

## **UC Merced**

### **Proceedings of the Annual Meeting of the Cognitive Science Society**

#### **Title**

Modelling the Correlation Between Two Putative Inhibition Tasks: An Analytic Approach

#### **Permalink**

<https://escholarship.org/uc/item/54d3z20g>

#### **Journal**

Proceedings of the Annual Meeting of the Cognitive Science Society, 32(32)

#### **ISSN**

1069-7977

#### **Authors**

Davelaar, Eddy  
Cooper, Richard

#### **Publication Date**

2010

Peer reviewed

# Modelling the Correlation Between Two Putative Inhibition Tasks: An Analytic Approach

Eddy J. Davelaar (E.Davelaar) and Richard P. Cooper (R.Cooper@bbk.ac.uk)

Department of Psychological Science, Birkbeck, University of London  
Malet Street, London WC1E 7HX, UK

## Abstract

A process of response inhibition is often held to be recruited in situations where it is necessary to withhold or inhibit a prepotent response. Individual differences in the efficacy of this function have been held to underlie individual differences in behaviour on tasks such as the Stroop colour-naming task and the stop-signal task. These claims, however, have been supported only with correlational analyses and informal argument. This paper considers the operationalisation of response inhibition by exploring existing mathematical and process models of both the Stroop and stop-signal tasks. We identify parameters that might underlie individual differences in the performance of the tasks and consider potential relations between those parameters. It is shown that (a) at least three potential inter-relations between parameters of the task models may lead to inter-task correlations, and (b) the observed correlations arise when attentional bias parameters in the models are equated but not when inhibition parameters are equated. We conclude that the ascription of such correlations to a process of response inhibition is premature.

**Keywords:** Cognitive control; Response inhibition; Stroop task; Stop signal task; Individual differences.

## Introduction

In much everyday behaviour, and in many psychological tasks, it is necessary to resist temptation or to avoid producing a prepotent response. Consider the well-known Stroop colour-naming task, where the subject is required to name the colour of the ink in which a word is printed. If the word is itself the name of a colour (e.g., RED printed in green ink) then the subject must actively or deliberately resist the temptation to read the word if they are to successfully name the ink colour.

It is commonly argued that the ability to inhibit a prepotent response is facilitated by a cognitive control process referred to as *response inhibition*. Critically, response inhibition is not a task-specific construct, limited to (e.g.) the Stroop task. Rather, it is held to be one of several general “executive” processes that are invoked across a range of tasks. Moreover, individual differences in the ability to inhibit a prepotent response are held to reflect individual differences in the efficacy of response inhibition. For example, in a well-known study of cognitive control by Miyake, Friedman and colleagues (2000), 137 subjects completed a battery of tasks, three of which were assumed specifically to tap response inhibition. Miyake and colleagues found significant pair-wise correlations in performance on the response inhibition tasks, and

confirmatory factor analysis supported their model of executive function as comprising at least three separable components, one of which was response inhibition.

The three tasks held by Miyake et al. to tap the latent construct were the Stroop task (as discussed above), the stop-signal task of Logan (1994), and an antisaccade task (Roberts et al, 1994). In the stop-signal task subjects complete a series of trials in which they must normally respond as quickly as possible to a stimulus (e.g., by indicating whether an auditorily presented noun denotes a type of animal). On a small proportion of trials the stimulus is followed by a second “stop” stimulus (e.g., a beep), indicating that on this particular trial a response should be withheld. In the antisaccade task trials involved visual presentation of a fixation point at the centre of a monitor screen. This was followed by a brief cue appearing to the left/right of the screen and then an even briefer target appearing on the opposite side of the screen. Subjects were required to make a choice decision based on a feature of the target. To do so, they needed to avoid making a saccade to the cue, as this would prevent them from being able to make a saccade back to the target before it was replaced by a mask. Response inhibition was indexed by Miyake et al. (2000) in the Stroop task by the difference in response times between incongruent and neutral trials. In the stop-signal task it was indexed by the number of stop trials on which a response was (incorrectly) produced. In the antisaccade task it was indexed by the proportion of correct target decisions. As noted above, significant pair-wise correlations were found between these measures. This result was effectively replicated in a subsequent study with 220 subjects which used the same tasks but slightly different dependent measures (Friedman & Miyake, 2004).

The studies of Miyake, Friedman and colleagues appear to provide strong support for the response inhibition construct and for its variability across individuals. However in both cases the evidence is purely correlational. Neither study attempts to provide a mechanistic account of response inhibition as it might be manifest in the various tasks. Clearly, if response inhibition is a cognitive control process that plays a causal role in the performance of the Stroop, stop-signal and antisaccade tasks (amongst others), then that process should be shared by computational accounts of the three tasks. Moreover, if the efficacy of that construct can vary across individuals, then that process should be parameterised in the computational accounts. Lastly, if pair-wise correlations in performance of the tasks are to be attributed at least in part to the efficacy of response

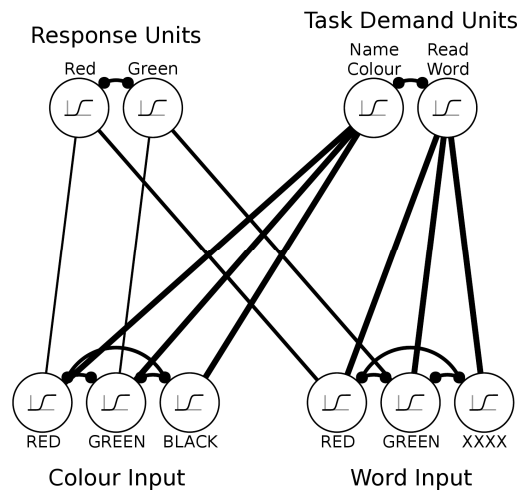


Figure 1: The architecture of the Cohen & Huston (1994) model of the Stroop task. The response function for each node is a sigmoid. Relative line thickness indicates connection strength. Lateral inhibition (shown by arrows with circular end points) operates between nodes in each group.

inhibition, then varying the response inhibition parameter in the computational accounts should also result in pair-wise correlations.

The difficulty, however, is that a cursory analysis of existing process models of the tasks used by Miyake, Friedman and colleagues suggests that their dependent measures are not obviously related to a common mechanism of response inhibition. Consider the widely accepted interactive activation model of the Stroop task of Cohen and colleagues (e.g., Cohen & Huston 1994; see Figure 1). In the model, interference on incongruent trials arises from competition between two response pathways – a word-reading pathway which is highly practiced and hence strong, and a colour-naming pathway which is less practiced and hence somewhat weaker. In order to generate a color-naming response on an incongruent trial it is necessary to selectively amplify the inputs to the color-naming pathway via task-demand units. This process, often referred to as attentional biasing, allows activation from the colour-naming pathway to dominate activation from the word-reading pathway. While individual differences in interference are not generally the focus of this model, they may be captured by assuming that individuals who show relatively little interference are better able to maintain strong excitation of the color-naming task-demand unit. This in turn might result either from greater input to the color-naming task demand unit from external sources (e.g., attentional processes) or conceivably from stronger lateral inhibition between task-demand units. Therefore in this model at least the dependent measure of Miyake et al (2000) indexes an aspect of task-demand, and not response inhibition.

The goal of this paper is to formalise this analysis and extend it to a second putative response inhibition task,

namely the stop-signal task, for which a relatively well-developed “off-the-shelf” computational account is also available (Boucher et al., 2007). We analyse potential sources of correlations in performance across the two tasks by couching both models within a common architecture. In so doing we question the standard concept of response inhibition and propose instead that correlations between performance on the Stroop and stop-signal tasks might be due to a somewhat different factor related to the strength or potency of the currently selected goal.

### The Task Models

In order to address the correlation between the Stroop and stop-signal tasks, we converged on an interactive activation architecture based on the existing published models. This architecture was then simplified to extract a small set of equations that relate the relevant parameters of cognitive control in these two tasks to the dependent measures used by Miyake et al. These equations were then used to generate distributions of the dependent measures by varying the critical parameters and calculating the resulting correlations.

### Stop-signal task

The version of the stop-signal task used by Miyake et al. consisted of two blocks. The principal task was an animacy-categorisation task. The first block only had categorisation trials and was intended to ensure that generating a response was indeed the prepotent response. The second block included 25% stop-trials. For our analytic modelling efforts the following components are relevant. First, the first block produced a mean response time. This was used on a subject-by-subject basis to adjust the onset of the stop-signal on stop-signal trials in the second block. For each subject this onset was their mean response time less 225ms. The stop-signal was therefore presented at (RT–225) ms post-stimulus. We assume a similar approach in the model. Second, the dependent variable was the proportion of categorisation responses generated on stop trials. This value represents errors due to failure to inhibit.

The architecture for our model is inspired by several preceding models. First, Boucher et al. (2007) used a simple interactive race model in which a “go” and a “stop” unit compete through lateral inhibition. Critical for their simulations is that the inhibition from stop unit to go unit is much larger than the reverse connection. This makes the model interactive for only a brief time. Second, the location of the units is downstream in the basal ganglia. This is also assumed in a related go/nogo-model of Frank and colleagues (2004). Third, in the go/nogo-model the nogo-signal comes through the subthalamic nucleus. This nucleus has been shown to form part of a response inhibition pathway that included the inferior frontal gyrus (Aron et al, 2004). It has been postulated that choice responses can be optimised through this pathway (Davelaar, 2009; Frank, 2006). This leaves us with the architecture shown in Figure 2a. It is assumed that the two units are located in the striatum and receive input from earlier processing levels

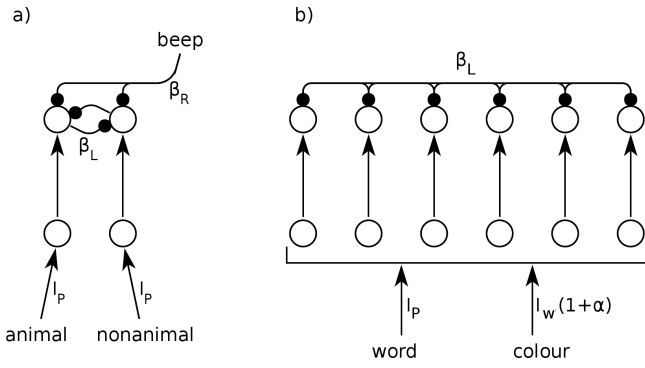


Figure 2: a) Basic architecture of the stop-signal model.  
b) Basic architecture of the Stroop model.

regarding the animacy of the stimulus. The two units are forward connected to two output units that are connected via lateral inhibition. These are assumed to be localised in the globus pallidus interna and substantia nigra. This is a straightforward model of response selection. The stop-signal is assumed to inhibit the final responses via the IFG  $\rightarrow$  STN  $\rightarrow$  SN/GPi-pathway. The strength of the response inhibition parameter  $\beta_R$  is one source of individual differences in stop-signal performance.

### Stroop task

The Stroop task used by Miyake et al. involved naming the colour of a word of which the ink could be in one of six colours. Relevant for the current analytic modelling effort is that the dependent measure is the difference in voice key response time between the mean RT on incongruent trials and neutral trials (which consisted of coloured asterisks). The architecture for the Stroop model follows the incarnations of Cohen and colleagues discussed above. In particular, compared to the neutral trial, an incorrect unit gets activated in response to reading the colour-word. The lateral inhibition between the response units slows down the responses in incongruent trials. Some extension to this model is needed, however. Recent analyses have shown that the Cohen models are unable to produce the correct relation between the stimulus-onset asynchrony in versions of the Stroop task when the word and its ink colour are presented asynchronously (Stafford & Gurney, 2007). The solution to this problem was to treat the output units of the Cohen model as the input units to the basal ganglia, i.e., the striatum (see Figure 2b). This automatically leads to a consistent architecture for both the Stroop task and the stop-signal task.

### Simplifying the Overarching Model

Some simplifications are necessary in order to provide an analysis of the correlation between the two tasks and the relevant parameters. First, we focus only on the abstracted basal ganglia pathway shown in Figure 2. Second, we assume no lateral inhibition in the input level and lateral inhibition of strength  $\beta_L$  at the output level. For the stop-signal model, an extra inhibitory connection of strength  $\beta_R$

to both units is assumed. Whereas in the stop-signal task, the animacy judgement is unambiguous (and prepotent), there is overlap in the Stroop task. This means that in the stop-signal task the only components doing the work are a single response unit and the  $\beta_R$ . In the Stroop task, there are two critical trial types. In the neutral condition the neutral response unit is activated in absence of any inhibition. In the incongruent condition the response unit receives input from the target channel and inhibition from the distractor channel via lateral inhibition. The amount of activation that goes through the target channel is under attentional control. Whereas earlier models of the Stroop task implemented a tradeoff in attention to both channels, recent functional imaging work did not find any support for a deactivation of the distractor channel (Egner & Hirsch, 2005). Instead, only a positive enhancing effect of attention was found in a Stroop-like task. Thus we assume an attentional parameter,  $\alpha$ , which enhances the target channel. We assume that the prepotent inputs for both tasks are identical and that the weaker target channel propagates a weaker signal.

This leads to the following equations that govern the input activation of the target unit in all tasks and conditions:

*Stop-signal task*

$$X(t) = \begin{cases} I_P(t) & \text{when } t < \overline{RT} - 225\text{ms} \\ I_P(t) - \beta_R & \text{when } t > \overline{RT} - 225\text{ms} \end{cases} \quad (1)$$

*Stroop task*

$$X(t) = \begin{cases} (1 + \alpha)I_w(t) & \text{neutral condition} \\ (1 + \alpha)I_w(t) - \beta_L I_P(t) & \text{incongruent condition} \end{cases} \quad (2)$$

In order to obtain response time, we assume a linear output activation function:

$$\frac{d}{dt} F(x) = X(t) \quad (3)$$

This choice is justified by the observation that simple and choice reaction time models operate optimally when they are in the linear part of a sigmoidal output function (Bogacz, et al., 2006). By assuming linear output activation functions, we thus assume optimal responding.

Finally, we assume that the response threshold,  $\theta$ , is the same for both tasks. For the stop-signal task, a response deadline is included of 1500ms (as used in Miyake et al., 2000).

Given the above assumptions, the response time in the stop-signal task is:

$$RT_{SS} = \begin{cases} \frac{\theta - I_P \cdot 225}{I_P - \beta_R} + 225 & \text{for } I_P > \beta_R \\ \infty & \text{otherwise} \end{cases} \quad (4)$$

This is tested against the response deadline. An erroneous response is produced if the response time is less than this deadline. The difference in RTs between incongruent and neutral trials in the Stroop task is:

$$\Delta RT_{Stroop} = \theta \cdot \left\{ \frac{1}{(1 + \alpha) \cdot I_W(t) - \beta_L \cdot I_P} - \frac{1}{(1 + \alpha) \cdot I_W(t)} \right\} \quad (5)$$

For both equations  $\theta$  was fixed at one and noise was added.

One immediate observation of interest is that architecturally, the mechanisms producing incorrect stop-trials and slowed down Stroop trials are not identical. In fact, Stroop performance is determined by the lateral inhibition between two information channels, whereas stop-errors are due to a pathway that inhibits both competing channels.

Our focus is on four parameters: the prepotent response parameter,  $I_P$ , the response inhibition parameter,  $\beta_R$ , the attention parameter,  $\alpha$ , and the lateral inhibition parameter,  $\beta_L$ . There are a number of constraints on the parameters and points to note. First, note that  $I_P$  is shared between the models and moreover that this is the *only* parameter that is shared. Thus, it is expected that this parameter will be the locus of (at least some of) the correlation between the two tasks. Second, the following constraints hold:

- $(1 + \alpha) \cdot I_W > \beta_L \cdot I_P$  in order to ensure that response accuracy in the Stroop task is above 50%
- $\beta_L < \beta_R$ . This is justified based on the findings of Boucher et al. (2007)
- $I_W < I_P$ , by definition

We focus on the following three potential sources of correlation between the proportion of stop-errors and the size of the Stroop interference effect:

**1. Pre-potency of input.** The pre-potency of the input,  $I_P$ , is an obvious choice from the architectural viewpoint, as it is the only parameter that features in both models. Therefore varying  $I_P$  across subjects should produce the positive correlation between the two tasks. The pre-potency, however, is not a factor that is mentioned as an executive function by Miyake et al (2000) and in fact would in most accounts be categorised as the parameter that has to be overcome via executive control.

**2. Correlated executive functions.** To overcome the pre-potency in the stop-signal task, response inhibition,  $\beta_R$ , is the relevant parameter, while for the Stroop task, the attentional control,  $\alpha$ , is the relevant parameter. Obviously, varying these parameters across subjects should not produce a correlation in performance measures. However, one could argue that executive functions are themselves partly correlated (as is done by many authors including Miyake et al., 2000). If this is the case, then a correlation between the two tasks may not be due to shared variance in inhibition parameters, but due to a correlation between the executive functions of inhibition and attentional focus. One possibility that we will come back to in the discussion is that both of these concepts might be subsumed under a more general notion of the strength or potency of the goal, as both tasks require the need to exert control based on the recognition of a stimulus (stop-signal or colour-word).

**3. Correlated inhibition.** Perhaps the most natural way of addressing the correlation is to assume that response inhibition in the stop-signal model,  $\beta_R$ , and lateral inhibition

in the Stroop model,  $\beta_L$ , are correlated. However, note that the dependent variables are such that greater (response) inhibition in the stop-signal task leads to fewer errors and hence lower levels of the dependent measure, whereas greater (lateral) inhibition in the Stroop task leads (perhaps counter-intuitively) to slower responses in the incongruent condition and higher levels of the dependent measure. Thus, correlated inhibition will lead to a correlation in the dependent measures, but this will be a *negative* correlation – not a positive one! Thus correlated inhibition can only result in the observed positive correlation between dependent measures on the stop-signal and Stroop tasks if the inhibition parameters are *negatively* correlated.

## Sampling Studies: Methods and Results

Several sampling studies were conducted based on the above analysis. The aim of these studies was to assess effects of the three potential sources of correlation identified above on the cross-task correlation in dependent measures. To this end, equations 4 and 5 were used to obtain dependent measures for each task as all parameters except  $I_W$  were varied uniformly using boundaries that (a) were found to be adequate to produce values for the dependent measures that were within the range of the actual empirical results and (b) adhered to the set of constraints above.  $I_W$  was fixed to 0.6. The choices of uniform distributions and the precise value of  $I_W$  are arbitrary and do not impact on the conclusions drawn from this work.

We imposed associations among parameters as follows:

1. To address the pre-potency view, only the  $I_P$  distribution was varied between-subjects and each subjects'  $I_P$  value was used in both task models. For each virtual subject, the other three parameters were randomly sampled 100 times corresponding to 100 trials within a task. The proportion of stop-errors was calculated as the mean number of times that a response time in the stop-signal task was shorter than 1500 ms. The Stroop mean interference was calculated over the 100 3-parameter combinations (together with the subject's  $I_P$ ). One hundred subjects were simulated and a Pearson product-moment correlation coefficient was calculated over the resulting set of 100 data pairs.
2. To address the correlated executive function view,  $\beta_R$  and  $\alpha$  were used as between-subjects parameters ( $I_P$  and  $\beta_L$  varied within-subjects). There were two versions: uncorrelated and correlated  $\beta_R$  and  $\alpha$ .
3. Finally to address the correlated inhibition view, we correlated  $\beta_R$  and  $\beta_L$  between-subjects ( $I_P$  and  $\alpha$  varied within-subjects).

In all cases additional noise was added to the correlated parameter in order to lower the resulting correlation in dependent measures and obtain a value of approximately 0.18 as found in the behavioural studies of Miyake et al. (2000).

Scatter plots showing the correlation between dependent measures for four situations are shown in Figure 3. Positive correlations can be obtained between the dependent measures either when  $I_P$  is fixed within-subjects (exploring




Figure 3: Scatter plots of dependent measures under different conditions. a) When  $I_p$  is fixed within-subjects (the pre-potency view) the correlation is positive. b) If  $\beta_R$  and  $\alpha$  vary within-subjects but are uncorrelated, then there is no correlation. c) The correlation is positive when  $\beta_R$  is correlated with  $\alpha$  (the correlated executive functions view). d) When  $\beta_R$  and  $\beta_L$  are positively correlated the correlation between dependent measures is negative.

the pre-potency view; Figure 3a), or when  $\beta_R$  is correlated with  $\alpha$  (exploring the correlated executive functions view; Figure 3c). If  $\beta_R$  and  $\alpha$  vary within-subjects but are uncorrelated, then there is no correlation between the dependent measures (Figure 3b). When  $\beta_R$  and  $\beta_L$  are correlated, then as anticipated the correlation between dependent measures is negative (Figure 3d).

## Discussion

We set out to address the correlation between two well-known tasks that have been discussed as tapping executive inhibition. Correlations between performance on the stop-signal and the Stroop tasks have been found in several behavioural studies and both tasks have been the subject of detailed computational modelling. The modelling has been at the same architectural level, thus allowing the integration of those models into a larger more general model. As the parameters in the models are tied to specific mechanisms, we can address the source of the correlation between the tasks at a parameter level without having to make imprecise verbal assumptions about the relation between mechanisms operating in the two tasks. The general model itself can be simplified without loss of argument and applied to the complex enterprise of not only modelling individual differences in task performance, but also the correlations among tasks.

If the argument is that co-variability in the stop-signal and Stroop task is due to shared variability of a single executive function referred to as response inhibition, then our results question this strong statement. First, the only mechanism that is truly shared between the tasks is the strength of the

pre-potent response channel. Given that this channel is the one that is the target of executive control and thus cannot be considered to be an executive control function itself, we see no basis to assume that a shared inhibition-type of executive function underlies the behavioural correlation. Second, the mechanisms that have been assumed and shown in modelling to be critical in overcoming the pre-potency are different between the tasks, thus a single inhibition-type of executive function is not an appropriate label. Instead, if these mechanisms are correlated, then a more appropriate label might be “goal potency”. We elaborate on this view below. Third, if the shared inhibition function is taken literally and the inhibition mechanisms are correlated, then the simulation suggests that a negative correlation should be found between the tasks. However, the behavioural studies show a positive correlation between the tasks. This is in the context of literature that claims a positive correlation between each task and a latent inhibition factor. These points together argue against the use of a response inhibition construct in the individual differences literature as a mechanistic explanation for the behavioural correlation.

We suggested that the correlation between the tasks is due to the potency or strength of the current goal. More specifically, the computational studies are consistent with either a unitary mechanism that affects the rate of activation accumulation or one that relates to the level of the maximum possible activation. Both of these are emergent from an activation-based framework in which perceptual or cognitive information is actively maintained through self-excitatory loops (Davelaar, et al., 2005). Whether they can be distinguished remains to be demonstrated. However we

note that in a further part Miyake et al.'s (2000) study, it was shown that the common factor underlying performance on the stop-signal and Stroop tasks dissociated from a factor common to performance on several other tasks that were held to require a further executive function, referred to as *task-shifting*. A full account must therefore relate, in computational terms, the function isolated in this study and a task-shifting function. This is particularly important as Gilbert and Shallice (2002) consider task-shifting in the context of the Stroop task, and account for it in a model closely related to the Cohen and Huston model that forms the basis for part of this work.

The idea of goal potency has some support from other areas of cognitive neuroscience. Thus, Duncan et al. (2008) refer to the inability to execute a goal on presentation of a stimulus, even though knowledge about the rules regarding stimulus and response is present, as *goal-neglect*. Individuals differ in the degree to which they exhibit goal-neglect. If goal-neglect (or a factor underlying it) lies behind our factor, then one would expect that the proportion of stop-errors and the size of the Stroop interference effect should both be positively correlated with measures of goal-neglect. We know of no study that has investigated the correlation between stop-errors and goal-neglect.

We have focused only on the stop-signal and the Stroop task. As noted in the introduction, Miyake et al. (2000) also considered the antisaccade task. This task requires an eye-movement away from a distractor stimulus when this stimulus appears. In the Miyake et al. study the dependent measure for this task was the proportion of correct trials. Thus, overcoming pre-potency increases the score. This is important, as for the stop-signal and Stroop tasks, overcoming the pre-potency decreases the corresponding dependent measure. Consequently one might expect a negative correlation between the measures. Instead a positive correlation was found between the antisaccade task and both tasks. This is inconsistent within the response inhibition view. However, processes of active maintenance or activation accumulation can account for positive correlation where overcoming prepotent responses would expect negative correlations. In all but the antisaccade task, the stimulus conveys information that is used in activation of the relevant goal. In the antisaccade task, the first stimulus is a distractor and does not convey positive information, while the second is the target. Therefore being able to quickly activate information will produce less accurate responses. This leads to more activation producing lower levels of the dependent measure (accuracy) in the antisaccade task, together with more activation leading to lower levels of the dependent measures in the stop-signal (proportion stop-errors) and Stroop (interference effect) tasks. Our argument therefore is that the latent factor in the Miyake et al. studies reflects an activation-based function, and not an inhibition function.

This work also demonstrates more generally the importance of using explicit formal analyses to uncover the mechanisms underlying latent cognitive constructs.

## References

- Aron, A. R., Robbins, T. W., & Poldrack, R. A. (2004). Inhibition and the right inferior frontal cortex. *Trends in Cognitive Sciences*, 8(4), 170-177.
- 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.
- Cohen, J. D. & Huston, T. A. (1994). Progress in the use of interactive models for understanding attention and performance. In *Attention and performance 15: Conscious and nonconscious information processing*. C. Umiltà & M. Moscovitch (Eds) (pp. 453-476). Cambridge, MA, US: The MIT Press.
- Davelaar, E. J. (2009). Conflict-monitoring and (meta)cognitive control. In J. Mayor, N. Ruh, & K. Plunkett (Eds.), *Connectionist models of behaviour and cognition II*. (pp. 241-252). Singapore: World Scientific.
- Davelaar, E. J., Goshen-Gottstein, Y., Ashkenazi, A., Haarmann, H. J., & Usher, M. (2005). The demise of short-term memory revisited: empirical and computational investigations of recency effects. *Psychological Review*, 112, 3-42.
- Duncan, J., Parr, A., Woolgar, A., Thompson, R., Bright, P., Cox, S., et al. (2008). Goal neglect and Spearman's g: Competing parts of a complex task. *Journal of Experimental Psychology: General*, 137(1), 131-148.
- Egner, T., & Hirsch, J. (2005). Cognitive control mechanisms resolve conflict through cortical amplification of task-relevant information. *Nature Neuroscience*, 8(12), 1784-1790.
- Frank, M. J. (2006). Hold your horses: A dynamic computational role for the subthalamic nucleus in decision making. *Neural Networks*, 19(8), 1120-1136.
- Frank, M. J., Seeberger, L. C., & O'Reilly, R. C. (2004). By carrot or by stick: Cognitive reinforcement learning in Parkinsonism. *Science*, 306(5703), 1940-1943
- Friedman, N. P., & Miyake, A. (2004). The relations among inhibition and interference control functions: A latent-variable analysis. *Journal of Experimental Psychology: General*, 133(1), 101-135.
- Gilbert, S. J., & Shallice, T. (2002). Task Switching: A PDP Model. *Cognitive Psychology*, 44(3), 297-337.
- 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.
- Miyake, A., Friedman, N. P., Emerson, M. J., Witzki, A. H., Howerter, A., & Wager T. D. (2000). The unity and diversity of Executive Functions and their contributions to complex "Frontal Lobe" tasks: A latent variable analysis. *Cognitive Psychology*, 41(1), 49-100.
- Stafford, T. & Gurney, K.N. (2007). Biologically constrained action selection improves cognitive control in a model of the Stroop task. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 362 (1485), 1671-1684.