

Effective Connectivity Reveals Important Roles for Both the Hyperdirect (Fronto-Subthalamic) and the Indirect (Fronto-Striatal-Pallidal) Fronto-Basal Ganglia Pathways during Response Inhibition

Sara Jahfari,¹ Lourens Waldorp,¹ Wery P. M. van den Wildenberg,¹ H. Steven Scholte,¹ K. Richard Ridderinkhof,^{1,2} and Birte U. Forstmann^{1,2}

¹Department of Psychology and ²Cognitive Science Center Amsterdam, University of Amsterdam, 1018 WB Amsterdam, The Netherlands

Fronto-basal ganglia pathways play a crucial role in voluntary action control, including the ability to inhibit motor responses. Response inhibition might be mediated via a fast hyperdirect pathway connecting the right inferior frontal gyrus (rIFG) and the presupplementary motor area (preSMA) with the subthalamic nucleus or, alternatively, via the indirect pathway between the cortex and caudate. To test the relative contribution of these two pathways to inhibitory action control, we applied an innovative quantification method for effective brain connectivity. Functional magnetic resonance imaging data were collected from 20 human participants performing a Simon interference task with an occasional stop signal. A single right-lateralized model involving both the hyperdirect and indirect pathways best explained the pattern of brain activation on stop trials. Notably, the overall connection strength of this combined model was highest on successfully inhibited trials. Inspection of the relationship between behavior and connection values revealed that fast inhibitors showed increased connectivity between rIFG and right caudate (rCaudate), whereas slow inhibitors were associated with increased connectivity between preSMA and rCaudate. In compliance, connection strengths from the rIFG and preSMA into the rCaudate were correlated negatively. If participants failed to stop, the magnitude of experienced interference (Simon effect), but not stopping latency, was predictive for the hyperdirect–indirect model connections. Together, the present results suggest that both the hyperdirect and indirect pathways act together to implement response inhibition, whereas the relationship between performance control and the fronto-basal ganglia connections points toward a top-down mechanism that underlies voluntary action control.

Introduction

How do we stop ongoing movements? Projections from frontal cortex to the globus pallidus pars interna/externa (GPi/e), via the striatum or the subthalamic nucleus (STN), and then back to the cortex (via the thalamus) are thought crucial to voluntary action control (Mink, 1996; Nambu et al., 2002; Aron, 2010; Ridderinkhof et al., 2011). However, the relative contribution of these specific pathways to inhibitory control remains elusive.

Neuronal recordings in monkeys (Isoda and Hikosaka, 2007, 2008), deep brain stimulations in Parkinson patients (van den Wildenberg et al., 2006), and functional or diffusion tensor imaging (Aron and Poldrack, 2006; Aron et al., 2007) suggest that response inhibition is initiated via a fast hyperdirect pathway in the right hemisphere connecting the inferior frontal gyrus (IFG),

the presupplementary motor area (preSMA), and the STN. Conversely, lesions in the rodent striatum (Eagle and Robbins, 2003), clinical groups showing reduced caudate activity (Booth et al., 2005), and functional connectivity data from healthy subjects performing a stop task (Li et al., 2008; Duann et al., 2009) point toward a prominent striatal role within the indirect pathway.

Here we specify the contribution of the hyperdirect and indirect pathways to the efficiency of implementing response inhibition by combining model comparison techniques (Burnham and Anderson, 2004; Mars et al., 2010) with functional and effective connectivity analysis to determine which pathway is best supported by neural activation patterns. To this end, we conducted a functional magnetic resonance imaging (fMRI) study using a combined Simon/stop task (see Fig. 1*B*). The designated response, indicated by the color of a go stimulus, competes with response tendencies elicited by the task-irrelevant spatial location that is either corresponding (C) or noncorresponding (NC) with the designated response. On some trials, the go stimulus is followed by a stop-signal tone instructing the subject to abort the go response.

Using a region-of-interest (ROI) approach including the right (r) IFG, preSMA, and the basal-ganglia nodes, we first set out to identify the anatomical model that best explains activation patterns

Received Oct. 6, 2010; revised March 18, 2011; accepted March 21, 2011.

Author contributions: S.J., W.P.M.v.d.W., and B.U.F. designed research; S.J. performed research; S.J. and L.W. contributed unpublished reagents/analytic tools; S.J., L.W., and H.S.S. analyzed data; S.J., L.W., K.R.R., and B.U.F. wrote the paper.

This work was supported by a Mozaiek grant (S.J.), VENI (B.U.F., W.P.M.v.d.W.), and VICI (K.R.R.) from the Netherlands Organization for Scientific Research.

Correspondence should be addressed to Sara Jahfari, Department of Psychology, University of Amsterdam, Roetersstraat 15, 1018 WB Amsterdam, The Netherlands. E-mail: S.Jahfari@uva.nl.

DOI:10.1523/JNEUROSCI.5253-10.2011

Copyright © 2011 the authors 0270-6474/11/316891-09\$15.00/0

on stop trials. To this end, we used a recently developed method (“ancestral graphs”) that can identify effective or functional connectivity, uses a model selection approach to identify the best-fitting model, and allows testing for possibly missing regions (Waldorp et al., 2011). Next, we verified that the overall connection strengths within the best-fitting stop model are higher when participants succeeded versus failed to inhibit their response. Previous work has shown close relationships between interference control and key cortical nodes active during response inhibition (Forstmann et al., 2008b,c). By using a combined Simon/stop task, we further examined the potential effect of interference on fronto-basal ganglia pathways active during response inhibition. Specifically, we predicted higher overall connection strengths on noncorresponding versus corresponding stop trials. Finally, we explored how individual connection strengths within the identified stop model relate to stopping latency [stop signal reaction time (SSRT)] and the amount of experienced conflict (interference effect). If response inhibition can be viewed as a top-down process, then especially the hyperdirect and indirect cortical connections projecting to the basal ganglia should predict behavior.

Materials and Methods

Participants. Twenty young adults (nine male; mean age, 23.55 years; range, 18–33 years) participated in this study. All participants provided written informed consent before the scanning session. The local ethics committee approved the experiment, and all procedures complied with relevant laws and institutional guidelines. All subjects had normal or corrected-to-normal vision.

Task and procedure. As shown in Figure 1B, we used a combination of the Simon task (Simon, 1967, 1990) and the stop-signal task (Logan and Cowan, 1984). To avoid stimulus repetitions, four different target colors were used. On go trials, 10 participants were instructed to press the left response button (with the left-hand index finger) if a yellow or blue square was presented and to press the right response button (with the right-hand index finger) if a red or green square was presented. The color mapping (e.g., blue/yellow, press left) was reversed for the remaining participants. All trials started with the presentation of a fixation cross, followed by a colored square that could appear on the left or right side of the screen. On C trials, responses were spatially compatible with the position of the target (e.g., blue square on the left side of the screen indicating a left-hand button press), whereas on NC trials, responses were opposite with respect to the target location (e.g., blue square on the right side of the screen indicating a left-hand button press). Instructions emphasized that participants should try to respond as fast as possible to the color of the presented stimuli while ignoring the location. In addition, participants were instructed to stop the indicated response if the go stimulus was followed by a sound (stop signal).

Each trial started with a jitter interval of 0, 500, 1000, or 1500 ms to obtain an interpolated temporal resolution of 500 ms. During this interval, a fixation cross was presented and participants were asked to maintain fixation. The colored square was then presented on the left or right side of the fixation cross for a period of 500 ms, indicating a color-cued response. On one-third of all trials, the go stimulus was followed by a stop signal. The stop-signal delay (SSD) between the go stimulus (the colored square) and the stop signal was dynamically

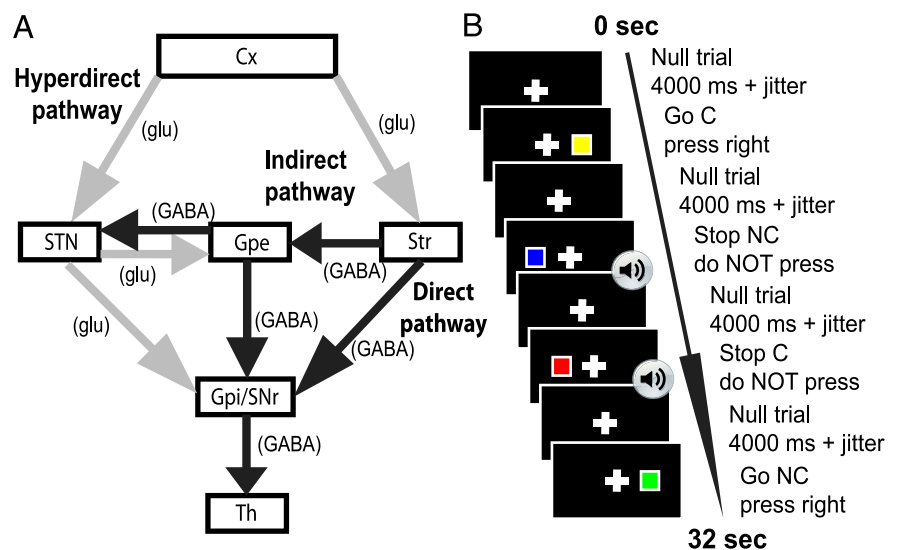


Figure 1. Schematic drawing of the fronto-basal ganglia model and the Simon/stop task. **A**, Fronto-basal ganglia network with the direct (go), the hyperdirect (fast stop), and the indirect (selective stop) pathways. Cx, Cortex; SNr, substantia nigra; Th, thalamus. Gray arrows represent excitatory connections; black arrows represent inhibitory connections. Figure reproduced from Nambu et al. (2002). **B**, Each trial started with the presentation of a fixation cross, followed by a colored square on the left or right side of the screen. Four different colored squares were used as target stimuli; each associated with a left-hand (e.g., yellow/blue) or right-hand (e.g., red/green) index finger press. During C trials, responses were spatially compatible with the position of the target, whereas on NC trials, responses were incompatible with the target location (e.g., blue square on the right side of the screen indicating a left-hand button press). On go trials, participants had a 1250 ms window to respond to the colored stimulus. On stop trials, a tone was played at some delay (SSD) after the presentation of the colored stimulus and instructed the participants to suppress the indicated response and not to press at all. To estimate the hemodynamic response for each trial, trials were separated by a null trial in which only the fixation cross was presented for a period of 4000 ms.

adjusted separately for C and NC trials according to the staircase method to ensure convergence to p (inhibit) of 0.5. For example, if a stop signal was presented on a C trial and the subject responded (“failed stop”), then the SSD for the C staircase was reduced by 50 ms on the subsequent C stop trial; if the subject did not respond (i.e., “successful stop”), then SSD was increased by 50 ms. Initial SSD set to 250 ms for both C and NC stop trials. Each trial had a fixed time length of 4000 ms. If participants had not responded within a time window of 1250 ms after go stimulus presentation, feedback stating “te langzaam” (“too slow” in Dutch) for a duration of 2000 ms was presented. To estimate the hemodynamic response per trial for each subject, each trial (i.e., stop, go, feedback) was followed by a null trial in which only the fixation cross was presented for 4000 ms.

All stimuli were presented on a black-projection screen that was viewed via a mirror system attached to the MRI head coil. Before the MRI session, participants performed a practice block of 36 trials to familiarize with the task. In the MRI scanner, participants subsequently performed two experimental blocks with a total of 192 go (96 corresponding; 96 noncorresponding), and 64 stop (32 corresponding; 32 noncorresponding) trials.

Behavioral analysis. Error rates and median reaction times (RTs) were calculated for correct C and NC go and failed stop trials. The overall interference effect (Simon effect) was computed by subtracting the median RT on C from median RT on NC go trials. For the stop task, SSRT was estimated separately for C and NC stop trials using the so-called “integration method” (Logan and Cowan, 1984; Verbruggen and Logan, 2009). Paired t tests were used to analyze the Simon interference effect on go and stop trials. Pearson’s correlations were used to examine the relationship between stop-signal response inhibition and interference control across subjects. All reported p values are two tailed, unless stated otherwise.

Magnetic resonance imaging scanning procedure. The fMRI data were acquired in a single scanning session, during two runs on a 3 T scanner (Philips). Each scanning run acquired 540 functional T2*-weighted echoplanar images [220² mm field of view (FOV); 96² in-plane resolu-

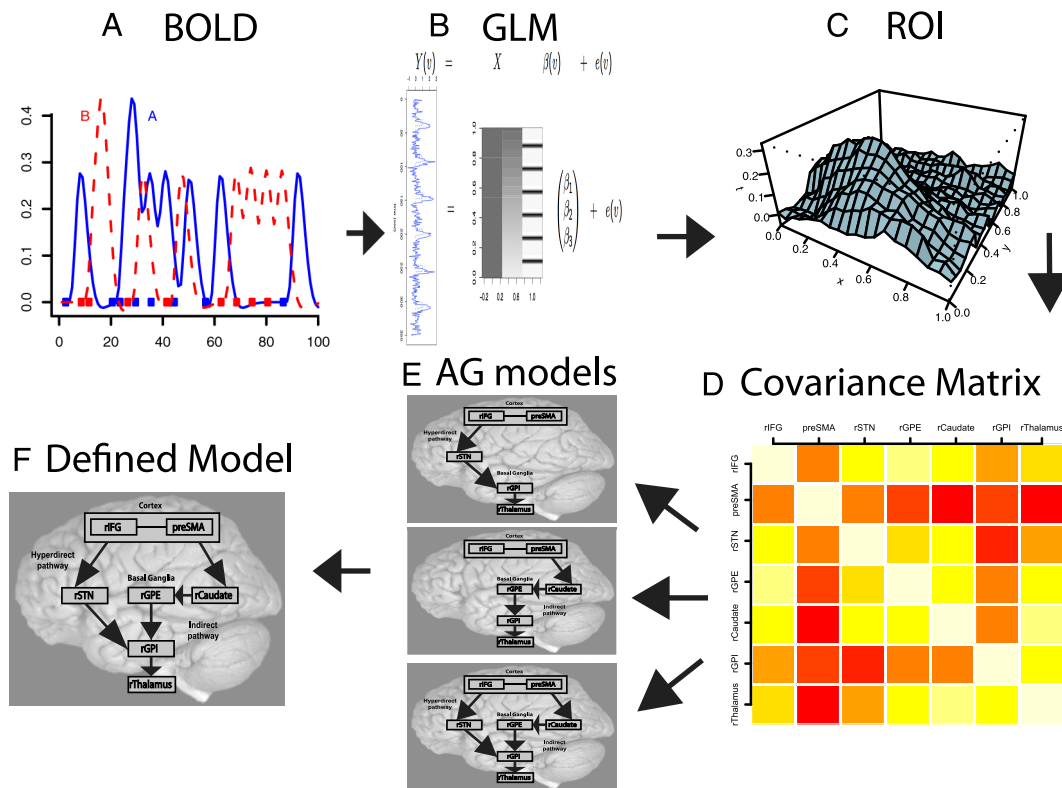


Figure 2. Graphical representation of all steps in determining an ancestral graph for connectivity (A). The event-related BOLD measurements are used as inputs to the general linear model (GLM) (B). C, The seven structural ROIs determine the number of nodes (variables) in the ancestral graph and are used to determine the amplitudes for each trial separately for each region, such that the covariance matrix for these regions can be determined based on trial data (D). The data covariance matrix for each condition is then compared with all defined ancestral graph models, and each model obtains an AIC score. Note that directed arrows represent effective connectivity and so infer directionality. Undirected lines (e.g., preSMA—IFG) represent functional connectivity and infer no directionality (E). The model with the lowest AIC values represents the group network best and is selected (F).

tion; 3 mm slice thickness; 0.3 mm slice spacing; repetition time (TR), 2000 ms; echo time (TE), 28 ms; flip angle (FA), 90°, ascending orientation]. For registration purposes, a three-dimensional T1 scan was acquired before the functional runs (T1; turbo field echo, 220² mm FOV; 256² in-plane resolution; 182 slices, 1.2 mm slice thickness; TR, 9.56 ms; TE, 4.6 ms; FA, 8°, coronal orientation).

Functional magnetic resonance analysis. Analysis was performed using FEAT (for fMRI Expert Analysis Tool) version 5.98, part of FSL for Functional MRI of the Brain (FMRIB) Software Library; www.fmrib.ox.ac.uk/fsl]. The first six volumes were discarded to allow for T1 equilibrium effects. The remaining images were then realigned to compensate for small head movements. The data were filtered in a temporal domain using a high-pass filter with a cutoff frequency of 100 s to correct for baseline drifts in the signal. Finally, the functional data were prewhitened using FSL.

All functional datasets were individually registered into three-dimensional space using the participant’s individual high-resolution anatomical images. The individual three-dimensional image was then used to normalize the functional data into Montreal Neurological Institute (MNI) space by linear scaling (affine transformation). The statistical evaluation was performed using the general linear model. The design matrix was generated with a hemodynamic response function and its first derivative with local autocorrelation correction (Woolrich et al., 2001). To replicate previous fMRI results using a stop-signal task, we computed the following contrasts: (1) successful stop_{null} – go_{null}, (2) failed stop_{null} – go_{null}, (3) NC_{null} – C_{null}, (4) C – null, and (5) NC – null. Higher-level analysis was performed using FLAME (for FMRIB Local Analysis of Mixed Effects) stage 1 and stage 2 with automatic outlier detection (Beckmann et al., 2003; Woolrich et al., 2004, 2009). For the whole-brain analysis, Z (Gaussianized T/F) statistic images were thresholded using clusters determined by $z > 3.1$ and $p = 0.05$ (using Gaussian random field theory).

Ancestral graphs and connectivity analysis. To explore the fronto-basal ganglia network dynamics during interference control and response inhibition, the fMRI data were modeled with a recently developed method called ancestral graphs (Waldorp et al., 2011). Figure 2 shows a graphical representation of all computational steps for ancestral graphs connectivity and model comparison. Ancestral graphs can account for functional or effective connectivity between brain areas by taking into account the distribution of blood oxygenation level-dependent (BOLD) activation per region/subject/trial. Importantly, this method infers connectivity by taking into account the distribution of BOLD activation per ROI over trials and therefore does not depend on the low temporal resolution of the time series in fMRI. A graphical model reflects the joint distribution of several neuronal systems with the assumption that, for each individual, the set of active regions is the same. The joint distribution (graphical model) of two nodes is estimated from the replications of condition-specific trials and not from the time series. With this method, we can infer three types of connections: (1) effective connectivity (directed connection, →), (2) functional connectivity (undirected connection, —), and (3) unobserved systems (bidirected connection, ↔). Directed connections are regression parameters in the usual sense (denoted by β), and undirected connections are partial covariances (unscaled partial correlations; denoted by λ). The bidirected connections refer to the covariance of the residuals from the regressions (denoted by ω). These three types of connection can be identified by modeling the covariance matrix as follows:

$$B^{-1} \begin{pmatrix} \Lambda^{-1} & 0 \\ 0 & \Omega \end{pmatrix} (B^{-1}),$$

where B contains the regression coefficients, Λ contains the partial covariances, and Ω contains the covariances between residuals. A random effects model is used to combine identical models across subjects to then compare different models over the whole group using Akaike’s informa-

tion criterion (AIC) (Akaike, 1973), which involves the log-likelihood $L(\theta)$ with q parameters collected in the vector θ for graph G_q :

$$AIC(G_q) = -2L(\theta) + 2q.$$

The graph with the lowest AIC value will be selected. To infer directions from the ancestral graph, it is required that a change in direction implies a change in probability distribution. This is not always the case. For example, a chain from A to B to C is in terms of conditional independencies equivalent to a chain with the directions reversed, that is from C to B to A (for more details, see Waldorp et al., 2011). Two equivalent models, such as those just mentioned, will result in the same AIC value, indicating that directionality cannot be inferred. The most important structure is when two arrowheads meet (a collider). This will always result in a change in AIC value. The causal interpretations of the connections from an ancestral graph that is the best model according to the AIC can be briefly described as follows:

$A \rightarrow B$: A is a cause of B,

$A \text{---} B$: A is a cause of B and/or B is a cause of A,

$A \leftrightarrow B$: there is a latent common cause of A and B.

For a more detailed description and cautions on causal interpretations, see Zhang (2008).

The main differences between ancestral graphs described by Waldorp et al. (2011) and dynamic causal modeling (DCM) or structural equation modeling (SEM) are as follows: (1) inference is based on trial-by-trial variation and not on the time series as in DCM or SEM because of the low-frequency sampling in fMRI, (2) both functional and effective connectivity can be represented in a single ancestral graph, which cannot be done in DCM or SEM, and (3) a common unobserved (latent) cause of a connection can be detected. The method of ancestral graphs relies on conditional independencies implied by the topology of the network. Therefore, different models (e.g., different directions of connections) result in different fits to the data. Currently, differences between models are characterized by AIC, which combines both accurate descriptive (for the data at hand) and predictive (for future data) value.

To obtain the single trial images for the ancestral graphs analysis, fMRI data processing was performed using FEAT version 5.98, part of FSL (www.fmrib.ox.ac.uk/fsl). Z (Gaussianized T/F) statistic images were thresholded at $p = 0.01$ (uncorrected). Based on previous reports on the involvement of right hemispheric areas during response inhibition and the specific regions within the hyperdirect and indirect pathway, we defined the following right hemisphere anatomical ROIs as key nodes for the stop trials: (1) rIFG [center of gravity (cog) 51, 19, 17], (2) rPreSMA (cog 9, 24, 50), (3) rGPi (cog 17, -6, -4), (4) rGPe (cog 20, -4, 0), (5) rCaudate (cog 13, 10, 10), (6) rThalamus (cog 11, -18, 7), and (7) rSTN (cog 8, -9, -11). Individual STN masks were derived from a previous study (Forstmann et al., 2010) using ultra-high 7 T scanning that allowed the manual segmentation of the STN of nine participants who did not take part in the present study. From these subjects, an STN template in MNI space was created and the center of gravity coordinate was derived (i.e., cog 8, -9, -11). Visual inspection of the registration ensured the precise localization of the STN.

For each ROI, with the exception of the STN mask, we obtained the standardized amplitude parameter of only the active voxels (as identified with F tests) per subject, per trial, and per condition. For successful stop trials, the average number of parameters per ROI (the average \pm SD number of trials per subject) over the C trials was 15.0 ± 3.3 and 14.3 ± 3.3 for NC trials. For failed stop trials, this average was 17.0 ± 3.3 on C and 17.8 ± 3.3 on NC trials. For the left-hand go trials, we obtained an average of 43.1 ± 5.5 standardized parameters per ROI on C trials and 43.4 ± 4.4 on NC trials. When subjects responded with their right hand on go trials, the average number of parameters over trials per ROI was 43.1 ± 5.5 of C trials and 43.0 ± 5.3 on NC trials. Error trials and misses were excluded from additional analysis.

The aim of our connectivity analysis was to examine how different brain areas work together to implement control during stop trials. To this

Table 1. Models specified for testing with ancestral graphs

Right hemisphere models	Specified path
Indirect model	
Ctx = rIFG	Ctx-rCaudate-rGPe-rGPi-rThalamus
Ctx = preSMA	
Ctx = rIFG + preSMA	
Hyperdirect model	
Ctx = rIFG	Ctx-rSTN-rGPi-rThalamus
Ctx = preSMA	
Ctx = rIFG + preSMA	
Hyperdirect-indirect model	
Ctx = rIFG	Ctx-rCaudate-rGPe-rGPi-rThalamus +
Ctx = preSMA	Ctx-rSTN-rGPi-rThalamus
Ctx = rIFG + preSMA	

Ctx, ROIs used as cortex areas with unique projections to the basal ganglia.

Table 2. Behavioral data for go and stop trials

	C	NC
Median go RT (ms)	613.7 \pm 157.1	624.6 \pm 156.3
Go errors (%)	10.0 \pm 9.4	9.6 \pm 8.7
Median failed-stop RT (ms)	579.6 \pm 146.0	594.3 \pm 139.7
SSD (ms)	358.4 \pm 139.3	367.8 \pm 165.9
P_{inhibit}	0.53 \pm 0.1	0.55 \pm 0.1
SSRT (ms)	263.9 \pm 76.2	265.4 \pm 85.4

Values are mean \pm SDs.

end, model fits were preformed on (1) successful stop C, (2) successful stop NC, (3) failed stop C, (4) failed stop NC, (5) correct go C, and (6) correct go NC. In the case of ancestral graphs with seven regions (the number of structural ROIs that we used), the number of possible models to test is between 1.05×10^9 and 4.04×10^{12} . This set is too large to search through completely. Alternatively, model comparison can be performed on a set of hypothesized models. Table 1 contains a graphical outline of all nine models that were specified to test the contribution of the hyperdirect and the indirect pathway during the process of response inhibition. The specified path refers to the definition of connections between ROIs. To compare the contribution of each model with the AIC criterion, all seven regions are always entered into the model, whereas only specific connections are specified. These specific connections, including directed (\rightarrow), undirected (---), and bidirected (\leftrightarrow) connections, determined the model.

Results

Behavioral performance

Median C and NC RT were calculated after removal of errors. RTs longer than 2.5 SD above the mean were discarded as outliers. This resulted in a data reduction of 4.7% for all go trials and 1.5% for all unsuccessful stop trials.

Table 2 gives an overview of the behavioral data for go and stop trials. There were very few omission errors on go trials (mean \pm SD, $0.26 \pm 0.36\%$, combined over C and NC trials). In general, participants responded marginally more slowly on NC than C go trials ($t_{(19)} = -1.99$, $p = 0.06$). However, interference did not affect the overall percentage of choice errors ($t_{(19)} = 0.23$, $p = 0.8$). This was probably because addition of stop signals to the Simon paradigm makes people more reluctant to give very fast responses. During stop trials, we found no additional effects of interference when inspecting RT on trials in which participants failed to stop ($t_{(19)} = -1.44$, $p = 0.16$), the SSD ($t_{(19)} = -0.54$, $p = 0.6$), and SSRT ($t_{(19)} = -0.09$, $p = 0.9$). A recent study, focusing on the effect of fMRI noise on performance, has shown a significant reduction of the mean Simon interference effect when participants were tested while hearing fMRI-scanner noise (Hommel et al., 2011). Possibly, the long intertrial intervals cho-

Table 3. Sequential behavioral analysis, from both go and stop trials

	c-C	c-NC	nc-C	nc-NC
Median go RT (ms)	634.8 ± 162.2	653.7 ± 161.1	647.1 ± 155.4	641.4 ± 152.0
Median failed-stop RT (ms)	585.5 ± 154.7	610.8 ± 134.7	564.7 ± 106.3	590.2 ± 140.5
SSRT (ms)	252.9 ± 84.1	281.3 ± 92.8	274.4 ± 92.2	255.5 ± 103.4

Reaction time on corresponding and noncorresponding trials when taking into account the correspondence of the previous ($n - 1$; indicated by lowercase letters) trial (c-C, corresponding–corresponding; c-NC, corresponding–noncorresponding; nc-C, noncorresponding–corresponding; nc-NC, noncorresponding–noncorresponding). Values are mean ± SDs.

sen to gain an optimal BOLD estimation per trail (4 s between each trial), together with the conduction of the experiment in the fMRI scanner, weakened the magnitude of the overall interference effect usually observed during this task on go and stop trials. Next, we tested whether the overall interference effect is related to the time needed to fully suppress a response. SSRT C, SSRT NC, and the average SSRT did not correlate significantly with the overall interference effect (all p values >0.30).

Previous work with the combined Simon/stop task has shown longer SSRTs for NC trials when the preceding trial (e.g., $n - 1$) is C (Verbruggen et al., 2005). In addition, the overall interference effect has been found to be stronger for trials that follow a corresponding trial. Table 3 gives an overview of the sequential analyses. Repeated-measures ANOVA indicated that the compatibility of trial $n - 1$ interacted with the interference effect on trial n ($F_{(1,9)} = 6.48, p < 0.05$) during go trials but not during successful ($F_{(1,9)} = 1.55, p = 0.23$) or failed ($F_{(1,9)} = 0.00, p \leq 1$) stop trials. In line with previous findings, inspection of the observed interaction on go trials revealed a significant interference effect after a C trial ($t_{(19)} = -2.23, p < 0.05$) but not after a NC trial ($t_{(19)} = 0.60, p = 0.56$). No significant interference effects were observed after C or NC trials for successful or failed stop trials (all p values ≥ 0.1).

Regional brain activations during interference and full response inhibition

To investigate the reliability of the fMRI dataset, before connectivity modeling, a conventional set of contrasts was performed to confirm the replication of typical findings (Aron and Poldrack, 2006; Forstmann et al., 2008b,c). In line with previous results, contrasting successful stop trials with go trials revealed activation in the rIFG, preSMA, left and right inferior parietal lobule, the left insula, left frontal pole, and the cingulate gyrus posterior division (Table 4). With the exception of the left frontal pole, similar activation patterns were obtained for the contrast: failed stop versus go. We did not find higher activation in the rSTN for the contrast successful or failed stop more than go. In line with our previous findings, a direct comparison between corresponding and noncorresponding trials did not yield significant clusters.

Ancestral graphs connectivity analysis during stop trials: the hyperdirect and indirect pathway

Because of a lack of sufficient numbers of successful stop trials necessary for modeling the data with the ancestral graphs method, one participant was removed from the connectivity analysis. Random effects analysis across the whole group with the ancestral graphs method pointed toward a right-hemisphere-centered model that involved both the hyperdirect and indirect pathways during successful and failed stop trials. Figure 3 shows the graphical outline of this hyperdirect–indirect model with functional connectivity between the rIFG and the preSMA, and effective connectivity from the rIFG and preSMA to the rSTN and rCaudate, from the rCaudate to the rGPe, from the rGPe

Table 4. Location of increased activation during successful and failed stop trials versus go trials

Anatomical area	Cluster size (mm ²)	MNI coordinates (mm)			Maximum effect size
		x	y	z	
Successful stop > go					
Inferior frontal gyrus	3619	46	26	12	5.97
Medial frontal cortex/preSMA	591	7	33	39	4.28
Inferior parietal lobe	2426	-58	-50	26	6.07
Inferior parietal lobe	4462	58	-44	21	6.61
Insula	398	-37	19	-10	3.94
Frontal pole	248	-38	52	-5	3.94
Cingulate gyrus posterior division	161	1	-28	27	4.39
Failed stop > go					
Inferior frontal gyrus	3630	48	21	8	6.97
Medial frontal cortex/preSMA	1708	6	32	40	5.68
Inferior parietal lobe	1619	-61	-47	25	6.18
Inferior parietal lobe	3405	59	-43	23	6.24
Insula	907	-41	18	-7	5.41
Cingulate gyrus posterior division	176	1	-25	29	4.53

MNI coordinates are those of the gravity point of each cluster. Cluster thresholding with $z > 3.1$ and $p < 0.05$, whole-brain corrected.




Figure 3. Scheme of the hyperdirect–indirect model. The hyperdirect–indirect model was identified as the best model with the ancestral graphs method to explain the pattern of activity in stop trials using a priori anatomically defined ROIs. Undirected lines represent functional connectivity between two regions. Directed black arrows represent effective connectivity between two regions. Functional connectivity was found between the rIFG and the preSMA (straight line). On the left side of this model, the hyperdirect pathway is displayed including both the rIFG and preSMA (cortex) projecting into the right STN. Right STN was found to have effective connectivity with the rGPe, which in turn projects into the thalamus. On the right side, the model is complemented with the indirect pathway. Again, the rIFG and preSMA project into the rCaudate that activates a more selective stopping route, via projections to the rGPe going into the rGPe and finally into the rThalamus.

and rSTN to rGPe, and from the rGPe to rThalamus. As shown in Table 5, both the hyperdirect and indirect pathways were first tested separately. In line with previous findings, the indirect pathway with effective connectivity between the cortex and caudate yielded a better fit than the hyperdirect pathway with direct connections between the cortex and the rSTN (Duann et al., 2009). However, the best-fitting model with the lowest AIC values for all stop conditions was obtained by combining the hyperdirect and the indirect pathways into one single model. Note that the AIC has a preference for simple models and penalizes for the increase

Table 5. Model fits for the indirect and hyperdirect pathways during successful stop and failed stop trials

	Successful stop				Failed stop				Correct go			
	C		NC		C		NC		C		NC	
	AIC	<i>n</i>	AIC	<i>n</i>	AIC	<i>n</i>	AIC	<i>n</i>	AIC	<i>n</i>	AIC	<i>n</i>
Indirect												
rIFG	2296.2	19	2235.7	19	2223.9	19	2235.8	19	6933.4	0	7272.1	0
preSMA	2306.2	19	2261.0	19	2248.0	19	2227.0	19	7004.7	0	7342.7	0
rIFG + preSMA	1997.1	19	1925.1	19	1934.2	19	1918.6	19	5685.7	1	5963.5	1
Hyperdirect												
rIFG	2564.1	19	2469.4	19	2498.4	19	2479.0	19	8132.2	0	8283.6	0
preSMA	2561.9	19	2463.4	19	2508.1	19	2478.5	19	8115.2	0	8279.2	0
rIFG + preSMA	2281.4	19	2152.6	19	2215.3	19	2192.1	19	6935.1	1	7048.8	0
Hyperdirect–indirect												
rIFG	2232.6	19	2154.6	19	2155.8	19	2181.3	19	6826.2	0	7124.7	0
preSMA	2240.5	19	2174.0	19	2189.6	19	2171.9	19	6880.6	0	7190.9	0
rIFG + preSMA	1903.1	19	1792.4	19	1850.9	19	1839.4	19	5529.3	1	5779.8	0

Lower AIC values indicate a better balance between the variance and bias of the estimated model connections. *n* indicates the number of subjects in which the defined model actually fitted the activity pattern in the a priori anatomically defined ROIs.

in complexity associated with combining the two pathways (i.e., $+2 \times$ number of connections). Therefore, a decrease of 80 points (when compared with the indirect pathway) in the AIC value is remarkable, because a decrease of $2 \times 3 = 6$ points already indicates that the model with the lowest AIC is the preferred model. Finally, to infer the specificity of these model fits, all nine stop models were also tested on go trials. As shown on the right side of Table 5, all models were found to have a large increase in the AIC values and showed a poor or no fit to the activity pattern of the ROIs during go trials. Interestingly, these findings support a large body of literature that indicate a specific role for the hyperdirect pathway with connections between the cortex and the rSTN (Aron and Poldrack, 2006; Aron et al., 2007; Isoda and Hikosaka 2008) and suggest that we use both the hyperdirect and the indirect pathway to stop a planned response.

Connection strengths in the hyperdirect–indirect model

In a next step, we used the normalized individual regression (effective connectivity) and correlation values (functional connectivity) to examine whether the connection strengths of the combined model were stronger during successful stop versus failed stop and during NC stop versus C stop trials. Repeated-measures ANOVA was used including the factors stop (successful stop vs failed stop), interference (C vs NC), and connection (the nine connections identified in the best-fitting model). There was a main effect of stop; overall, the connection strengths in the hyperdirect–indirect model were higher when participants managed to inhibit their response (mean \pm SD, -0.011 ± 0.004 ; $F_{(1,18)} = 8.4$, $p = 0.01$) compared with failed stop trials (-0.064 ± 0.02). There were no additional main effects or interactions. These findings are partly in line with our predictions and show that successful response inhibition, but not interference, on noncorresponding stop trials is related to stronger connections within the hyperdirect–indirect model. Note that, for the current design (with long intertrial intervals for an optimal bold estimation per trial), no interference effects were obtained during stop trials. Therefore, more effective interference modulations replicating previous interference effects on SSRT might be more sensitive to modulations in the hyperdirect–indirect model.

Behavior and individual connections strengths in the hyperdirect–indirect model

In a final and more explorative step, we examined how individual behavioral performance in conflict resolution and response inhi-

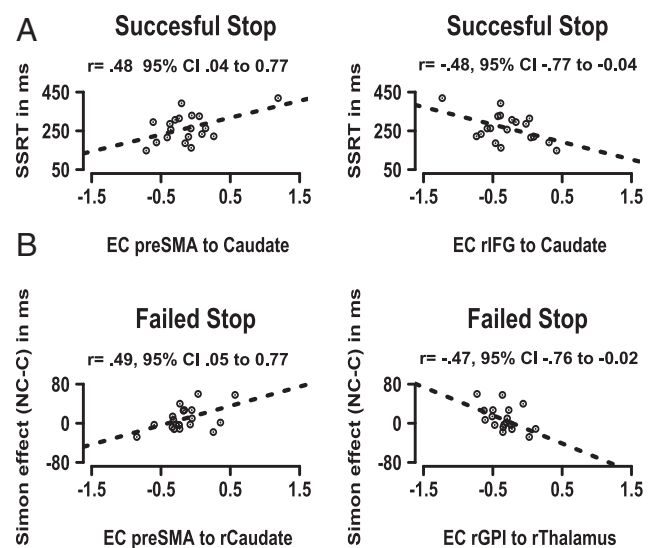


Figure 4. Individual connection strengths in the hyperdirect–indirect model are predicted by behavior. **A**, During successful stop trials, SSRT was predictive for effective connectivity (EC) strengths between the cortex (rIFG and preSMA) and the rCaudate. Left, Positive correlation between SSRT and the connection strength between preSMA to the rCaudate. Right, Negative correlation between SSRT and the connection strengths between rIFG to the rCaudate. **B**, When participants failed to inhibit their response on stop trials (failed stop), the amount of experienced interference (Simon effect) was found to have a positive relationship with the connection strength between preSMA to rCaudate (bottom left) and a negative relationship with the connection strength between rGPI to rThalamus. CI, Confidence interval.

biton might relate to the individual connection strengths in the hyperdirect–indirect model. Because we found no interference effects on stopping behavior or on the hyperdirect–indirect model connections, SSRT and the connection strengths of the hyperdirect–indirect model were collapsed over correspondence conditions.

As displayed in Figure 4A, the time needed to fully stop a response (SSRT) was predictive for the strength of effective connectivity from the preSMA to the rCaudate and from the rIFG to the rCaudate (respectively, $t_{(17)} = 2.2$, $p < 0.05$ and $t_{(17)} = -2.1$, $p < 0.05$; both two-sided, robust regressions). Note, however, that the direction of the relationship from the cortex going into the rCaudate differs between the preSMA and the rIFG. Participants who had higher effective connectivity from the rIFG to the rCaudate needed less time to suppress a response, whereas par-




Figure 5. Effective connectivity (EC) strengths from the rIFG and the preSMA projecting into the rCaudate are closely related during successful stop trials (left) and failed stop trials (right). CI, Confidence intervals.

Participants who had a stronger connectivity from the preSMA to the rCaudate needed more time to stop a response. There were no additional significant relationships between SSRT and the connection strengths of the hyperdirect–indirect model during successful or failed stop trials. No relationships between the overall Simon effect and connection strengths were observed during successful stop trials. On failed stop trials, however, the amount of interference experienced on go trials, as reflected in the Simon effect, was predictive for the connections from the preSMA to the rCaudate ($t_{(17)} = 2.3$, $p < 0.05$, two-sided, robust regression) and from the rGPi to the rThalamus rCaudate ($t_{(17)} = -2.5$, $p < 0.05$, two-sided, robust regression) but in opposite directions (Fig. 4B).

To further examine the reversed relationship between SSRT and the cortico-caudate connections, a final analysis focused on the relationship between the two effective connectivity pathways projecting into the rCaudate. As shown in Figure 5, during successful stop trials, a strong negative relationship was found between the two cortico-caudate connections ($t_{(17)} = -8.0$, $p < 0.0001$, two-sided, robust regression). Although weaker, this negative relationship was also observed during failed stop trials ($t_{(17)} = -3.2$, $p < 0.01$, two-sided). These findings support the reversed relationship between SSRT and connection strengths from both the rIFG and preSMA into the rCaudate, respectively, and suggest diverse but closely synchronized roles for the communication between cortex and the basal ganglia during the process of response inhibition.

Discussion

In the present study, we examined the specific roles and contribution of the hyperdirect and indirect pathways during full response inhibition of planned actions. Using a recently developed method for connectivity analysis of fMRI data that can test functional or effective connectivity in specifically defined models, we found that the hyperdirect and indirect pathways complement each other on stop trials to implement response inhibition. We also found that the overall connection strength of this hyperdirect–indirect model was higher when response inhibition was successful compared with unsuccessful inhibition.

On successful stop trials, the individual efficacy to fully stop a response (SSRT) was differentially predictive for the effective connectivity from cortical areas (rIFG and preSMA) to the rCaudate. Participants who had higher connection strengths between the preSMA and the rCaudate needed more time to stop a response, whereas higher connection strengths between the rIFG and the rCaudate were related to shorter SSRT. This relationship was further supported by the strong negative relationship between the two cortico-caudal projections. On failed stop trials, we found that the amount of experienced interference as indexed by the Simon effect was predictive for effective connectivity in the

hyperdirect–indirect model. Participants who experienced more interference also had a stronger connection between the preSMA and the rCaudate, whereas participants who experienced less interference were found to have a stronger connection between the rGPi and rThalamus.

A large and consistent body of literature has identified the rIFG, preSMA, and the basal ganglia as key areas in the process of motor control or response inhibition (Chambers et al., 2009; Swann et al., 2009; Jahfari et al., 2010; Mars et al., 2010), yet in humans a direct comparison of dominant models that could best explain the activity pattern within and among these areas were lacking. In the current study, we aimed to fill this gap by using an ROI approach in combination with the ancestral graphs method to test specific hypotheses about two well-defined fronto-basal ganglia pathways associated with the process of response inhibition and control. To examine which model best explains the activation pattern in the ROIs, the hyperdirect and indirect pathways were first tested as separate models. On stop trials, results indicated a good fit for both the hyperdirect and indirect pathway separately to explain the variance in ROIs, whereas the AIC identified the indirect pathway as a better model when both pathways were tested separately. We obtained the best model with the lowest AIC values and the best fits by combining the hyperdirect and the indirect pathway into one model. This finding is remarkable because the model selection method used by the ancestral graphs (i.e., the AIC method) has a preference for simple models. Based on these findings, we propose that both the hyperdirect and indirect pathways play a crucial role in the process of response inhibition and work together to implement control on trials in which participants are required to fully stop a planned response.

At this point, only right-hemispheric ROIs were selected to test the relative contribution of the hyperdirect or indirect pathway. This was done because of a loss of power that would have been incurred by including additional left hemispheric nodes. Although an extensive amount of literature has pointed toward the involvement of right-hemispheric nodes during the process of response inhibition (Aron et al., 2006, 2007; Swann et al., 2009; Jahfari et al., 2010), the current experiment cannot make any inferences about hemispheric differences.

In addition, the overall model strengths of the hyperdirect–indirect model were found to be higher for successful stopping compared with failed stopping. In line with these findings, previous studies have shown higher activity in the rIFG in the beta-frequency band for successful stopping compared with failed stopping (Schmajuk et al., 2006; Swann et al., 2009). The current findings indicate that, when the communication between the key areas in the fronto-basal ganglia network is more efficient, suppression of basal ganglia output with downstream inhibitory effects on the primary motor cortex (M1) is successful (van den Wildenberg et al., 2010).

In contrast to our initial predictions, we did not obtain any interference effects on the time needed to stop a response or on the overall model connections of the hyperdirect–indirect model. Previous work with the Simon/stop task has shown interference effects on SSRT but only when the previous trial was C and did not prime the need to increase the level of control (Verbruggen et al., 2005). In the current experiment, we did not obtain any interference effects on SSRT. We note that the main goal of this experiment was to examine which fronto-basal ganglia model best explained the observed pattern of activation during response inhibition. The long intertrial interval (i.e., 4 s between each trial) that was chosen to gain an optimal BOLD estimation for connec-

tivity analysis may have reduced the magnitude of the sequential interference effects previously observed during stop trials. Therefore, more research on the role of interference during stop trials is needed to understand the interplay between these two processes.

Lately, a growing body of literature has focused on the specific roles of the cortex areas deemed relevant for response inhibition (for an overview, see Mostofsky and Simmonds, 2008; Chambers et al., 2009; Ridderinkhof et al., 2011). Recent transcranial magnetic stimulation (TMS) studies indicated that the rIFG has an inhibitory effect on M1 when participants have to cancel an initiated action (Buch et al., 2010) and that the ventral part of this region might play an important role in the updating of action plans (Verbruggen et al., 2010). At the same time, other studies have pointed out that the rIFG might serve an attentional role (Sharp et al., 2010), whereas the preSMA resolves the conflict between the go and stop signal in communication with the rSTN (Nachev et al., 2007; Isoda and Hikosaka, 2008). Shifting the focus to the preparation of action plans and decision-making, recent literature has linked the activity of (but also the structural connection between) the preSMA and the striatum to adjustments in response caution (Forstmann et al., 2008a, 2010), in particular when participants make speeded responses. In addition, the level of activity within the preSMA has been linked to the process of inappropriate response activation (Forstmann et al., 2008c).

Findings from the current study show that higher connection strengths between the preSMA and the rCaudate are related to longer inhibition times on successful stop trials and to a greater experience of interference on trials in which participants fail to stop. At the same time, higher connection strengths between the rIFG and the rCaudate were related to shorter inhibition times. Together, these findings suggest that the connection between the preSMA and the rCaudate might serve as a fast activator of the automatic go response (captured by the location of the stimulus), whereas the connection from the rIFG to caudate helps to update the relevance for an inhibitory act. Within this view, a strong connection between the preSMA and the rCaudate could lead to a failed attempt, or more time needed (SSRT), to prevent the advanced motor plan during stop trials from being executed, but results in a greater experience of interference during NC trials, because the primed automatic response (i.e., triggered by the color) interferes with the desired response (i.e., triggered by location). In line with this interpretation, a recent TMS study showed that, on switch trials, during which an already prepared response was to be inhibited and replaced by an opposite response, the preSMA had an excitatory influence on the primary motor cortex (pointing toward an updating role for action plans), whereas the effect of the rIFG on the primary motor cortex occurred later in time and was inhibitory in nature (Neubert et al., 2010). Alternatively, the preSMA–caudate connection and its negative correlation with SSRT could represent the greater demand on selection of the correct course of action (in this case, action inhibition) in the face of competing alternatives. The data do not allow us to discern between these alternatives.

We note that, although the hyperdirect pathway was found to be an essential part of the best model to explain the activity pattern in the ROIs, no relationships were found between connection strengths in this pathway and the two behavioral indices. The lack of these associations might be related to the speed with which the hyperdirect pathway is thought to operate when exerting global response inhibition. However, more research with standard stop-signal paradigms and techniques that use a high temporal resolution is needed to shed light on this matter.

To summarize, findings from the current study indicate that both the hyperdirect (fronto-subthalamic) and the indirect (fronto-striatal-pallidal) pathway play a crucial role in the process of response inhibition. The close relationship between behavioral indices of cognitive control and individual connection strengths from the cortex projecting into the basal ganglia point toward a top-down controlled inhibitory process, in situations in which we need to withdraw a planned action. Moreover, as the inter-region coupling of the two inhibitory pathways increases, suppression of the motor output is more likely to succeed.

References

- Akaike H (1973) Information theory and an extension of the maximum likelihood principle. In: *Proceedings of the second international symposium on information theory* (Petrov BN, Csaki F, eds), pp 267–281. Budapest: Akademiai Kiado.
- Aron AR, Poldrack RA (2006) Cortical and subcortical contributions to stop signal response inhibition: role of the subthalamic nucleus. *J Neurosci* 26:2424–2433.
- Aron AR, Behrens TE, Smith S, Frank MJ, Poldrack RA (2007) Triangulating a cognitive control network using diffusion-weighted magnetic resonance imaging (MRI) and functional MRI. *J Neurosci* 27:3743–3752.
- Aron AR (2010) From reactive to proactive and selective control: developing a richer model for stopping inappropriate responses. *Biol Psychiatry*. Advance online publication. Retrieved April 12, 2011. doi:10.1016/j.biopsych.2010.07.024.
- Beckmann CF, Jenkinson M, Smith SM (2003) General multilevel linear modeling for group analysis in fMRI. *Neuroimage* 20:1052–1063.
- Booth JR, Burman DD, Meyer JR, Lei Z, Trommer BL, Davenport ND, Li W, Parrish TB, Gitelman DR, Mesulam MM (2005) Larger deficits in brain networks for response inhibition than for visual selective attention in attention deficit hyperactivity disorder (ADHD). *J Child Psychol Psychiatry* 46:94–111.
- Buch ER, Mars RB, Boorman ED, Rushworth MF (2010) A network centered on ventral premotor cortex exerts both facilitatory and inhibitory control over primary motor cortex during action reprogramming. *J Neurosci* 30:1395–1401.
- Burnham KP, Anderson DR (2004) Multimodel inference: understanding AIC and BIC in model selection. *Sociol Methods Res* 33:261–304.
- Chambers CD, Garavan H, Bellgrove MA (2009) Insights into the neural basis of response inhibition from cognitive and clinical neuroscience. *Neurosci Biobehav Rev* 33:631–646.
- Duann JR, Ide JS, Luo X, Li CS (2009) Functional connectivity delineates distinct roles of the inferior frontal cortex and presupplementary motor area in stop signal inhibition. *J Neurosci* 29:10171–10179.
- Eagle DM, Robbins TW (2003) Inhibitory control in rats performing a stop-signal reaction-time task: effects of lesions of the medial striatum and d-amphetamine*1. *Behav Neurosci* 117:1302–1317.
- Forstmann BU, Dutilh G, Brown S, Neumann J, von Cramon DY, Ridderinkhof KR, Wagenmakers EJ (2008a) Striatum and pre-SMA facilitate decision-making under time pressure. *Proc Natl Acad Sci U S A* 105:17538–17542.
- Forstmann BU, Jahfari S, Scholte HS, Wolfensteller U, van den Wildenberg WP, Ridderinkhof KR (2008b) Function and structure of the right inferior frontal cortex predict individual differences in response inhibition: a model-based approach. *J Neurosci* 28:9790–9796.
- Forstmann BU, van den Wildenberg WP, Ridderinkhof KR (2008c) Neural mechanisms, temporal dynamics, and individual differences in interference control. *J Cogn Neurosci* 20:1854–1865.
- Forstmann BU, Anwander A, Schäfer A, Neumann J, Brown S, Wagenmakers EJ, Bogacz R, Turner R (2010) Cortico-striatal connections predict control over speed and accuracy in perceptual decision making. *Proc Natl Acad Sci U S A* 107:15916–15920.
- Hommel B, Fischer R, Colzato LS, van den Wildenberg WPM, Cellini C (2011) The effect of fMRI (noise) on cognitive control. *J Exp Psychol Human Percept Perform*, in press.
- Isoda M, Hikosaka O (2007) Switching from automatic to controlled action by monkey medial frontal cortex. *Nat Neurosci* 10:240–248.
- Isoda M, Hikosaka O (2008) Role for subthalamic nucleus neurons in switching from automatic to controlled eye movement. *J Neurosci* 28:7209–7218.

- Jahfari S, Stinear CM, Claffey M, Verbruggen F, Aron AR (2010) Responding with restraint: what are the neurocognitive mechanisms? *J Cogn Neurosci* 22:1479–1792.
- Li CS, Yan P, Sinha R, Lee TW (2008) Subcortical processes of motor response inhibition during a stop signal task. *Neuroimage* 41:1352–1363.
- Logan G, Cowan WB (1984) On the ability to inhibit thought and action: a theory of an act of control. *Psychol Rev* 91:295–327.
- Mars RB, Shea NJ, Kolling N, Rushworth MFS (2010) Model-based analyses: promises, pitfalls, and example applications to the study of cognitive control. *Q J Exp Psychol (Colchester)* 29:1–16.
- Mink JW (1996) The basal ganglia: focused selection and inhibition of competing motor programs. *Prog Neurobiol* 50:381–425.
- Mostofsky SH, Simmonds DJ (2008) Response inhibition and response selection: two sides of the same coin. *J Cogn Neurosci* 20:751–761.
- Nachev P, Wydell H, O'neill K, Husain M, Kennard C (2007) The role of the pre-supplementary motor area in the control of action. *Neuroimage* 36:T155–T163.
- Nambu A, Tokuno H, Takada M (2002) Functional significance of the cortico-subthalamo-pallidal “hyperdirect” pathway. *Neurosci Res* 43:111–117.
- Neubert FX, Mars RB, Buch ER, Olivier E, Rushworth MF (2010) Cortical and subcortical interactions during action reprogramming and their related white matter pathways. *Proc Natl Acad Sci U S A* 107:13240–13245.
- Ridderinkhof KR, Forstmann BU, Wylie SA, Burle B, and van den Wildenberg, WPM (2011) Neurocognitive mechanisms of action control: resisting the call of the sirens. *Wiley Interdiscip Rev Cogn Science* 2:174–192.
- Schmajuk M, Liotti M, Busse L, Woldorff MG (2006) Electrophysiological activity underlying inhibitory control processes in normal adults. *Neuropsychologia* 44:384–395.
- Sharp DJ, Bonnelle V, De Boissezon X, Beckmann CF, James SG, Patel MC, Mehta MA (2010) Distinct frontal systems for response inhibition, attentional capture, and error processing. *Proc Natl Acad Sci U S A* 107:6106–6111.
- Simon JR (1967) Ear preference in a simple reaction-time task. *J Exp Psychol* 75:49–55.
- Simon JR (1990) The effect of an irrelevant directional cue on human information processing. In: *Stimulus–response compatibility: an integrated perspective* (Proctor RW, Reeve TG, eds), pp 31–88. Amsterdam: Noord-Holland.
- Swann N, Tandon N, Canolty R, Ellmore TM, McEvoy LK, Dreyer S, DiSano M, Aron AR (2009) Intracranial EEG reveals a time- and frequency-specific role for the right inferior frontal gyrus and primary motor cortex in stopping initiated responses. *J Neurosci* 29:12675–12685.
- van den Wildenberg WP, van Boxtel GJ, van der Molen MW, Bosch DA, Speelman JD, Brunia CH (2006) Stimulation of the subthalamic region facilitates the selection and inhibition of motor responses in Parkinson's disease. *J Cogn Neurosci* 18:626–636.
- van den Wildenberg WP, Burle B, Vidal F, van der Molen MW, Ridderinkhof KR, Hasbroucq T (2010) Mechanisms and dynamics of cortical motor inhibition in the stop-signal paradigm: a TMS study. *J Cogn Neurosci* 22:225–239.
- Verbruggen F, Logan GD (2009) Models of response inhibition in the stop-signal and stop-change paradigms. *Neurosci Biobehav Rev* 33:647–661.
- Verbruggen F, Liefvooghe B, Notebaert W, Vandierendonck A (2005) Effects of stimulus-stimulus compatibility and stimulus-response compatibility on response inhibition. *Acta Psychol (Amst)* 120:307–326.
- Verbruggen F, Aron AR, Stevens MA, Chambers CD (2010) Theta burst stimulation dissociates attention and action updating in human inferior frontal cortex. *Proc Natl Acad Sci U S A* 107:13966–13971.
- Waldorp L, Christoffels I, van de Ven V (2011) Effective connectivity of fMRI data using ancestral graphs: dealing with missing regions. *Neuroimage* 54:2695–2705.
- Woolrich MW, Ripley BD, Brady M, Smith SM (2001) Temporal autocorrelation in univariate linear modeling of FMRI data. *Neuroimage* 14:1370–1386.
- Woolrich MW, Behrens TE, Beckmann CF, Jenkinson M, Smith SM (2004) Multilevel linear modelling for FMRI group analysis using Bayesian inference. *Neuroimage* 21:1732–1747.
- Woolrich MW, Jbabdi S, Patenaude B, Chappell M, Makni S, Behrens T, Beckmann C, Jenkinson M, Smith SM (2009) Bayesian analysis of neuroimaging data in FSL. *Neuroimage* 45 [Suppl 1]:S173–S186.
- Zhang J (2008) On the completeness of orientation rules for causal discovery in the presence of latent confounders and selection bias. *Artif Intell* 172:1873–1896.