Mirabella, Pani, & Ferraina, 2011) and basal ganglia nuclei (Schmidt, Leventhal, Mallet, Chen, & Berke, 2013) decays in response to the stop-signal within the time required to cancel the movement (Schall & Boucher, 2007). Transcranial magnetic stimulation studies in humans have shown similar dynamics for primary motor cortex excitability (Coxon, Stinear, & Byblow, 2006; Van de Wildenberg et al., 2009). Human imaging, stimulation, and lesion studies suggest that simple stopping relies on the inferior frontal cortex (Aron, Fletcher, Bullmore, Sahakian, & Robbins, 2003; Aron & Poldrack, 2006; Chambers et al., 2006; Verbruggen, Aron, Stevens, & Chambers, 2010), pre-supplementary motor area (Chen, Muggleton, Tzeng, Hung, & Juan, 2009; Li, Huang, Constable, & Sinha, 2006), and basal ganglia structures such as the striatum (Zandbelt, Bloemendaal, Hoogendam, Kahn, & Vink, 2012; Zandbelt & Vink, 2010) and subthalamic nucleus (Aron & Poldrack, 2006; Jahfari et al., 2011; Van de Wildenberg et al., 2006).

Notwithstanding the deep insights this research has yielded, simple stopping is limited as a model of cognitive control in daily life and psychiatric disorders. Simple stopping evokes control that is non-selective; after the stop-signal subjects have to stop all their planned actions. In reality, most situations require selective stopping, a more flexible form of control. It comprises control that is targeted at specific actions or triggered by specific stimuli and has been studied with selective stopping tasks. Selective stopping research has been suggested to not only have greater ecological validity, but also greater clinical relevance (Aron, 2011). However, two main factors currently complicate the interpretation of data from selective stopping research.

Firstly, although selective stopping research has relied on the independent race model, it is uncertain whether the model extends from simple stopping to selective stopping. One reason is that studies to date have reported tests of the predictions of the independent race model incompletely or not at all. Another reason is that the available data on tests of the predictions provide mixed evidence. For example, although at least one selective stopping

study showed that all three predictions of the race model held (Smittenaar, Guitart-Masip, Lutti, & Dolan, 2013), many others reported violations of at least one of the predictions (Bissett & Logan, 2014; De Jong, Coles, & Logan, 1995; Dimoska, Johnstone, & Barry, 2006; Van de Laar, Van den Wildenberg, Van Boxtel, & Van der Molen, 2010; Verbruggen & Logan, 2015) in as many as 61% of participants (Bissett & Logan, 2014). A final reason is that tests of the race model's predictions are often performed for the group as a whole rather than for each individual separately, which may mask violations occurring in a subset of individuals. To illustrate, one selective stopping study reported that one of the predictions (stop-respond RT should be shorter than no-signal RT) held for the group as a whole, but also reported that the very same prediction was violated in 32% of the participants (Sebastian et al., 2015).

Together, this state of affairs is unsatisfactory, because if it turns out that the independent race model does not extend generally from simple stopping to selective stopping, then SSRT estimates reported by previous selective stopping studies may be invalid and conclusions derived from them may be flawed. To address this problem, a systematic investigation of predictions of the independent race model across different forms of selective stopping and performed at the individual level is necessary. This will help clarify how often violations of predictions of the independent race model occur.

Secondly, it is unclear whether selective stopping is a homogeneous or a heterogeneous construct. As pointed out by Bissett and Logan (2014), selective stopping research has used a heterogeneous set of tasks, yet all of them are called selective stopping, as if selective stopping is a homogeneous construct. In some tasks (e.g., Aron & Verbruggen, 2008; Coxon et al., 2016; Coxon, Stinear, & Byblow, 2009; Majid, Cai, Corey- Bloom, & Aron, 2013; Smittenaar et al., 2013) participants are instructed to stop certain actions (e.g., a left-hand response), while continuing others (e.g., a right-hand response). We call this action-selective (AS) stopping (note that Bissett and Logan (2014) have called this unconditional motor-selective stopping). In other tasks (e.g., Bissett & Logan, 2014; Dimoska et al., 2006; Van de Laar et al., 2010; Ruiter, Oosterlaan, Veltman, Van den Brink, & Goudriaan, 2012; Sebastian et al., 2015; Sharp et al., 2010; Verbruggen & Logan, 2015) participants are instructed to stop to certain signals, while ignoring others. We call this stimulus-selective (SS) stopping. It remains unclear whether AS and SS stopping tap into the same form of control, as these tasks have never been compared directly.

On the one hand, several findings seem to suggest that they do involve the same form of control, including SSRTs that are in the same range, response strategies that show striking resemblances (Bissett & Logan, 2014; Macdonald, Stinear, & Byblow, 2012), and activation in seemingly similar brain regions, such as ventrolateral frontal cortex, dorsal frontal cortex, and basal ganglia (Coxon et al., 2016, 2009; Majid et al., 2013; Ruiter et al., 2012; Sebastian et al., 2015; Sharp et al., 2010; Smittenaar et al., 2013). On the other hand, violations of the independent race model have mainly been reported for SS stopping (Bissett & Logan, 2014; Dimoska et al., 2006; Van de Laar et al., 2010; Verbruggen & Logan, 2015) and rarely for action- selective stopping (De Jong et al., 1995). Moreover, AS stopping may involve at least two cognitive steps before the inhibition of a response is initiated (discriminating the signal and selecting the action to be canceled) and SS stopping may involve only one (discriminating the signal). Consequently, AS stopping may rely more heavily on motor-related brain regions, such as the pre-supplementary motor area (pre-SMA), the supplementary motor area (SMA), the dorsal premotor area (PMd), and possibly the basal ganglia. To address this problem, a direct comparison of AS and SS stopping tasks in terms of behavioural and neural measures of stopping is necessary.

In the present functional magnetic resonance imaging (fMRI) experiment, we used a task with AS and SS stopping intermixed, allowing us to compare these forms of selective control both behaviourally and neurobiologically. We tackle the problems described above by addressing two research questions:

1. Does the independent race model extend to selective stopping?
2. Is selective stopping a homogeneous or heterogeneous construct?

If the independent race model does not extend to selective stopping, we would expect to find that stopping performance would violate any of the three qualitative predictions of the race model. Alternatively, if the independent race model does extend to selective stopping, we would expect to find that stopping performance is in line with all three of the model's qualitative predictions.

If selective stopping is a homogeneous construct, then we would expect that AS and SS stopping would not differ in terms of stopping performance and brain activation measures of selective stopping. Alternatively, if selective stopping

is a heterogeneous construct, then we would expect that AS and SS stopping would differ in terms of stopping performance and brain activation measures of selective stopping.

Methods

Pre-registration

This project has been pre-registered at the Open Science Framework (www.osf.io) on May 30, 2016. In the pre-registration, all the methods, procedures, outcome measures and confirmatory analyses are described in detail. Additional, not pre-registered, analyses were exploratory, rather than confirmatory, and are indicated as such in the text below. The document is available on request and it will be made publicly accessible upon publication of this research. At the time of pre-registration, four datasets had already been collected, but not analysed.

There were four deviations from the pre-registration. Firstly, 24 subjects were included for this thesis, rather than the pre-registered 30. Before publication of this project as a research article, additional subjects will be scanned until thirty subjects are included in the analyses. Secondly, we used a Bayesian repeated-measures ANOVA instead of Bayesian logistic regression in the analysis of the effect of t_d on P_r (see 'Tests of independent race model predictions' section). Thirdly, in the analysis of the effect of t_d on stop-respond RT we added a restriction to the Bayesian repeated-measures ANOVA model that was not in the pre-registration (see 'Tests of independent race model predictions' section). Fourthly, in the fMRI analysis, only six motion regressors were used instead of the pre-registered twenty-four see 'Functional MRI analyses' section).

Participants

Twenty-five healthy participants volunteered for this study. One participant was excluded after the practice session (see 'Practice session' section), bringing the number of participants to twenty-four (mean age 23.8 years, range 18-33; 17 females). All participants had normal or corrected to normal vision and did report no history of neurological or psychiatric illness or claustrophobia. Written informed consent was obtained from all participants. The study procedures were in accordance with the Declaration of Helsinki and have been approved by the local Institutional Review Board (Committee on Research Involving Human Subjects Arnhem-Nijmegen, registered under CMO2014/088).

Task

We used a mixed AS and SS stop task (Fig. 1). All stimuli were presented in the centre of the screen on a grey background. The fixation stimulus was a white cross, subtending 3˚ along its vertical axis. The primary (go) stimulus was a white Hiragana character, subtending 6˚ along its vertical axis. On each trial it was chosen from a set of two. The secondary stimulus was a playing card suit symbol subtending 3˚ along the vertical axis and presented on top of the primary stimulus at 80% opacity. It was chosen from a set of four: an orange diamond, a cyan heart, a yellow spade, or a purple club.

Each trial started with the presentation of the fixation stimulus for 200 ms. The fixation cross was immediately followed by the primary go stimulus, which remained on the screen for 1200 ms, regardless of response time. Following go stimulus offset, feedback was presented for 500 ms. The next trial started after a further 200 ms, during which a blank screen was shown.

The primary task was to respond to the identity of the Hiragana character. Both characters required a bimanual response. This kept the primary task the same throughout the experiment in order to keep AS and SS stop trials as similar as possible. One character was mapped onto the two upper keys of a response pad; the other character was mapped onto the two lower keys. The character-to-key mapping was counterbalanced across participants. Participants pressed the two upper keys with their left and right middle fingers and the lower keys with their left and right index fingers. Participants were instructed to respond as accurate, fast, and simultaneously as possible.

On 40% of the trials, one of the four secondary stimuli (signals) was presented. The other 60% percent were no-signal (NS) trials. Three of the four signals indicated that an adjustment of the response to the go-stimulus was required. There were two versions of the task, counterbalancing the stimulus-to-signal mapping across participants. One stimulus (orange diamond/purple club) acted as the stop-left (SL) signal; it instructed participants to stop their left-hand response, but not their right-hand response. A second stimulus acted as the stop-right (SR) signal (cyan heart/yellow spade); it instructed participants to stop their right-hand response, but not their left-hand response. A third stimulus acted as the stop-both (SB) signal (yellow spade/cyan

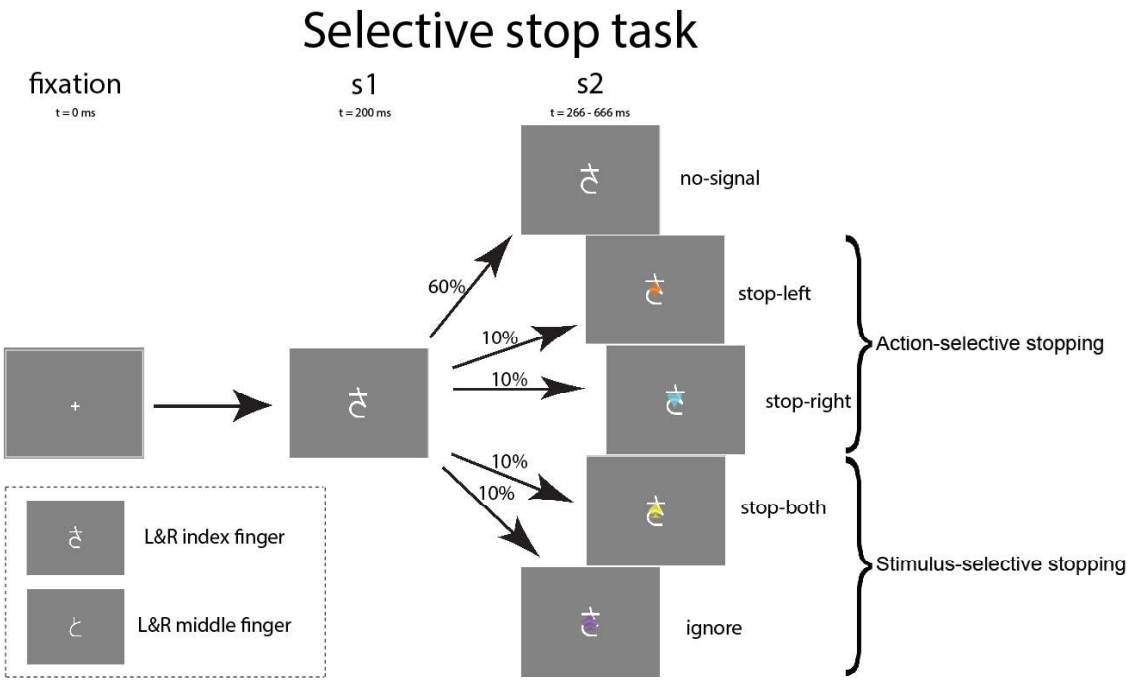

Fig. 1. Schematic of the selective stopping task. Each trial began with 200 ms fixation, followed by one of two go stimuli (s1). The main task was to quickly respond by pressing a button with either both index fingers or both middle fingers, depending on the identity of the go stimulus. On 40% of the trials the go stimulus was followed by one of four secondary stimuli (s2; a colored playing card suit symbol) after a delay of 66, 166, 266, 366 or 466 ms. Two served as action-selective stop-signals, indicating to stop the left- hand response but continue the right-hand response, or vice versa. The other two served as stimulus-selective stop-signals; one indicated to stop both responses, the other signal was to be ignored and the response should be carried out as normal.

heart); it instructed participants to stop both their left- and right-hand response. The fourth stimulus acted as the ignore (IG) signal (purple club/orange diamond); it had to be ignored and the response had to be continued as planned. The primary and secondary stimulus were separated by a stimulus onset asynchrony (i.e., the stop-signal delay, t_d). There were five t_d , each occurring with equal probability: 66 ms, 166 ms, 266 ms, 366 ms, and 466 ms. These values were based on pilot data.

Procedure

Practice session. The experiment began with both written and verbal instructions. Next, the task was practiced in three stages. In stage 1, participants performed one block of 50 no-signal trials to acquaint them with the go task. In stage 2, participants performed one AS stopping block (only SL or SR signals could appear) and one SS stopping block (only SB or IG signals could appear), to acquaint them with both the AS and SS stop task, while maintaining go task performance. The blocks consisted of 100 trials each and the order of the blocks was counterbalanced across participants. In stage 3, participants performed one block of 100 trials in which AS and SS stop trials were intermixed, to acquaint them with the task as it would be performed in the scanner. After each practice block participants were provided with feedback on their performance. The practice block had to be repeated if one of the criteria listed in Table 1 was not met. Each block could be repeated up to five times. If any of these criteria was not met after five subsequent repetitions, the experiment was terminated and the participant was excluded. One participant was excluded for this reason. Upon successful completion of the practice session, the fMRI session was scheduled for another day.

Functional MRI session. Participants performed two runs of 6 experimental blocks while being scanned with fMRI, using an event-related design. Each block consisted of 100 trials and ended with a 12.6-second feedback screen on the task performance (identical to the practice session). If one of the performance criteria listed in Table 1 was not met on five subsequent blocks, the participant

Table 1.

Performance criteria. NS = no-signal; SL = stop-left; SR = stop-right; SB = stop-both; IG = ignore

Outcome measure	Performance criteria
Mean NS RT	< 650 ms
NS choice performance (i.e., correct response to the go stimulus	>= 85% correct
Mean difference between left- and right hand RT on NS trials	< 50 ms
SL trials	$20\% < P(respond \mid\mid SLsignal) < 80\%$
SR trials	$20\% < P(respond \mid\mid SRsignal) < 80\%$
SB trials	$20\% < P(respond \mid\mid SBsignal) < 80\%$
IG trials	$P(respond \mid\mid IGsignal) >= 80\%$

would be excluded. One participant was excluded for this reason.

The trials were presented in a pre-determined pseudo-random sequence. In order to determine the optimal trial order, we generated 100,000 pseudo-random trial sequences and selected the two sequences (one for each fMRI run) with the highest detection efficiency for the contrast between AS and SS stopping and the lowest variance inflation factor, as determined with MATLAB-software for optimisation of fMRI designs (Wager & Nichols, 2003).

Data acquisition

Task performance data. The experiment was run in PsychoPy (version 1.83.04) in Windows 7 Enterprise OS, on a DELL PRECISION T3500 computer. Visual stimuli were projected on a screen positioned 75 cm from the subject and were viewed through a mirror mounted on the head coil. Responses were collected using two MR-compatible response pads (Current Designs, Inc; Philadelphia, PA, USA), one for each hand.

Neuroimaging data. The experiment was performed on a 3.0 T Siemens Magnetom Skyra MRI scanner (Siemens Medical Systems, Erlangen, Germany) at the Donders Institute. Images were acquired using a 32-channel head coil. During task performance, a total of 1214 images with blood-oxygen level-dependent (BOLD) contrast were acquired in 2 runs, using a whole-brain T2*-weighted gradient echo multi-echo echo planar imaging (EPI) sequence (34 slices per volume; transversal acquisition; repetition time = 2100 ms; echo times 8.5 ms, 19.3 ms, 30 ms, and 41 ms; field of view = 224 x 244 mm; flip angle = 90˚; 64 x 64 matrix; 3.5 mm in-plane resolution; 3 mm slice thickness; 0.5 mm slice gap). The first 6 scans of each run were discarded to allow T1 saturation to reach equilibrium. Before the first task run, 30 images were acquired during resting-state, with the same pulse sequence, for determining optimal weighting of echo times for each voxel. Between the two task runs, a whole-brain structural image was acquired for within-subject registration purposes, using a T1-weighted magnetisation prepared, rapid-acquisition gradient echo sequence (192 sagittal slices; repetition time = 2300 ms; echo time = 3.03 ms; field of view = 256 x 256 mm; flip angle = 8˚; 256 x 256 matrix; 1.0 mm in-plane resolution; 1.0 mm slice thickness).

Data analysis

Bayesian hypothesis testing. Behavioural data were analysed with a Bayesian hypothesis testing approach, based on the Bayes factor (Kass & Raftery, 1995). Bayesian hypothesis testing is comparative in nature and the Bayes factor quantifies the support that the data provide for one hypothesis (e.g., the null hypothesis, H_0) over another (e.g., the alternative hypothesis, H_1). This approach has several advantages over classical hypothesis testing based on the *p*-value. Most importantly, it allows for obtaining evidence both in favor and against H_0, rather than against H_0 only. This is particularly relevant in this study, because we investigate whether the independent race model does (H_1) or does not (H_0) extend to selective stopping and

whether AS and SS stopping is (H_0) or is not (H_1) a homogeneous construct. In addition, the support for one hypothesis over another is provided as a continuous measure (the Bayes factor) instead of a forced, all-or-none, decision.

The Bayes factor describes the relative probability of the data under competing hypotheses. In Bayesian hypothesis testing, the relative odds of the hypotheses themselves are evaluated:

$$\underbrace{\frac{P(H_0|Data)}{P(H_1|Data)}}_{posterior\ odds} = \underbrace{\frac{P(Data|H_0)}{P(Data|H_1)}}_{Bayes\ Factor\ (B_{01})} \times \underbrace{\frac{P(H_0)}{P(H_1)}}_{prior\ odds}$$

Here, the prior and posterior odds describe the beliefs about the hypotheses before and after observing the data, respectively. The primary measure of interest, however, is the Bayes factor that quantifies the change in odds from prior to posterior. In other words, the Bayes factor describes how the evidence from the data should change our beliefs. The Bayes factor prefers the hypothesis under which the observed data are most likely. To illustrate, if $B_{01} = 5$, the data are five times as likely to have occurred under H_0 than under H_1; if $B_{01} = 1$, the data provide equal support for H_0 and H_1. While the Bayes factor is easy to understand, it can be useful to summarise its value in words. For this purpose, we used a set of labels listed in Table 2, proposed by Wetzels and Wagenmakers (2012).

Table 2.

Bayes factor (B_{01}) categories, as proposed by Wetzels and Wagenmakers (2012).

B_{01}	Interpretation
> 100	Extreme evidence for H_0
30 – 100	Very strong evidence for H_0
10 – 30	Strong evidence for H_0
3 – 10	Moderate evidence for H_0
1 – 3	Anecdotal evidence for H_0
1	No evidence
1/3 – 1	Anecdotal evidence for H_1
1/10 – 1/3	Moderate evidence for H_1
1/30 – 1/10	Strong evidence for H_1
1/100 – 1/10	Very strong evidence for H_1
<1/100	Extreme evidence for H_1

Bayesian hypothesis testing requires specification of priors, which describe the distribu- tion of effect size. Prior distributions should be specified for both the null hypothesis and the alternative hypothesis. Here, we attempted to specify prior distributions that convey little information while maintaining desirable properties by placing priors on standardised effect sizes (δ) for H_0 and H_1, as well as on the variance (σ^2; Rouder, Speckman, Sun, Morey, & Iverson, 2009). The standardised effect size is assumed to be 0 under H_0 and distributed according to a Cauchy distribution with scale parameter $\tau = 0.5$ under H_1. The prior for σ^2 is less important, because it is the same for both hypotheses and cancels out in the Bayes factor. Following Rouder, Morey, Speckman, and Province (2012), we assume that σ^2 follows an inverse chi-square distribution with one degree of freedom.

In this study, we used Bayesian t-tests (Rouder et al., 2009) and Bayesian repeated- measures ANOVAs (Rouder et al., 2012). These tests involve running separate test models, one including the independent variable as a factor and one null model in which the only factor is the between-subject variance. Model comparison determines which model is the most likely, given the data. This determines whether or not the independent variable has an effect on the dependent variable.

For the Bayesian analyses, we used the Bayes factor package in the software R (https://cran.r-project.org/web/packages/BayesFactor/index.html) and JASP (https://jasp- stats.org/).

Behavioural analyses. The primary outcome measures in the behavioural analyses were P_r, no-signal RT and stop- respond RT for the different stop-signals. Only stop-respond trials with a bimanual response were included, because stop-respond trials with a unimanual response do not necessarily reflect failures of stopping: they may occur when participants successfully stop both prepared responses, but then erroneously discriminate the signal (e.g., confusing stop-left for stop- right) and execute the wrong response. In addition, trials with RTs faster than 150 ms were considered anticipated and were excluded from the analyses.

Tests of independent race model predictions. To test whether the independent race model extends to selective stopping, we tested whether AS and SS stopping performance was in line with the three predictions of the model. We tested these predictions for AS and SS stopping separately, both at the individual level and group level.

To test the first prediction (P_r increases with t_d) at the individual level, we plotted the inhibition functions of each subject. At the group level, we analysed P_r with a Bayesian repeated-measures ANOVA, with t_d as a factor. For the second prediction (stop-respond RT is faster than no-signal RT), we analysed the RTs at the individual level with Bayesian independent t-tests, with trial outcome (no-signal or stop-respond trial) as a factor. At the group level, we analysed the RTs with two Bayesian paired t-tests, one for AS and one for SS stopping. We tested the third prediction (stop-respond RT increases with t_d) by analysing stop-respond RT with a Bayesian repeated-measures ANOVA, with t_d as a factor. For this analysis (both at the individual level and the group level), the factor t_d had three levels rather than five: short delay (66, 166, 266 ms), intermediate delay (366 ms) and long delay (466 ms). Pooling of stop-respond RTs over the first three delays was necessary, because of the uneven distribution of stop-respond trials; there are more stop-respond trials at longer t_d.

There were two deviations from the pre-registration in these analyses. Firstly, we did not use Bayesian logistic regression for the analyses of the effect of t_d on P_r, because Bayesian logistic regression had not yet been implemented in the Bayes factor package for R. Secondly, we added a restriction to the models in the Bayesian repeated-measures ANOVAs. We described in the pre-registration that the data support an effect of the independent variable (t_d here; selective stopping type in next section's analyses) if the winning ANOVA model contains the independent variable as a factor. However, this model only supports a main effect, not necessarily in the right direction. In order to find evidence for or against an increase in P_r or stop-respond RT with increasing td, we created an order-restricted model. The order-restricted model took 10,000 samples from the posterior distribution of the full model (the not-order-restricted model) and computed the frequency of the correct ordering (higher P_r and longer RTs at longer t_d). We also report the results of the full model.

Behavioural tests of differences between selective stopping types. There were three behavioural analyses (group level) to test whether selective stopping is a behaviourally homogeneous or heterogeneous construct. Firstly, we analysed the P_r with a Bayesian repeated-measures ANOVA, with selective stopping type (AS or SS) as a factor. Secondly, we analysed the stop-respond RT with a Bayesian repeated-measures ANOVA, with selective stopping type as a factor. Thirdly, we analysed the two SSRTs using a Bayesian paired t-test, with selective stopping type as a factor. For the third analysis, subjects whose performance was not in line with all the predictions of the independent race model for both AS and SS stopping were excluded.

Functional MRI analyses

Image data were preprocessed using Statistical Parametric Mapping 12 (SPM12) software running in MATLAB (Mathworks Inc., Natick, MA, USA). The T1-weighted anatomical scans were skull-stripped using the FSL Brain Extraction Tool (Smith, 2002). The multi- echo images were combined with the PAID method (Poser, Versluis, Hoogduin, & Norris, 2006) using custom-written MATLAB software. The anatomical images were co-registered to the mean functional images using the normalised mutual information criteria method (Studholme, Hill, & Hawkes, 1999). The anatomical and functional images were then normalised to the standard Montreal Neurological Institute 152 (MNI 152) template. The normalised functional images were spatially smoothed using a 6-mm full-width at half-maximum Gaussian kernel.

First-level statistical analysis involved a mass-univariate approach based on general lin- ear models in SPM12. Each subject's whole-brain BOLD data were modeled with a general linear model, including 15 event-related predictors. These were the five different trial types, all subdivided in three ways based on the trial outcome. NS and IG trials were divided into 'fast', 'slow' and 'other' response trials, wherein the 'other' category contained all the incorrect responses. SL, SR and SB trials were subdivided into 'stop-respond-bimanual', 'stop-inhibit' and 'stop-respond-other' response trials. For all regressors, except the three NS regressors, we included a demeaned parametric modulator coding for stimulus onset asynchrony be- tween the go-stimulus and the signal (i.e., t_d). In addition, to account for residual head motion effects, we included the six motion parameters from the realignment procedure in the statistical model. Note that we did not include the first and second order derivatives, as is described in the pre-registration. Taken together, we included a total of 33 regressors per run (i.e., 15 predictors + 12 parametric modulators + 6 motion parameters).

The regressors were created by convolving the delta functions coding for go-stimulus onset with a canonical hemodynamic response function. Time series statistical analysis was performed using restricted maximum likelihood. Low frequency

drifts were controlled using a discrete cosine transform with a cutoff of 128 seconds. Serial correlations in the fMRI signal were estimated using restricted maximum likelihood estimates of variance components using a first-order autoregressive model. The resulting non-sphericity was used to form maximum- likelihood estimates of the activations. Time series statistical analyses were performed using restricted maximum likelihood.

We specified first-level contrasts to isolate activation associated with AS stopping, SS stopping and the difference between them. The contrasts and their purposes are listed in Table 3. The contrasts control for the attentional capture of the stop-signal by subtracting activity on ignore trials from activity on stop trials. A salient signal also occurred on ignore trials, but no inhibition was required, so only attention-related but not inhibition- related activation is subtracted. For this subtraction, we used only ignore trials on which a fast response was made, because on slow ignore trials a STOP process may have been activated temporarily (Bissett & Logan, 2014). The contrasts also control for the difference in speed of the GO process between signal and no-signal trials by first subtracting activation on no-signal trials from both stop- and ignore-related activity. To control for activation associated with the execution of a unimanual response on AS stop trials, we used conjunction analyses (Nichols, Brett, Andersson, Wager, & Poline, 2005) in the second-level statistics to test for activations occurring in both the $S_{AS,left}$ and $S_{AS,right}$ contrasts and in both the $S_{AS,left} - S_{SS}$ and $S_{AS,right} - S_{SS}$ contrasts.

Second-level statistical analyses consisted of two region-of-interest (ROI) analyses and one whole-brain analysis on the first-level contrasts.

First, we analysed brain activation in predefined ROIs using a Bayesian hypothesis testing approach based on the Bayes factor. This allowed us to compare activation in the same way as we compared task performance and find evidence in favor of and against the null hypothesis. ROIs were defined as 6-mm spheres around local maxima in key regions of inhibitory control reported by previous fMRI studies of AS and SS stopping (Table 4): inferior frontal gyrus (IFG), inferior frontal junction (IFJ), striatum (Str), pre-supplementary motor area (pre-SMA), supplementary motor area (SMA), dorsal premotor area (PMd), and primary motor cortex (M1). For each region, we also defined a 6-mm sphere ROI in the other hemisphere, flipping the sign of the x-coordinate. From these ROIs we extracted the mean activation level (i.e. parameter estimate) in the first-level contrasts. We then used Bayesian one-sample t-tests (t-test value = 0) for identification of activation associated with AS stopping ($B_{01} < 1$ in both the $S_{AS,left}$ and $S_{AS,right}$ contrasts) and SS stopping (< 1 in the SSS contrast) and to identify potential differences in activation in the ROIs between AS and SS stopping ($B_{01} < 1$ in the $S_{AS,left} - S_{SS}$ and $S_{AS,right} - S_{SS}$ contrasts).

Second, we performed another ROI analysis, now using a classical hypothesis testing approach and more broadly defined ROIs, based on probabilistic anatomical atlases. The classical hypothesis testing approach allowed us to analyse the fMRI data in a more common framework. The more broadly defined ROIs enabled us to assess activation within key areas of inhibitory control that fell outside the spheres the first ROI analysis focused on. The broad, anatomical

Table 3.

Contrasts created at the first level and their purpose. The subtractions within parentheses control for differences in the speed of the GO process between signal and no-signal trials. The sub- tractions between parentheses control for the attentional capture of the stop-signal.

Contrast	Purpose
$S_{AS,left} = (SL_{stop\text{-}inhibit} - NS_{correct\text{-}slow}) - (IG_{correct\text{-}fast} - NS_{correct\text{-}fast})$	Isolate activation associated with AS stopping
$S_{AS,right} = (SR_{stop\text{-}inhibit} - NS_{correct\text{-}slow}) - (IG_{correct\text{-}fast} - NS_{correct\text{-}fast})$	Isolate activation associated with AS stopping
$S_{SS} = (SB_{stop\text{-}inhibit} - NS_{correct\text{-}slow}) - (IG_{correct\text{-}fast} - NS_{correct\text{-}fast})$	Isolate activation associated with SS stopping
$S_{AS,left} - S_{SS} = SL_{stop\text{-}inhibit} - SB_{stop\text{-}inhibit}$	Isolate differences between AS and SS stopping
$S_{AS,right} - S_{SS} = SR_{stop\text{-}inhibit} - SB_{stop\text{-}inhibit}$	Isolate differences between AS and SS stopping

ROIs are listed in Table 5. The probabilistic map of each of the subregions was thresholded at 25% and the resulting maps were combined into one binary mask. This mask was used for small-volume correction for multiple comparisons.

Within these broad anatomical ROIs we used a one-sample t-test in SPM12 on the SSS contrast for the identification of activation associated with SS stopping. We used conjunction analyses to identify activation associated with AS stopping ($S_{AS,left} \cap S_{AS,right}$) and to identify potential differences in activation in the ROIs between AS and SS stopping ($S_{AS,left} - S_{SS} \cap S_{AS,right} - S_{SS}$). We report activations that were significant at $p < .001$ uncorrected for multiple comparisons and that survived small-volume correction at $p < .05$ after family wise error (FWE) correction for multiple comparisons.

Third, to explore activation associated with AS and SS stopping outside key areas of inhibitory control, we performed whole-brain voxel-wise random effects analyses. Again, we used a conjunction analysis to isolate AS stopping-related activation ($S_{AS,left} \cap S_{AS,right}$), a one-sample t-test on the S_{SS} contrast for SS stopping-related activation and a conjunction analysis for differences between AS and SS stopping related activation ($S_{AS,left} - S_{SS} \cap S_{AS,right}$

Table 4.

ROIs defined on the basis of local maxima reported by previous fMRI studies of selective stopping

ROI	MNI [x,y,z] coordinates	Reference	Contrast in reference
IFG	52, 10, 6	Majid et al. (2013)	MaybeStopRight+Left (Stop > Go)
IFJ	45, 8, 25	Sebastian et al. (2015)	AttentionalCapture > Go
Str	9, 6, 0	Ruiter et al. (2012)	SuccessfulInhibition > Control
Pre-SMA	20, 6, 62	Sharp et al. (2010)	Stop > Continue
SMA	15, -2, 72	Coxon et al. (2016)	StopLeftGoRight > StopAll ∩ StopLeftGoRight > Go ∩ StopRightGoLeft > StopAll ∩ StopRightGoLeft > Go
PMd	28, -2, 65	Coxon et al. (2016)	StopLeftGoRight > StopAll ∩ StopLeftGoRight > Go ∩ StopRightGoLeft > StopAll ∩ StopRightGoLeft > Go
M1	-36, -24, 63	Ruiter et al. (2012)	SuccessfulInhibition > Control

Table 5.

ROIs based on probabilistic neuroanatomical atlases.

ROI	Subregions included	Reference
vlFC	ventral premotor, inferior frontal junction, 44v, 44d, 45, inferior frontal sulcus	F.-X. Neubert, Mars, Thomas, Sallet, & Rushworth (2014)
dFC	supplementary motor area, pre-supplementary motor area, and dorsal premotor area	J. Sallet et al. (2013)
Str	executive, rostral motor, caudal motor	Tziortzi et al. (2014)
M1	area 4a, area 4p	Geyer et al. (1996)

$- S_{SS}$). We report activations that were significant at $p < .001$ uncorrected for multiple comparisons and that survived cluster-level correction at $p < .05$ family wise error-corrected for multiple comparisons.

Results

Behaviour

We tested the three predictions of the independent race model at the group level and at the individual level, separately for AS and SS stopping. Subsequently, we tested to what extent stopping performance differed between AS and SS stopping. One participant failed to meet all pre-set performance criteria (Table 1), hence this dataset was excluded, resulting in a sample of 23 participants in the behavioural analyses.

Figure 2 summarises response times and response probabilities for the main trial types.

P_r increased with t_d. Figure 3 depicts individual and group mean inhibition functions (P_r over the five t_d) for AS and SS stopping. Clearly, P_r increased with t_d in all individuals. Indeed, the analyses at the group level confirmed that, both for AS and SS stopping, the data were much more likely under a model including t_d than a null model that did not include t_d as a factor (AS stopping, $B_{01} = 9.12e-54$; SS stopping, $B_{01} = 1.40e-55$). The data were also more likely under an order-restricted model, in which P_r increases with td, than the null model (AS stopping, $B_{01} = 8.06e-56$; SS stopping, $B_{01} = 1.19e-57$).

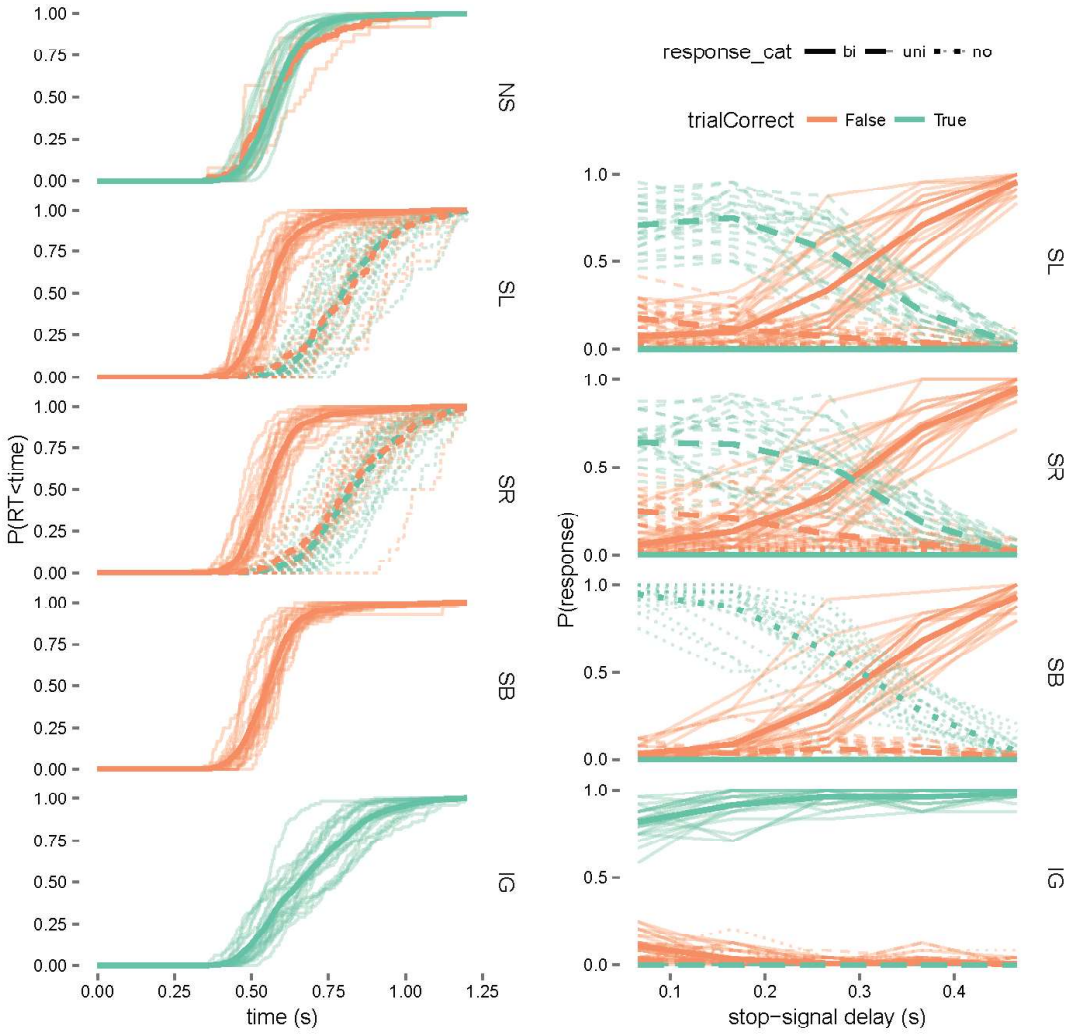

Fig. 2. Overall task performance. The left column of figures shows the cumulative distribution functions of RTs for the five trial types. The right column of figures shows the cumulative probabilities of a response over the signal delays (t_d). Faded lines represent individual subjects; bold lines represent the group mean. Solid lines represent bimanual responses, broken lines represent unimanual responses.

Difference between stop-respond RT and no-signal RT. Figure 4A and C display the relationship between mean stop-respond RT versus the mean no-signal RT. Figure 4B and D show the boxplots of inividual RTs for these trials. At the group level, the stop-respond RT was faster than no-signal RT (M = 583 ms) in both AS stopping (M = 566 ms, B_{01} = 0.001) and SS stopping (M = 562 ms, B_{01} = 4.74e − 05). At the individual level, there was more than anecdotal evidence (B_{01} < 1/3) that stop-respond RT was faster than no-signal RT for nine subjects in AS stopping and for eight subjects in SS stopping.

Stop-respond RT did not increase with t_d. The mean stop-respond RT of each subject in the three t_d bins (short, intermediate, long) is displayed in Figure 5A.

At the group level, the data were much more likely under the full model, which included t_d as a factor, than the null model that did not include t_d , in AS stopping (B_{01} = 0.001). The data for SS stopping were only slightly more likely under the full model than under the null model (B_{01} = 0.735). However, the order-restricted model showed that stop-respond RT did not increase with increasing td, neither in AS stopping (B_{01} = Inf, meaning that none of the 10,000 samples of the posterior distribution of the full model had the correct ordering) nor in SS stopping (B_{01} = 21.89).

At the individual level, there was more than anecdotal evidence for the full model being more likely (B_{01} < 1/3) in four subjects in AS stopping and four subjects in SS stopping. The order-restricted model of increasing stop-respond RT with increasing t_d was supported with more than anecdotal evidence in two of those subjects in AS stopping and in all four of the subjects in SS stopping.

The distribution of the log10 transformed Bayes factors (log10 [B_{01}]) for the full model and order-restricted model are shown for in AS and SS stopping in Figure 5B.

Race model predictions taken together. In total, two subjects performed in line with all three predictions of the independent race model in AS stopping and six subjects performed in line with all three predictions in SS stopping (B_{01} < 1; negative log10[B_{01}] in Fig. 4 and 5). Thus, at least one of the predictions of the independent race model was violated in 91% of the individuals in AS stopping and in 74% of the individuals in SS stopping. There were no subjects that performed in line with all three predictions in both AS and SS stopping.

Little difference in behaviour between selective stopping types. Figure 6 provides a clear visual comparison of the inhibition functions and stop-respond RTs of AS and SS stopping.

The two selective stopping types only differed in the P_r . Both a full model and an order- restricted model that included selective stopping type as a factor were more likely than a null model without it as a factor. The effect was only small, however. Model comparison showed that the models including selective stopping type as a factor were 1.13 times more likely.

There was no effect of selective stopping type on the stop-respond RT. The full model without including selective stopping type as a factor was 3.56 times more likely than with it as a factor. The order-restricted models with and without selective stopping type as a factor both returned infinite Bayes factors, because there was no ordered effect of t_d on stop-respond RT at the group level.

Lastly, we intended to analyse the effect of selective stopping type on the SSRT. However, not one subject performed in line with all three predictions of the independent race model in both

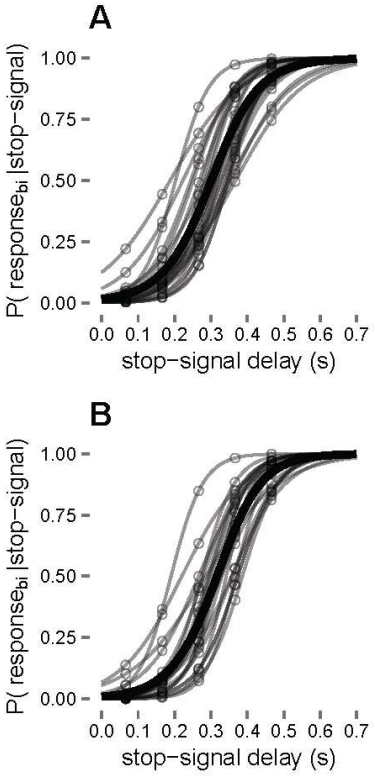

Fig. 3. Inhibition functions for the individual subjects and the group mean for AS (**A**) and SS stopping (**B**). Faded lines and open dots represent individual subjects; bold line represents the group mean.

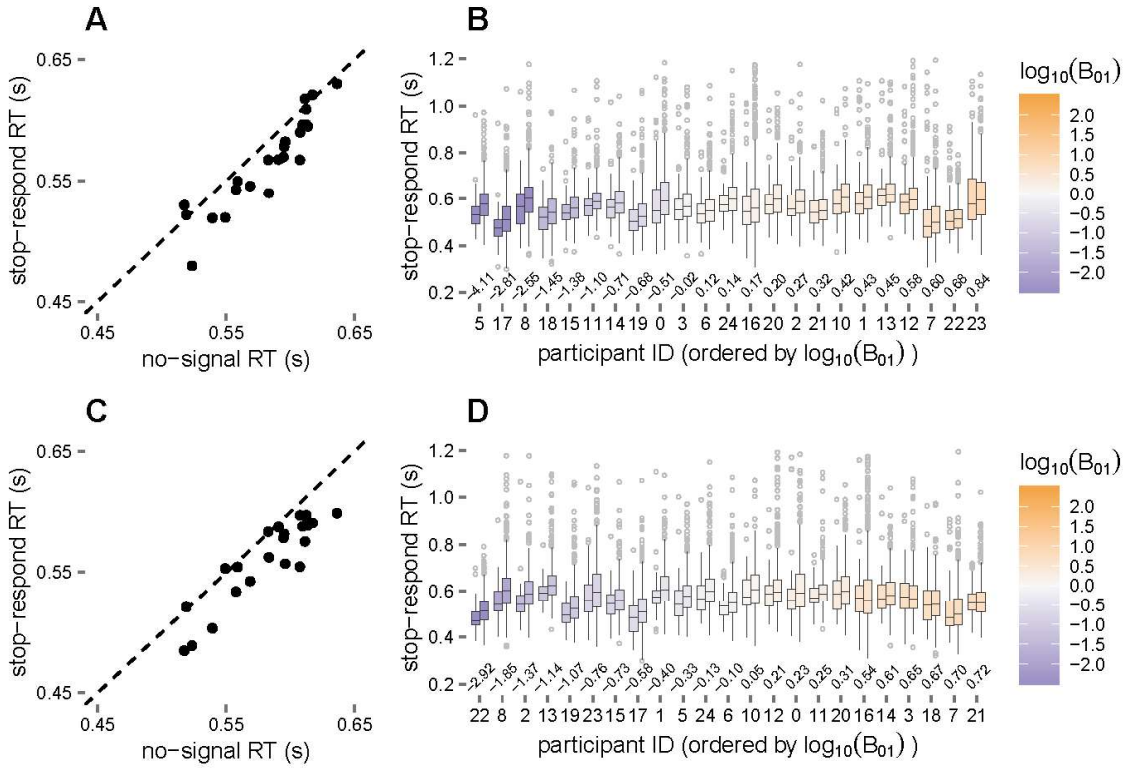

Fig. 4. Relation between stop-respond RT and no-signal RT. Panels **A** and **C** show the individual mean stop-respond against mean no-signal RT for AS and SS stopping, respectively. Panels **B** and **D** show boxplots of the distributions of stop-respond (left) and no-signal (right) RTs for AS and SS stopping, respectively. The values in panels **B** and **D** represent log10 Bayes factors, i.e. log10(B_{01}). Values bigger than 0.5 and smaller than -0.5 indicate more than anecdotal evidence for H_0 and H_1, respectively.

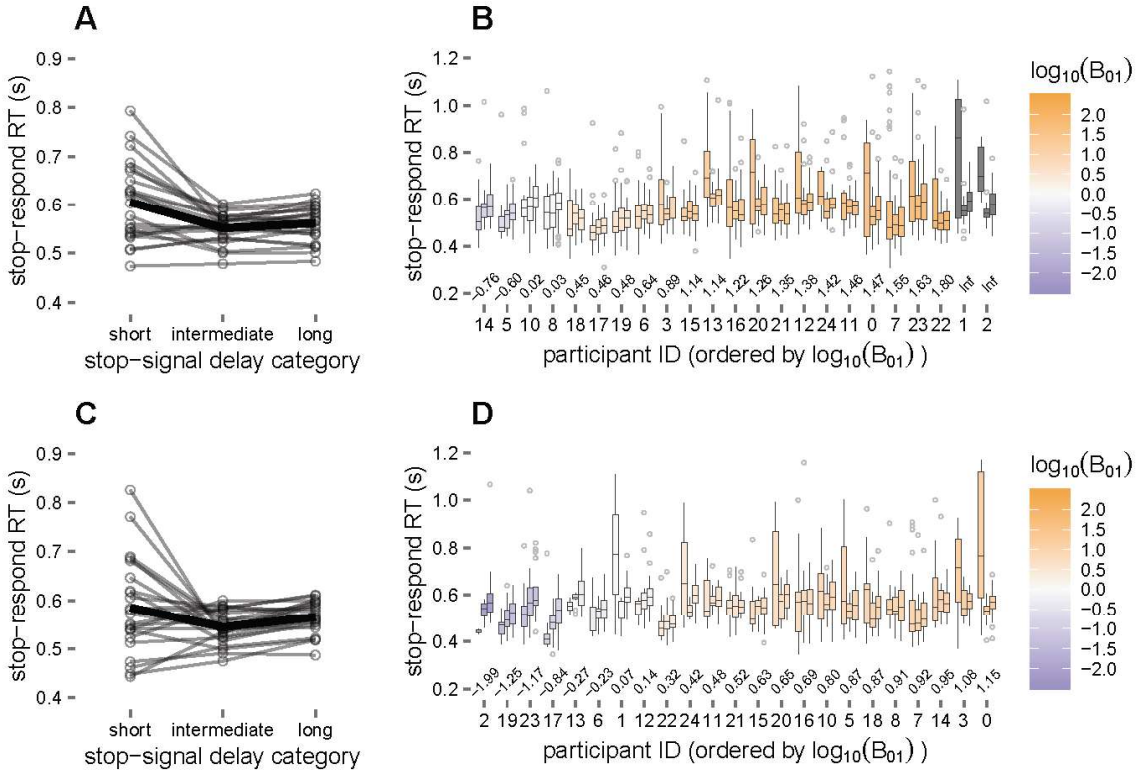

Fig. 5. Effect of stop-signal delay on stop-respond RT. Panels **A** and **C** show the individual mean stop-respond RT for the three td categories in AS and SS stopping, respectively. Panels **B** and **D** show boxplots of the distributions of stop-respond RTs for short (left), intermediate (middle) and long (right) delays in AS and SS stopping, respectively. Conventions as in Figure 4.

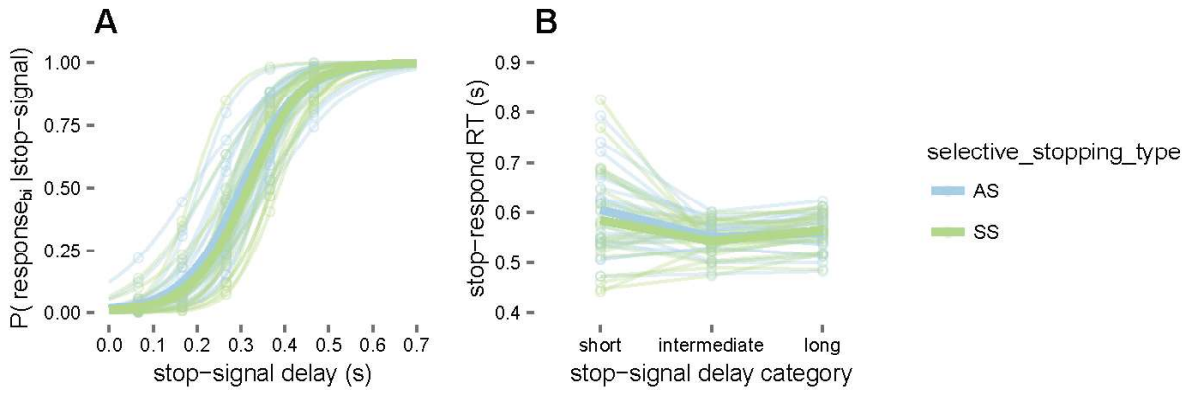

Fig. 6. Comparison of the inhibition functions (**A**) and stop-respond RTs in the three t_d categories (**B**) for AS and SS stopping.

AS and SS stopping. Thus, we could not estimate SSRT reliably, hence we could not analyse the effect of selective stopping type on the SSRT.

Functional Magnetic Resonance Imaging

We excluded two additional datasets from the fMRI analyses because of excessive head movement in the scanner, resulting in a final sample of 21 participants.

To identify brain activations related to stopping in simple stop-signal tasks, previous studies typically used the contrast stop > go. Figure 7 shows the whole-brain activation maps of the contrasts $SB_{inhibit} - NS_{slow}$, $SL_{inhibit} - NS_{slow}$, $SR_{inhibit} - NS_{slow}$ and $IGf_{ast} - NS_{fast}$.

These contrasts are comparable with the simple stop > go contrast, with the exception that they also control for the speed of the GO process. We found that AS and SS stopping, like simple stopping, activate a network of regions in the frontal and parietal lobe as well as the basal ganglia, suggesting that the task manipulation worked.

For the analyses below, we used contrasts that control for both the attentional capture of the salient signal and the speed of the GO process (Table 2). We applied these contrasts to isolate AS and SS stopping-related activations in predefined functional ROIs, broader anatomical ROIs and at the whole-brain level.

Functional ROI analyses. First, we tested which of the predefined ROIs were (de)activated in association with selective stopping, using Bayesian one-sample t-tests. Figure 8 shows boxplots of the contrast estimates in the ROIs, colour-coded for the evidence the data provide for the null versus the alternative hypothesis.

AS stop trials were associated with deactivation of the contra-lateral M1 (i.e. stop-right trials deactivated left M1 and vice versa). Two ROIs showed activation associated with AS stopping: left PMd and left SMA. For these ROIs the Bayes factor supported activation ($B_{01} < 1$) in both AS stopping contrasts. There was evidence for absence of AS stopping-related activation ($B_{01} > 1$) in the left and right IFG, left and right IFJ and the left and right striatum. The other ROIs were activated in one of the AS stopping contrasts but not the other, providing mixed evidence.

SS stop trials were associated with deactivation of both motor cortices. Four other ROIs showed deactivation associated with SS stopping: left IFJ, right PMd, left and right striatum. In the other ROIs there was evidence for no SS stopping-related activation.

There was a difference between activations associated with AS stopping and SS stopping in the left IFJ, left and right PMd, left pre-SMA, left SMA and left and right striatum. For these ROIs there was evidence for a difference ($B_{01} < 1$) in both contrasts subtracting SS stopping-related activations from AS stopping-related activations. There was evidence for no effect of selective stopping type in the right IFG and the right pre-SMA. The other three ROIs were activated in one of the contrasts but not the other, providing mixed evidence for an effect of selective stopping type.

The contrast estimates in the left PMd and left SMA were positive in the contrasts subtracting SS stopping-related activations from AS stopping-related activations and they were activated in both AS stopping contrasts. Thus, the left PMd and left SMA were activated during AS stopping and their activation was greater in AS stopping than in SS stopping.

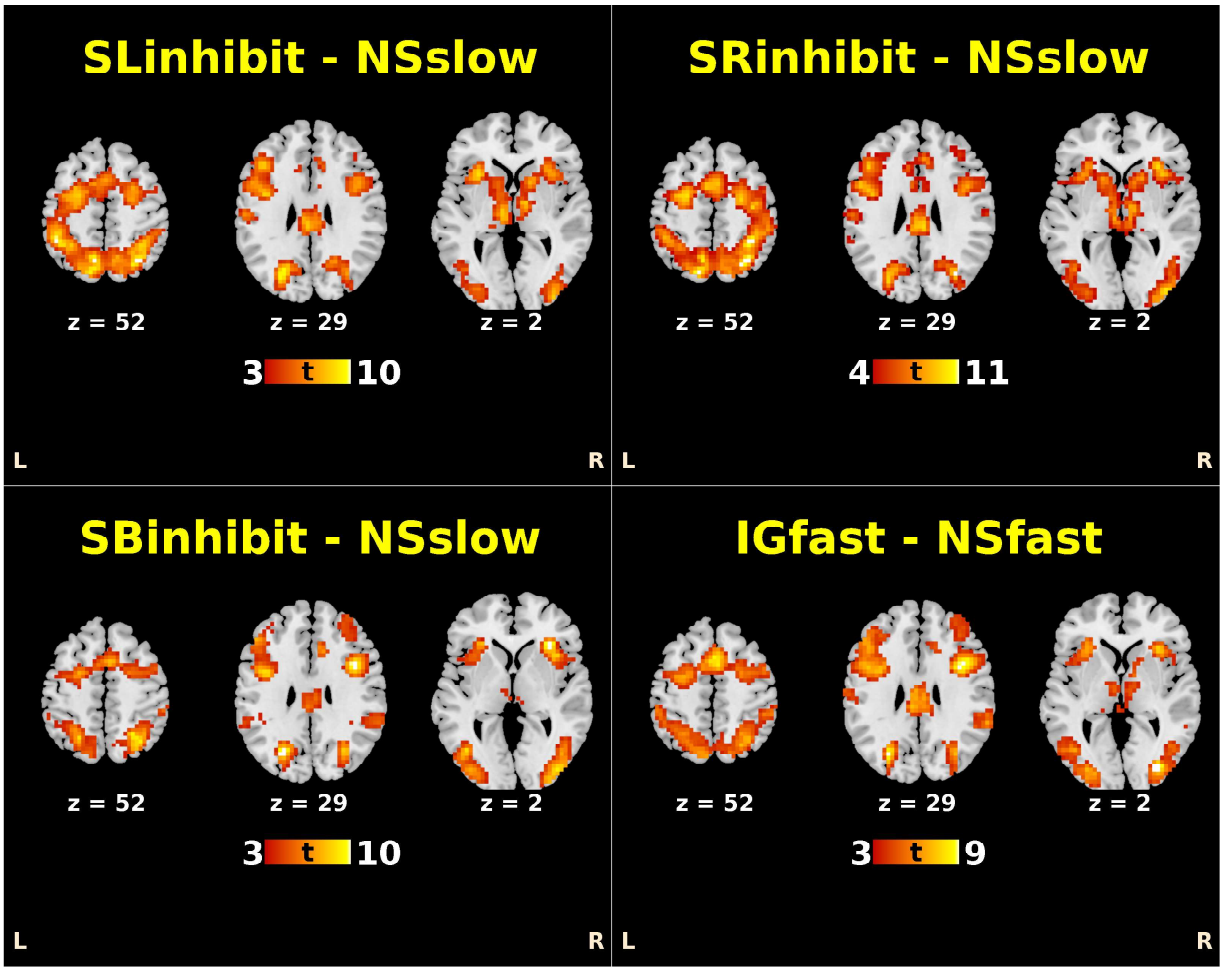

Fig. 7. Whole-brain activation on signal trials. One-sample t-tests of the contrasts $SB_{inhibit} - NS_{slow}$, $SL_{inhibit} - NS_{slow}$, $SR_{inhibit} - NS_{slow}$ and $IG_{f\,ast} - NS_{fast}$. The activations are significant at the cluster-level (cluster-defining threshold $p < .001$ uncorrected; cluster probability $p < .05$, family wise error-corrected for multiple comparisons) and are overlaid on a template brain in MNI space (neurological orientation).

Broad anatomical ROI analyses. Next, we investigated selective stopping-related brain activation in more broadly defined ROIs (Table 5), based on probabilistic anatomical atlases, using classical hypothesis testing. Local maxima of activations in the conjunctions of the AS stopping contrasts and the contrasts subtracting SS stopping-related activations from AS stopping-related activations are shown in Table 6 and 7. Figure 9 shows the masked and small-volume corrected activation maps.

The left and right PMd showed significant activation associated with AS stopping ($p < .001$ uncorrected and $p < .05$ FWE small-volume correction for multiple comparisons) and there was significantly greater activation in AS stopping than in SS stopping in the left and right PMd, the left PMv, and the left SMA ($S_{AS,left} - S_{SS} \cap S_{AS,right} - S_{SS}$). There were no significant activations in the SS stopping contrast in the anatomical ROIs.

Whole-brain analyses. We also applied the same contrast and conjunctions at the whole-brain level to identify activations associated with selective stopping outside the key inhibitory control areas.

Besides the earlier reported activations in the left and right PMd, the whole-brain analysis revealed that AS stopping was associated with activations in the left superior and inferior parietal lobule (Table 8, Fig. 10A). Furthermore, activations associated with AS stopping were greater than SS stopping-related activations in superior prefrontal areas, including the left and right PMd, the left precentral gyrus, superior and inferior parietal areas and the left thalamus and right cerebellum (Table 9, Fig. 10B). There were no significant activations in the SS stopping contrast at the whole-brain level.

Exploratory analyses

We performed two exploratory analyses. First, we compared the observed stop-respond RTs with the

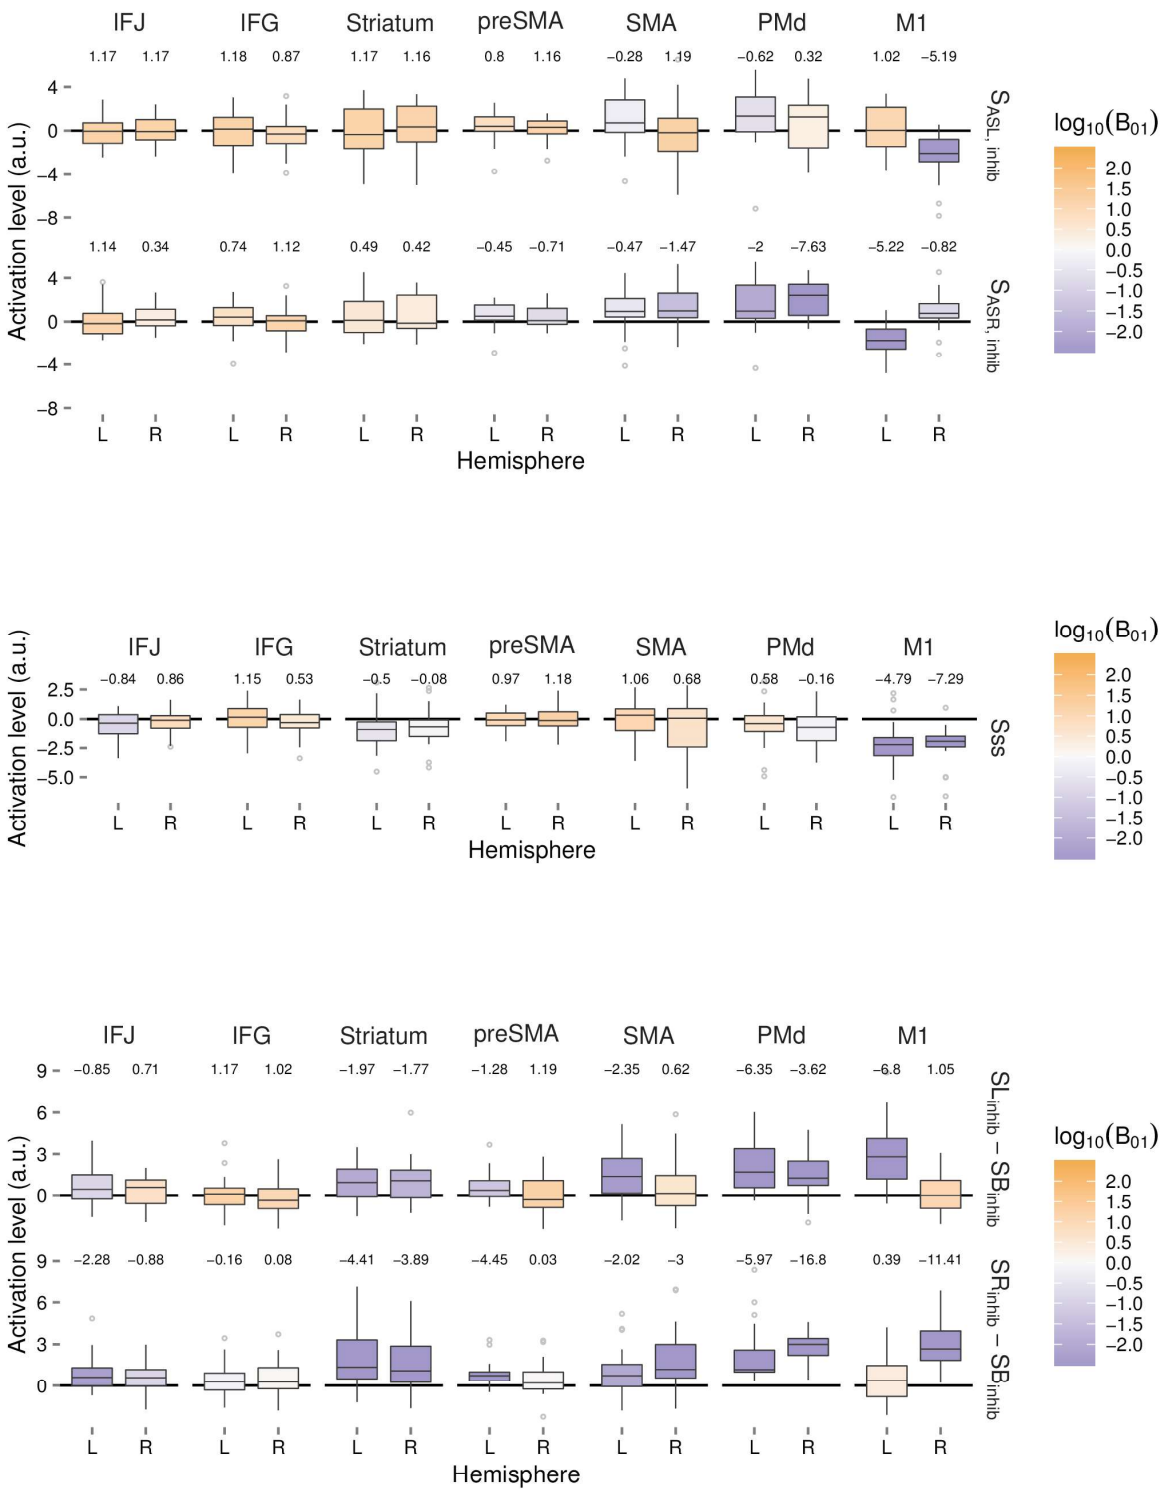

Fig. 8. Boxplots of the activation levels (a.u.) in the predefined functional ROIs. The upper panel shows the activation in the AS stopping contrasts: $S_{AS,left}$ and $S_{AS,right}$. The middle panel shows the activation in the SS stopping contrast: SSS. The lower panel shows the activation in the contrasts subtracting SS stopping-related activations from AS stopping-related activations: $S_{AS,left} - S_{SS}$ and $S_{AS,right} - S_{SS}$. The boxplots are color coded with the corresponding Bayes fit'tors. Conventions of the color coding and values as in Figure 4.

Table 6.

Local maxima of brain activations in the anatomical ROIs during AS stopping (*SAS,left* ∩ *SAS,right*) in MNI x-, y-, and z- coordinates with associated Z-score and small-volume corrected p- value at $p < .05$ FWE-corrected.

Region	Hemisphere	x	y	z	Voxel Z value	*pz*
PMd	L	-24	-4	56	4.76	.004
PMd	R	26	-7	52	4.62	.007

Table 7.

Local maxima of differences in brain activations in the anatomical ROIs between AS and SS stopping (*SAS,left – SSS* ∩ *SAS,right – SSS*) in MNI x-, y-, and z-coordinates with associated Z-score and small-volume corrected p-value at $p < .05$ FWE- corrected.

Region	Hemi-sphere	x	y	z	Voxel Z value	*pz*
PMd	L	-24	-7	56	5.90	.000
PMd	R	26	-7	52	5.57	.000
PMv	L	-55	7	28	5.41	.000
SMA	L	-2	0	52	4.54	.01

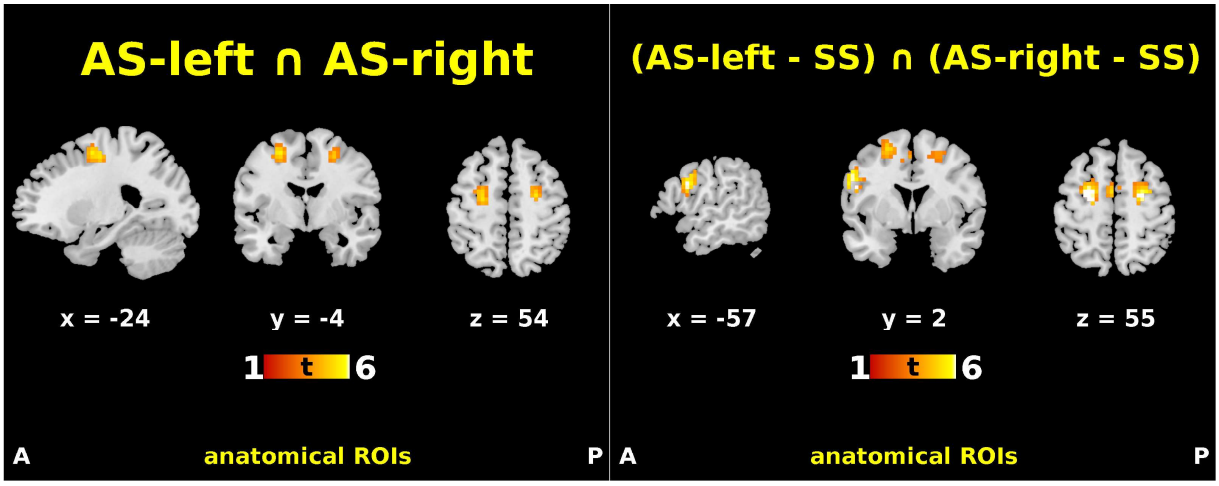

Fig. 9. Brain regions significantly activated in the broad anatomical ROIs in the conjunctions $S_{AS,left} \cap S_{AS,right}$ (left) and $S_{AS,left} - S_{SS} \cap S_{AS,right} - S_{SS}$ (right). Activations are significant at $p < .001$ uncorrected and $p < .05$ FWE small-volume corrected for multiple comparisons, and are overlaid on a template brain in MNI space (neurological orientation).

Table 8.

Local maxima of brain activations (whole-brain analysis) during AS stopping (*SAS, left* ∩ *SAS, right*) in MNI x-, y-, and z-coordinates with associated Z-score ($p < .001$ uncorrected) and cluster size in number of voxels (k; $p < .05$ FWE-corrected).

Region	Hemisphere	x	y	z	Voxel Z value	pz	Cluster size (k)	pk
Superior Parietal Lobule (Precuneus)	L	-10	-63	52	5.45	.000	66	.003
Superior Frontal Gyrus (PMd)	L	-24	-4	56	4.76	.000	55	.007
Superior Frontal Gyrus (PMd)	R	26	-7	52	4.62	.000	35	.045
Inferior Parietal Lobule	L	-41	-38	49	4.46	.000	80	.001

stop-respond RTs that the independent race model predicts (Fig. 11). We performed this analysis, because performance was in line with the second prediction of the independent race model at the group level, yet more than half of the individuals violated it. The model predicts that the mean stop-respond RT of an individual corresponds to that individual's mean of the fast bin of no-signal RTs. In our study, the mean stop-respond RT of most subjects was only slightly faster than their overall mean no-signal RT (fast and slow combined). The independent race model predicts a much larger difference (60 ms on average), as can be seen in Figure 11.

Second, we investigated why the SS stopping contrast yielded no significant activations, by examining the two contrasts from which S_{SS} is built up: $(SB_{stop-inhibit} - NS_{correct-slow})$ and $(IG_{correct-fast} - NS_{correct-fast})$. Subtracting the second subcontrast from the first was supposed to control for activations associated with the attentional capture of the salient stop-signal. However, examination of the two contrasts revealed that both activated the same brain regions (Fig. 12A), to the same degree (Fig. 12B). Thus, subtracting the two subcontrasts for the main S_{SS} contrast completely canceled out the activations.

Discussion

The past decade has seen a surge of interest in selective stopping. Researchers studying se- lective stopping have relied on the independent race model of simple stopping for estimation of the primary outcome measure of stopping, the stop-signal reaction time (SSRT). Further- more, they have investigated selective stopping with a heterogeneous set of tasks, including action-selective stop tasks probing control of specific actions and stimulus-selective stop tasks examining control triggered by specific stimuli. However, it remains unclear whether the independent race model extends to selective stopping and whether selective stopping is a homogeneous or heterogeneous construct. Here, we addressed these important gaps by testing whether selective stopping performance is in agreement with predictions of the independent race model, and by comparing action- and stimulus-selective stopping in terms of performance and brain activation.

We found violations of the predictions of the independent race model in almost all subjects in both AS and SS stopping, suggesting that the model does not apply to selective stopping. Our behavioural and neuroimaging results further suggest that AS and SS stopping were not different in terms of stopping.

Selective stopping involves a race, but not an independent race

We found striking differences between the results of the tests of the independent race model's predictions at the group level and at the individual level. For the group as a whole, selective stopping performance was in agreement with two predictions of the independent race model: the probability of responding increased with stop-signal delay and response times were on average faster on stop-respond than no-signal trials. These findings are in line with previous selective stopping studies that tested predictions of the independence race model (Aron & Verbruggen, 2008; Sebastian et al., 2015; e.g., Smittenaar et al., 2013). However, we found that a third prediction of the model was violated:

Table 9.

Local maxima of differences in brain activations (whole- brain analysis) between AS and SS stopping (*SAS,left – SSS* ∩ *SAS,right – SSS*) in MNI x-, y-, and z-coordinates with associated Z-score ($p < .001$ uncorrected) and cluster size in number of voxels (k; $p < .05$ FWE-corrected).

Region	Hemi-sphere	x	y	z	Voxel Z value	*pz*	Cluster size (k)	p_K
Inferior Parietal Lob-ule	L	-38	-38	42	6.42	.000	613	.000
Superior Frontal Gyrus (PMd)	L	-24	-10	60	6.02	.000	193	.000
Precentral Gyrus	L	-55	4	28	5.75	.000	83	.001
Superior Frontal Gyrus (PMd)	R	26	-7	52	5.57	.000	110	.000
Thalamus	L	-13	-21	10	5.00	.000	129	.000
Superior Parietal Lobule	R	15	-66	63	4.82	.000	73	.002
Cerebellum	R	18	-66	-21	4.60	.000	98	.000
Posterior-Medial Frontal	L	-2	0	52	4.54	.000	42	.023
Inferior Parietal Sul-cus	R	36	-35	38	4.47	.000	54	.008

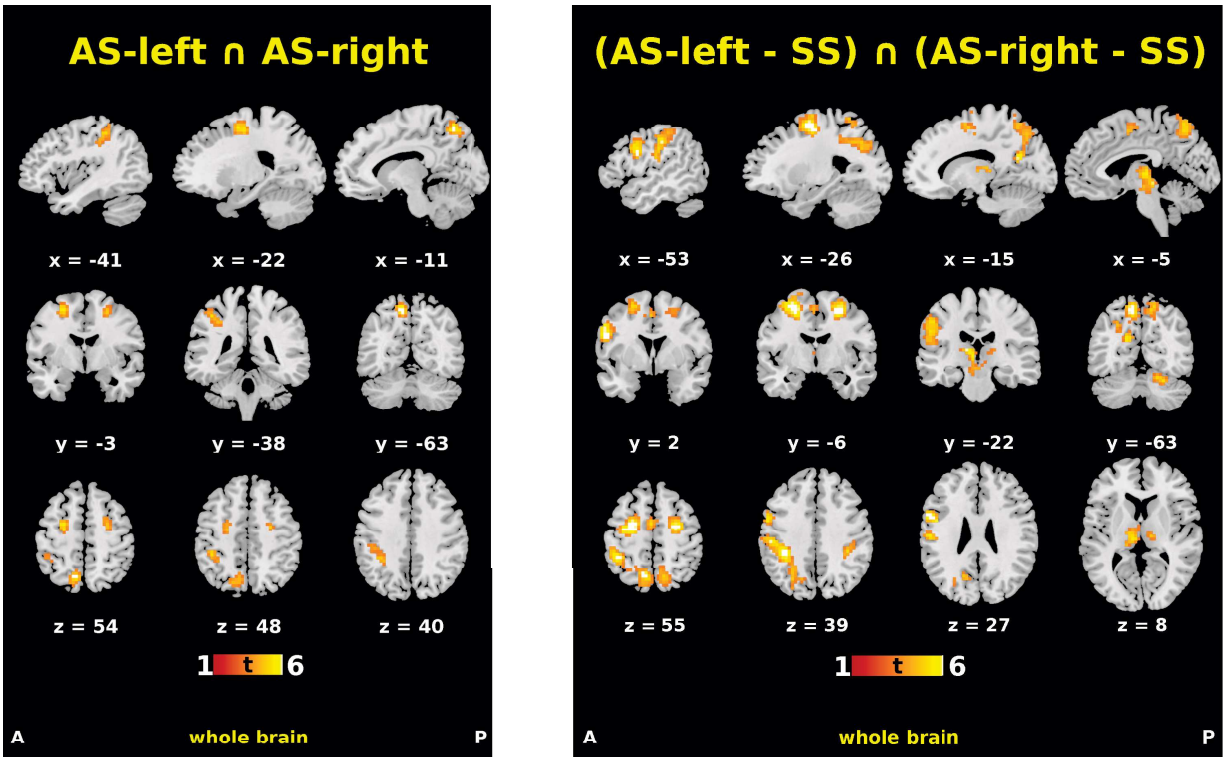

Fig. 10. Brain regions significantly activated at the whole-brain level in the conjunctions $S_{AS,left} \cap S_{AS,right}$ (left) and $S_{AS,left} - S_{SS} \cap S_{AS,right} - S_{SS}$ (right). The activations are significant at the cluster-level (cluster-defining threshold $p < .001$ uncorrected; cluster probability $p < .05$, family wise error-corrected for multiple comparisons) and are overlaid on a template brain in MNI space (neurological orientation).

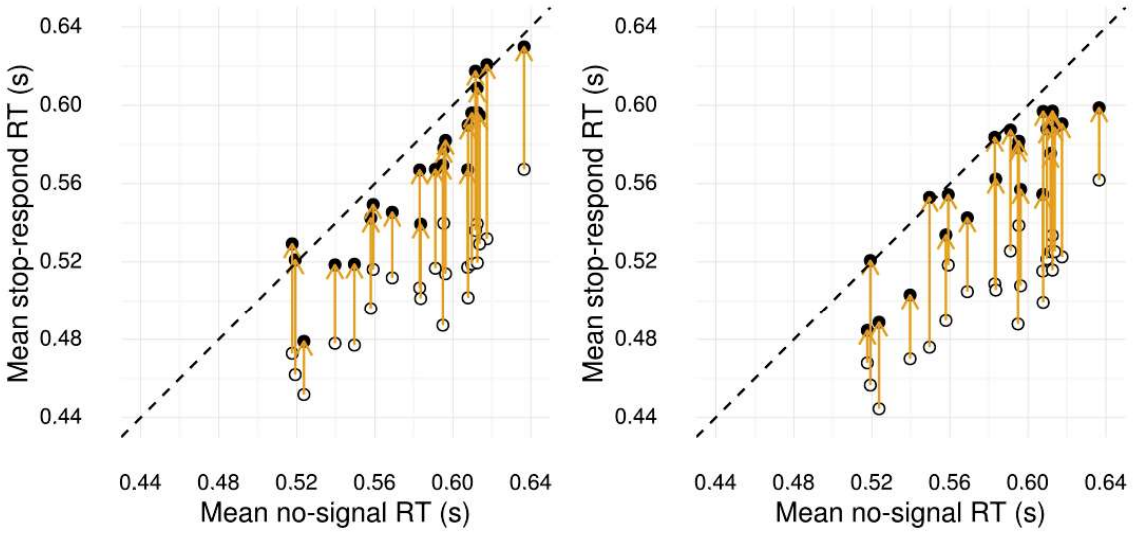

Fig. 11. Difference between the observed stop-respond RTs and stop-respond RTs predicted by the independent race model. Solid dots represent the observed mean RTs, open dots represent the stop-respond RTs predicted by the independent race model, given the observed no-signal RTs.

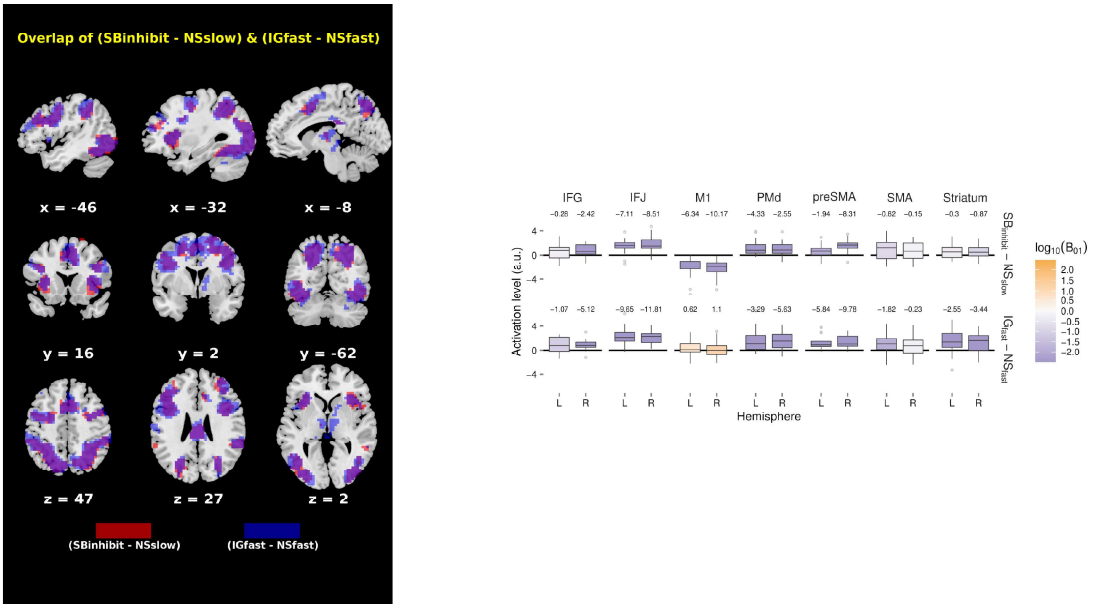

Fig 12. Activations in the S_{SS} subcontrasts $SB_{stop-inhibit} - NS_{correct-slow}$ and $IG_{correct-fast} - NS_{correct-fast}$. The left figure shows the overlap of the activations in the two contrasts in purple. A mask was created of each activation map and they are displayed on top of each other at 0.5 opacity. The right figure shows that activation levels (a.u.) within the functional ROis are similar in the two contrasts.

stop-respond RT did not increase with stop-signal delay. Although one study reported selective stopping performance in line with this prediction (Smittenaar et al., 2013), other studies did not test this assumption.

Violations were more dramatic at the individual level. Although, the prediction that the probability of responding increases with stop-signal delay held for all subjects, at least one of the other two predictions was violated in 96% of the subjects in AS stopping and 78% of the subjects in SS stopping. These results support previous work suggesting that stopping is indeed a race between a GO and a STOP process (Boucher, Palmeri, Logan, & Schall, 2007; Logan & Cowan, 1984; Mallet et al., 2016; Ramakrishnan, Sureshbabu, & Murthy, 2012; Schmidt et al., 2013), but that the assumption of independence between the two processes is violated in selective stopping (Bissett & Logan, 2014; De Jong et al., 1995).

Our results show that violations of the assumptions of the independent race model in individuals can be completely masked by the performance of the group as a whole. To illustrate, Figure 11 shows that even though the observed mean stop-respond RTs were much slower than the independent race model predicts, they were generally still faster than the mean no-signal RTs. Given the data of the group as a whole, the prediction that stop- respond RT is faster than no-signal RT was a thousand times more likely to be true than not in AS stopping, and even more in SS stopping. Nonetheless, testing this prediction for each individual unveiled that it was violated in about two thirds of the subjects in both AS and SS stopping. Estimated SSRTs for these subjects would be invalid, because the SSRT cannot be reliably estimated if this prediction is violated (Logan et al., 2014; Verbruggen & Logan, 2015).

Taken together, the results demonstrate that the independent race model's assumptions of independence often do not hold. The impact of these violations should not be underestimated, because nearly all studies of response inhibition rely on the independent race model, as they use SSRT as an outcome measure. Yet, only few studies report tests of the model's predictions and those that do report tests have performed them at the group level (e.g., Sebastian et al., 2015; Smittenaar et al., 2013). We urge users of stop tasks in general and selective stop tasks in particular to assess and report tests of the independent race model at the individual level and calculate estimates of SSRT if and only if individual datasets meet all qualitative predictions.

Action-selective and stimulus-selective stopping: more similar than different

Figure 6 shows that stopping performance was nearly identical for the two selective stopping types. Our results suggest that AS and SS stopping form a homogeneous construct, and neither involve selective stopping.

The evidence is two-fold. First, there was distinct response slowing on both AS and SS signal trials. The continued response on AS stop trials and the continued responses on ignore trials were both slower than no-signal RTs (see Fig. 2). Such response slowing has been reported before in both AS stopping (e.g., Aron & Verbruggen, 2008; Cai, George, Verbruggen, Chambers, & Aron, 2012; Coxon et al., 2012) and SS stopping (e.g., Bissett & Logan, 2014; Sebastian et al., 2015; Sharp et al., 2010). If subjects had selectively stopped the responses, the continued and ignore RTs should not differ from no-signal RTs.

Second, there was no difference between brain activation on stop-both trials and ignore trials. An exploratory analysis into the SS stopping contrast revealed that stop-both and ignore trials activated the same brain regions to the same degree (see Fig. 12). That implies that there was inhibition-related activity not just on stop trials, but also on ignore trials. Taken together, these results suggest that, instead of stopping selectively, subjects globally inhibited all responses when a signal occurred, and subsequently selectively re-initiated, or released inhibition of, the correct response (Aron & Verbruggen, 2008; Bissett & Logan, 2014).

There was a difference between AS and SS stopping, however, in the comparison of their associated brain activations. The functional and anatomical ROI analyses showed that the left and right PMd, the left PMv and the left SMA were activated in AS stopping and more so than in SS stopping (see Fig. 7, Fig. 8 and Table 7). Thus, AS stopping seemed to rely more on brain regions associated with motor planning than SS stopping. We speculate that these activations reflect action reprogramming on AS stop trials, rather than a difference in response inhibition.

If subjects applied non-selective, global inhibition in both AS and SS stopping, then AS and SS stop trials did not differ in terms of stopping. The difference that then remains lies in the re-initiation or continuation of the correct response. In SS stopping, on ignore trials, that correct response is the same as the initial response to the go-stimulus, but on AS stop trials action reprogramming is required: instead of a bimanual response, now only a left-hand or only a right-hand response must be made. In the fMRI contrasts that we used, the activity on ignore trials was subtracted from the activity on stop trials. Since, there appeared to be inhibition-related activity on ignore trials, this activity was subtracted from the inhibition- related activity on the stop trials. Thus, the AS stopping activations that survived the subtraction may reflect the action reprogramming that is required on AS stop trials but not on SS stop trials, rather than a difference in response inhibition. Previous findings of areas associated with action reprogramming are in line with the current observations of the PMd, PMv and SMA in the ROI analyses (Buch, Mars, Boorman, & Rushworth, 2010; Chambers et al., 2007; Coxon et al., 2016; Mirabella et al., 2011) and the precentral, and superior and inferior parietal areas in the whole-brain analysis (Mars, Piekema, Coles, Hulstijn, & Toni, 2007).

Conclusion

Nearly all selective stopping research has relied on the independent race model of simple stopping. In addition, selective stopping has been investigated with a heterogeneous set of tasks, including action-selective and stimulus-selective stopping paradigms, implicitly assuming that selective stopping is a homogeneous construct. However, it has been unclear whether the independent race model extends to selective stopping and whether selective stopping is a homogeneous or heterogeneous construct.

Our findings suggest that selective stopping can be modeled as a race, but not as an independent race. We found violations of the independent race model's assumptions in nearly every subject in both action- and stimulus-selective stopping. These individual violations were almost completely masked by the performance at the group level. We therefore urge users of stop tasks in general and selective stop tasks in particular to assess and report tests of the independent race model at the individual level. The results further suggest that action-selective and stimulus-selective stopping form a homogeneous construct, as subjects appeared to stop non-selectively rather than selectively on both trial types.

References

Aron, A. R., Fletcher, P. C., Bullmore, E. T., Sahakian, B. J., & Robbins, T. W. (2003). Stop-signal inhibition disrupted by damage to right inferior frontal gyrus in humans. *Nature Neuroscience, 6*(2), 115–116.

Aron, A. R., & Poldrack, R. A. (2006). Cortical and subcortical contributions to stop signal response inhibition: Role of the subthalamic nucleus. *The Journal of Neuroscience, 26*(9), 2424–2433.

Aron, A. R., & Verbruggen, F. (2008). Stop the presses dissociating a selective from a global mechanism for stopping. *Psychological Science, 19*(11), 1146-1153.

Aron, A. R. (2011). From reactive to proactive and selective control: Developing a richer model for stopping inappropriate responses. *Biological Psychiatry, 69*(12), e55–68.

Bissett, P. G., & Logan, G. D. (2014). Selective stopping? Maybe not. *Journal of Experimental Psychology: General, 143*(1), 455–472.

Boucher, L., Palmeri, T. J., Logan, G. D., & Schall, J. D. (2007). Inhibitory control in mind and brain: An interactive race model of countermanding saccades. *Psychological Review, 114*(2), 376–397.

Buch, E. R., Mars, R. B., Boorman, E. D., & Rushworth, M. F. S. (2010). A network centered on ventral premotor cortex exerts both facilitatory and inhibitory control over primary motor cortex during action reprogramming. *The Journal of Neuroscience, 30*(4), 1395–1401.

Cai, W., George, J. S., Verbruggen, F., Chambers, C. D., & Aron, A. R. (2012). The role of the right presupplementary motor area in stopping action: Two studies with event-related transcranial magnetic stimulation. *Journal of Neurophysiology, 108*(2), 380–389.

Chambers, C. D., Bellgrove, M. A., Stokes, M. G., Henderson, T. R., Garavan, H., Robertson, I. H., Mattingley, J. B. (2006). Executive "brake failure" following deactivation of human frontal lobe. *Journal of Cognitive Neuroscience, 18*(3), 444–455.

Chambers, C. D., Bellgrove, M. A., Gould, I. C., English, T., Garavan, H., McNaught, E., … Mattingley, J. B. (2007). Dissociable mechanisms of cognitive control in prefrontal and premotor cortex. *Journal of Neurophysiology, 98*(6), 3638–3647.

Chen, C. Y., Muggleton, N. G., Tzeng, O. J. L., Hung, D. L., & Juan, C. H. (2009). Control of prepotent responses by the superior medial frontal cortex. *NeuroImage, 44*(2), 537–545.

Coxon, J. P., Goble, D. J., Leunissen, I., Impe, A. V., Wenderoth, N., & Swinnen, S. P. (2016). Functional brain activation associated with inhibitory control deficits in older adults. *Cerebral Cortex, 26*(1), 12–22.

Coxon, J. P., Impe, A. V., Wenderoth, N., & Swinnen, S. P. (2012). Aging and inhibitory control of ac- tion: Cortico-subthalamic connection strength predicts stopping performance. *The Journal of Neuroscience, 32*(24), 8401–8412.

Coxon, J. P., Stinear, C. M., & Byblow, W. D. (2006). Intracortical inhibition during volitional inhibition of prepared action. *Journal of Neurophysiology, 95*(6), 3371–3383.

Coxon, J. P., Stinear, C. M., & Byblow, W. D. (2009). Stop and go: The neural basis of selective movement prevention. *Journal of Cognitive Neuroscience, 21*(6), 1193–1203.

De Jong, R., Coles, M. G., & Logan, G. D. (1995). Strategies and mechanisms in nonselective and selective inhibitory motor control. *Journal of Experimental Psychology. Human Perception and Performance, 21*(3), 498–511.

Dimoska, A., Johnstone, S. J., Barry, R. J., & Clarke, A. R. (2003). Inhibitory motor control in children with attention-deficit/hyperactivity disorder: Event-related potentials in the stop-signal paradigm. *Biological Psychiatry, 54*(12), 1345–1354

Dimoska, A., Johnstone, S. J., & Barry, R. J. (2006). The auditory-evoked n2 and p3 components in the stop-signal task: Indices of inhibition, response-conflict or error-detection? *Brain and Cognition, 62*(2), 98–112.

Gauggel, S., Rieger, M., & Feghoff, T. A. (2004). Inhibition of ongoing responses in patients with parkinson's disease. *Journal of Neurology, Neurosurgery, and Psychiatry, 75*(4), 539–544.

Hanes, D. P., Patterson, W. F., & Schall, J. D. (1998). Role of frontal eye fields in countermanding saccades: Visual, movement, and fixation activity. *Journal of*

Neurophysiology, 79(2), 817–834.

Jahfari, S., Waldorp, L., Van den Wildenberg, W. P. M., Scholte, H. S., Ridderinkhof, K. R., & Forstmann, B. U. (2011). Effective connectivity reveals important roles for both the hyperdirect (fronto- subthalamic) and the indirect (fronto-striatal-pallidal) fronto-basal ganglia pathways during response inhibi- tion. *The Journal of Neuroscience, 31*(18), 6891–6899.

Janssen, T. W. P., Heslenfeld, D. J., Van Mourik, R. , Logan, G. D., & Oosterlaan, J. (2015). Neural correlates of response inhibition in children with attention-deficit/ hyperactivity disorder: A controlled version of the stop-signal task. *Psychiatry Research, 233*(2), 278–284.

Kass, R. E., & Raftery, A. E. (1995). Bayes factors. *Journal of the American Statistical Association, 90*(430), 773–795.

Li, C.-s. R., Huang, C., Constable, R. T., & Sinha, R. (2006). Imaging response inhibition in a stop-signal task: Neural correlates independent of signal monitoring and post-response processing. *The Journal of Neuroscience, 26*(1), 186–192.

Logan, G. D., & Cowan, W. B. (1984). On the ability to inhibit thought and action: A theory of an act of control. *Psychological Review, 91*(3), 295–327.

Logan, G. D. (1994). On the ability to inhibit thought and action: A users' guide to the stop signal paradigm. In D. Dagenbach & T. H. Carr (Eds), *Inhibitory processes in attention, memory, and language* (pp.189–239). Academic Press.

Logan, G. D., Van Zandt, T., Verbruggen, F., & Wagenmakers, E.-J. (2014). On the ability to inhibit thought and action: General and special theories of an act of control. *Psychological Review, 121*(1), 66–95.

Macdonald, H. J., Stinear, C. M., & Byblow, W. D. (2012). Uncoupling response inhibition. *Journal of Neurophysiology, 108*(5), 1492–1500.

Majid, D. S. A., Cai, W., Corey-Bloom, J., & Aron, A. R. (2013). Proactive selective response suppression is implemented via the basal ganglia. *Journal of Neuroscience, 33*(33), 13259–13269.

Mallet, N., Schmidt, R., Leventhal, D., Chen, F., Amer, N., Boraud, T., & Berke, J. D. (2016). Arkypallidal cells send a stop signal to striatum. *Neuron, 89*(2), 308–316.

Mars, R. B., Piekema, C., Coles, M. G. H., Hulstijn, W., & Toni, I. (2007). On the programming and reprogramming of actions. *Cerebral Cortex, 17*(12), 2972–2979.

Mirabella, G., Pani, P., & Ferraina, S. (2011). Neural correlates of cognitive control of reaching movements in the dorsal premotor cortex of rhesus monkeys. *Journal of Neurophysiology, 106*(3), 1454–1466.

Nichols, T., Brett, M., Andersson, J., Wager, T., & Poline, J.-B. (2005). Valid conjunction inference with the minimum statistic. *NeuroImage, 25*(3), 653–660.

Paré, M., & Hanes, D. P. (2003). Controlled movement processing: Superior colliculus activity associ- ated with countermanded saccades. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience, 23*(16), 6480–6489.

Poser, B. A., Versluis, M. J., Hoogduin, J. M., & Norris, D. G. (2006). BOLD contrast sensitivity enhancement and artifact reduction with multiecho EPI: Parallel-acquired inhomogeneity-desensitized fMRI. *Magnetic Resonance in Medicine, 55*(6), 1227–1235.

Ramakrishnan, A., Sureshbabu, R., & Murthy, A. (2012). Understanding how the brain changes its mind: Microstimulation in the macaque frontal eye field reveals how saccade plans are changed. *The Journal of Neuroscience, 32*(13), 4457–4472.

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*(2), 225–237.

Rouder, J. N., Morey, R. D., Speckman, P. L., & Province, J. M. (2012). Default Bayes factors for ANOVA designs. *Journal of Mathematical Psychology, 56*(5), 356–374.

Ruiter, M. B. de, Oosterlaan, J., Veltman, D. J., Van den Brink, W., & Goudriaan, A. E. (2012). Similar hyporesponsiveness of the dorsomedial prefrontal cortex in problem gamblers and heavy smokers during an inhibitory control task. *Drug and Alcohol Dependence, 121*(1), 81–89.

Sallet, J., Mars, R. B., Noonan, M. P., Neubert, F. X., Jbabdi, S., O'Reilly, J. X., & Rushworth, M. F. (2013). The organization of dorsal frontal cortex in humans and macaques. *Journal of Neuroscience, 33*(30), 12255–12274.

Schall, J. D., & Boucher, L. (2007). Executive control of gaze by the frontal lobes. *Cognitive, Affective, & Behavioral Neuroscience, 7*(4), 396–412.

Schmidt, R., Leventhal, D. K., Mallet, N., Chen, F., & Berke, J. D. (2013). Canceling actions involves a race between basal ganglia pathways. *Nature Neuroscience, 16*(8), 1118–1124.

Sebastian, A., Jung, P., Neuhoff, J., Wibral, M., Fox, P. T., Lieb, K., … Mobascher, A. (2015). Dissociable attentional and inhibitory networks of dorsal and ventral areas of the right inferior frontal cortex: A combined task-specific and coordinate-based meta-analytic fMRI study. *Brain Structure & Function, 221*(3), 1635-1651.

Sharp, D. J., Bonnelle, V., De Boissezon, X., Beckmann, C. F., James, S. G., Patel, M. C., & Mehta, M. A. (2010). Distinct frontal systems for response inhibition, attentional capture, and error processing. *Proceedings of the National Academy of Sciences, 107*(13), 6106–6111.

Smith, S. M. (2002). Fast robust automated brain extraction. *Human Brain Mapping, 17*(3), 143–155.

Smittenaar, P., Guitart-Masip, M., Lutti, A., & Dolan, R. J. (2013). Preparing for selective inhibi- tion within frontostriatal loops. *Journal of Neuroscience, 33*(46), 18087–18097.

Studholme, C., Hill, D., & Hawkes, D. (1999). An overlap invariant entropy measure of 3D medical image alignment. *Pattern Recognition, 32*(1), 71–86.

Thakkar, K. N., Schall, J. D., Boucher, L., Logan, G. D., & Park, S. (2011). Response inhibition and response monitoring in a saccadic countermanding task in schizophrenia. *Biological Psychiatry, 69*(1), 55–62.

Van de Laar, M. C., Van den Wildenberg, W. P. M., Van Boxtel, G. J. M., & Van der Molen, M. W. (2010). Processing of global and selective stop signals: Application of donders' subtraction method to stop-signal task performance. *Experimental Psychology, 57*(2), 149–159.

Van de Laar, M. C., Van den Wildenberg, W. P. M., Van Boxtel, G. J. M., & Van der Molen, M. W. (2011). Lifespan changes in global and selective stopping and performance adjustments. *Frontiers in Psychology, 2,* 357.

Van den Wildenberg, W. P. M., Van Boxtel, G. J. M., Van der Molen, M. W., Bosch, D. A., Speelman, J. D., & Brunia, C. H. M. (2006). Stimulation of the subthalamic region facilitates the selection and inhibition of motor responses in parkinson's disease. *Journal of Cognitive Neuroscience, 18*(4), 626–636.

Van den Wildenberg, W. P. M., Burle, B., Vidal, F., Van der Molen, M. W., Ridderinkhof, K. R., & Hasbroucq, T. (2009). Mechanisms and dynamics of cortical motor inhibition in the stop-signal paradigm: A TMS study. *Journal of Cognitive Neuroscience, 22*(2), 225–239.

Verbruggen, F., & Logan, G. D. (2008). Response inhibition in the stop-signal paradigm. *Trends in Cognitive Sciences, 12*(11), 418–424.

Verbruggen, F., & Logan, G. D. (2015). Evidence for capacity sharing when stopping. *Cognition, 142,* 81–95.

Verbruggen, F., Aron, A. R., Stevens, M. A., & Chambers, C. D. (2010). Theta burst stimulation dissociates attention and action updating in human inferior frontal cortex. *Proceedings of the National Academy of Sciences of the United States of America, 107*(31), 13966–13971.

Wager, T. D., & Nichols, T. E. (2003). Optimization of experimental design in fMRI: a general framework using a genetic algorithm. *NeuroImage, 18*(2), 293–309.

Wetzels, R., & Wagenmakers, E. J. (2012). A default bayesian hypothesis test for cor- relations and partial correlations. *Psychonomic Bulletin & Review, 19*(6), 1057–1064

Zandbelt, B. B., & Vink, M. (2010). On the role of the striatum in response inhibition. PLoS ONE, 5(11), e13848.

Zandbelt, B. B., Bloemendaal, M., Hoogendam, J. M., Kahn, R. S., & Vink, M. (2012). Transcranial magnetic stimulation and functional MRI reveal cortical and subcortical interactions during stop-signal response inhibition. *Journal of Cognitive Neuroscience, 25*(2), 157–174.

Zandbelt, B. B., Buuren, M. van, Kahn, R. S., & Vink, M. (2011). Reduced proactive inhibition in schizophrenia is related to corticostriatal dysfunction and poor working memory. *Biological Psychiatry, 70*(12), 1151–1158.