C C R C M , N M S , R I b

Christopher H. Chatham[1]*, Eric D. Claus[2], Albert Kim[3], Tim Curran[3], Marie T. Banich[4], Yuko Munakata[3]

1 Department of Cognitive, Linguistic and Psychological Sciences, Brown University, Providence, Rhode Island, United States of America, **2** The Mind Research Network, Albuquerque, New Mexico, United States of America, **3** Department of Psychology and Neuroscience, University of Colorado, Boulder, Colorado, United States of America, **4** Institute of Cognitive Science, University of Colorado, Boulder, Colorado, United States of America

Abstract

The inhibition of unwanted behaviors is considered an effortful and controlled ability. However, inhibition also requires the detection of contexts indicating that old behaviors may be inappropriate – in other words, inhibition requires the ability to monitor context in the service of goals, which we refer to as context-monitoring. Using behavioral, neuroimaging, electrophysiological and computational approaches, we tested whether motoric stopping per se is the cognitively-controlled process supporting response inhibition, or whether context-monitoring may fill this role. Our results demonstrate that inhibition does not require control mechanisms beyond those involved in context-monitoring, and that such control mechanisms are the same regardless of stopping demands. These results challenge dominant accounts of inhibitory control, which posit that motoric stopping is the cognitively-controlled process of response inhibition, and clarify emerging debates on the frontal substrates of response inhibition by replacing the centrality of controlled mechanisms for motoric stopping with context-monitoring.

Citation: Chatham CH, Claus ED, Kim A, Curran T, Banich MT, et al. (2012) Cognitive Control Reflects Context Monitoring, Not Motoric Stopping, in Response Inhibition. PLoS ONE 7(2): e31546. doi:10.1371/journal.pone.0031546

Editor: Sam Gilbert, University College London, United Kingdom

Received April 12, 2011; **Accepted** January 13, 2012; **Published** February 27, 2012

Funding: This study was funded by the National Institute of Mental Health (http://www.nimh.nih.gov) with a P50 center grant (079485-02). The funders had no role in study design, data collection and analysis, decision to publish, or preparation of the manuscript.

Competing Interests: The authors have declared that no competing interests exist.

* E-mail: chathach@gmail.com

Introduction

Inhibition is critical for enabling controlled behavior: bad habits, unfamiliar situations, and dangerous environments often require that default behaviors be stopped and more context-appropriate actions performed [1]. Inhibitory control has been statistically equated with the behavioral and genetic variance common across multiple tests of cognitive and behavioral control [2–3]. Moreover, a particular domain of inhibitory control – response inhibition – has been exempted from the skepticisms surrounding other domains of inhibition [4], and specifically linked to the functioning of a particular frontal region (the right ventrolateral prefrontal cortex; rVLPFC). In this way, the study of response inhibition has supported theorizing that similar mechanisms may enable the inhibition of thoughts and emotions. Thus, modern theorizing is largely consistent with a hypothesis proposed 130 years ago: that ''the centers of inhibition being thus the essential factor of attention, constitute the organic basis of all the higher intellectual faculties'' [5].

However, effective inhibitory control not only requires actually stopping unwanted actions, thoughts, or emotions – it also requires the efficient detection of those contexts that indicate the need for these forms of stopping. To use an example from the domain of response inhibition, one's goal may be to cross a street; this requires actually crossing the street, and stopping these motor actions if oncoming traffic is approaching, but to do so the environment must be monitored so that this motoric stopping can

motor responses on that trial. In contrast, in the ''Double Go Task,'' Signal trials require subjects to repeat their response for that trial as quickly as possible (see methods and Text S1). Thus, both tasks require monitoring for the context that signals what actions should be executed, but only the Stop Task explicitly requires motor actions to be stopped.

The cognitive control required for response inhibition is thought to rely on the prefrontal cortex, to be most crucial at the moment when motoric stopping is required, to be associated with substantial mental effort, to be recruited in a goal-directed fashion, and to support consistent individual differences. We assess each of these characteristics of cognitive control via behavioral, computational, hemodynamic, electrophysiological and pupillometric techniques to determine whether context-monitoring or motoric stopping may reflect the cognitively-controlled process recruited during response inhibition. Convergent evidence of this kind is necessary for making broad claims about the content of cognitive control because cognitive control cannot be unambiguously defined on the basis of any of these characteristics in isolation (e.g., neither prefrontal recruitment nor mental effort alone are sufficient). In addition, this convergent evidence allows us to make multiple points of contact with prior uses of these techniques in the domain of response inhibition, as outlined below.

We used functional magnetic resonance imaging (fMRI) to assess the recruitment of the prefrontal cortex in our tasks. Numerous previous fMRI studies have demonstrated transient activation within the right ventrolateral prefrontal cortex (rVLPFC) and the adjoining anterior insula during trials that require motoric stopping [1–2,6,8–12]. Collectively, this and related evidence has been interpreted to indicate that the rVLPFC is a dedicated substrate for inhibition, and that this function may also be deployed proactively to support behaviors like ''responding with restraint'' [12–13]. Alternatively, it is possible that these hemodynamic patterns reflect the context moni

To foreshadow our results, our results uniformly suggest that, during response inhibition, cognitive control is primarily engaged for the purpose of monitoring the environmental context in the service of goals, rather than for motoric stopping *per se*.

Results

Univariate fMRI Results

First, we found that context-monitoring rather than stopping explained the transient prefrontal contribution to response inhibition. Accounts which posit that motoric stopping is the controlled process during response inhibition tasks predict rVLPFC activation only in the Stop task, but event-related fMRI revealed that the Stop and Double Go tasks activated completely overlapping regions of prefrontal cortex (Fig. 2A), consistent with the tasks' shared context-monitoring demands. Specific regions of interest (ROIs) in the rVLPFC and interconnected subthalamic nucleus (STN) that have been proposed to be specific to the motoric stopping demands were uniformly more strongly recruited on Signal trials in the Double Go Task (Fig. 2B&C; STN: $t(17) = 5.49$, $p<.0001$; BA44: $t(17) = 5.08$, $p<.0001$; BA45: $t(17) = 2.83$, $p = .012$; BA47: $t(17) = 2.5$, $p = .023$), challenging any characterization of these areas as specialized for motoric stopping. A significantly different pattern was observed in areas thought to have a more general attentional role (e.g., the temporo-parietal junction; TPJ [19–20]; $F(1,17) = 31.57$, $p<.0001$), such that both tasks recruited this area equivalently. This equal recruitment of the TPJ across tasks indicates that decreased recruitment of the rVLPFC in the Stop task cannot be explained by globally-decreased activation during that task (e.g., as might result from fatigue; see also discussion in Text S1). Moreover, the increased recruitment of rVLPFC during the Double Go task is consistent with several recent findings, which also demonstrate that tasks involving both context-monitoring and response commission are associated with increased rVLPFC activity relative to tasks involving both context-monitoring and a demand to stop motor actions [7–10] (but see [6] and discussion, below).

Our hybrid fMRI design also allowed us to assess the extent to which neural regions were recruited in a sustained fashion across all trials within the Stop and Double Go tasks. Such sustained activity is potentially a hallmark of proactive context-monitoring processes. Indeed, this analysis revealed sustained hemodynamics in the rVLPFC during both tasks at the timescale of seconds-to-minutes (Fig. 2D), consistent with their shared sustained context-monitoring demands. In contrast, accounts positing that motoric stopping is the cognitively-controlled process during response inhibition predict no sustained rVLPFC activity in the Double Go task, since only response commission is required by that task, and ''responding with restraint'' is unnecessary.

Multivariate Pattern Analysis

We next leveraged multi-voxel pattern analysis to determine whether the sam3(5)1999-1.1y8n-378.d

This pattern is wholly consistent with the idea that similar context-monitoring processes are elicited by Signal trials within both tasks.

In a second multi-voxel pattern analysis, subject-specific classifiers were trained to decode the multivariate patterns that differentiate Double Go_{Signal} and Double $Go_{No\text{-}Signal}$ trials. Classifiers generalized this training on the Double Go task to correctly identify $Stop_{Signal}$ trials with 92–97% accuracy in the rVLPFC, significantly higher than

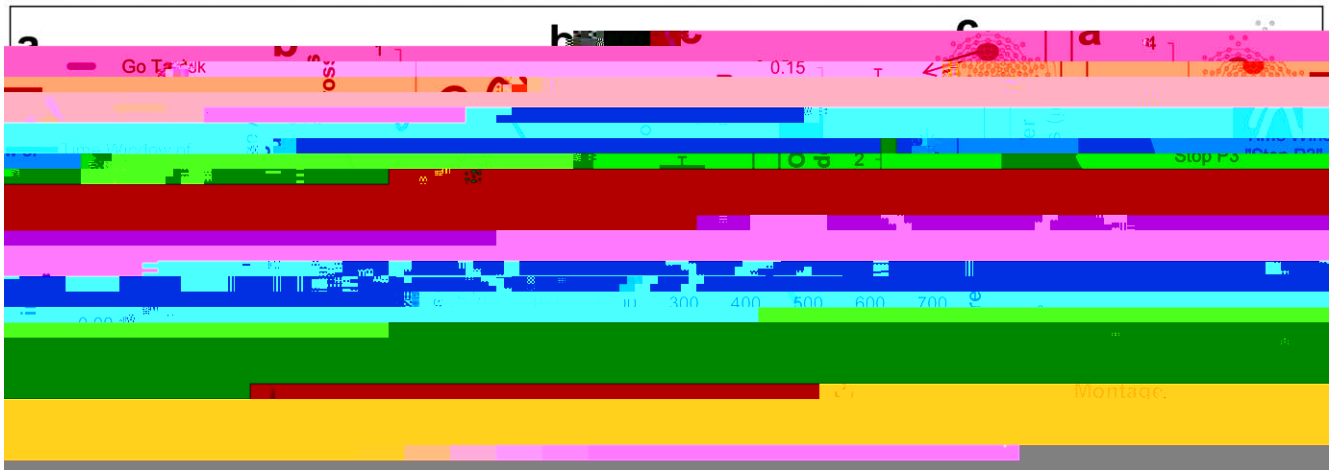

Figure 4. Prefrontal event-related potentials do not strongly distinguish the tasks. A prefrontal positivity peaking around 300 ms, known as the "Stop P3," has been previously associated with stopping, but this component (darkened region of A) was significantly enhanced in the (Double) Go task. Individual differences in voltage were also highly correlated across tasks, indicating substantial overlap in the underlying cortical processes (B). Moreover, prefrontal correlations between the scalp voltage recorded across tasks were disproportionately increased following the presentation of the signal, relative to the increase in occipital correlations observed at the same time (C). This difference indicates increased cross-task similarity in prefrontal processing specifically at signal onset.
doi:10.1371/journal.pone.0031546.g004

appear to engage in reflexive stopping even on the Double Go task, where such stopping runs contrary to instructed goals. Specifically, although Double Go_{Signal} trials require that subjects commit a subset of the motor responses required on Double $Go_{No\text{-}Signal}$ trials, subjects were nonetheless slower to provide even their first response to stimuli when they were followed by the signal than when they appeared alone (Double $Go_{Signal}{}^{1st\ RT}$ > Double $Go_{NoSignal}{}^{Only\ RT}$; $t(148) = 9.59$, $p < .0005$; Fig. 6A). To the extent that this behavioral slowing in the Double Go task reflects some transient stopping, it runs contrary to subjects' goals in the Double Go task and therefore might not be engaged in a controlled or goal-directed manner.

On the other hand, the presence of goal-inconsistent slowing during the Double Go task does not by itself refute the idea that motoric stopping can be a controlled process in this task or in others. Indeed, one alternative interpretation of this slowing is that it does in fact reflect a controlled and goal-directed process: it may be an attempt to stop or replace the motor plan required on Double $Go_{No\text{-}}$

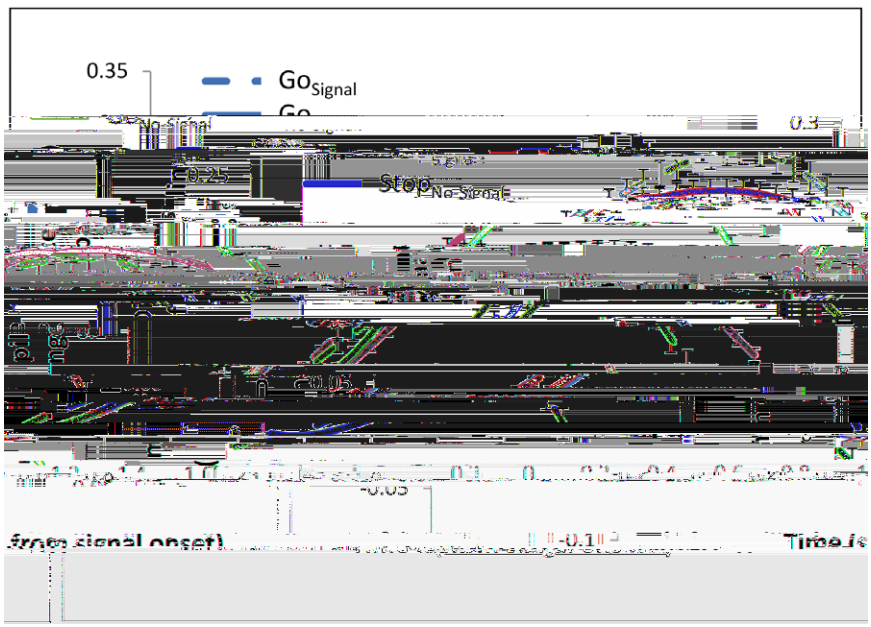

Figure 5. Patterns of mental effort assessed via pupillometry indicate that effort matches demands on context-monitoring, not stopping, and is modulated by the relevance of the infrequent stimulus to the planned response. In particular, stopping a response ($Stop_{Signal}$ trials) was associated with more mental effort was required by monitoring for the appearance of stimuli that would demand stopping ($Stop_{No\text{-}Signal}$ trials) than by stopping itself ($Stop_{Signal}$ trials) or by monitoring for the appearance of stimuli that would demand an additional act of going ($Go_{No\text{-}Signal}$ trials).
doi:10.1371/journal.pone.0031546.g005

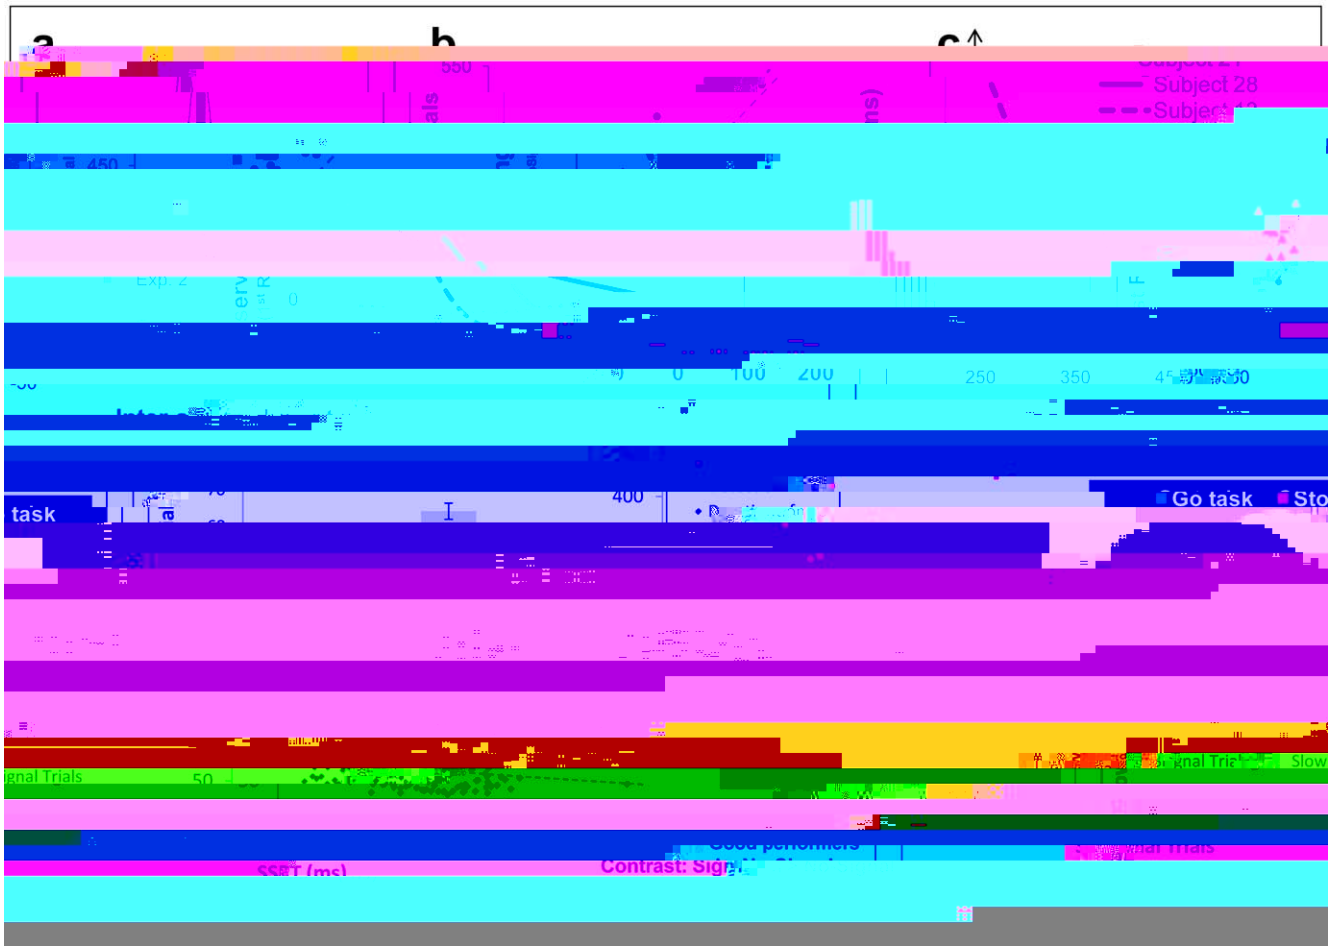

Figure 6. Mixture model analyses separate slowed from unslowed trials in the Go task, and demonstrate this slowing is not the source of the commonality across tasks.

Signal trials (i.e. the motor plan for ''respond once'' is stopped or replaced with the motor plan for ''respond twice''). We assessed this possibility with a model-based decomposition of subjects' behavior; however, the results of this analysis argue against this possibility, and further show that the efficiency of subjects' context-monitoring, rather than the efficiency of motoric stopping or motor plan replacement, shares a closer relationship with SSRT.

To assess the alternative accounts, we developed a formal model of context-monitoring and stopping by building on the classic race model of the Stop task [7] (see also Text S1 and Figure S7A) in order to precisely estimate the duration of motoric slowing experienced by subjects in the Double Go task, as well as exactly which trials underwent such slowing (Figure S7B). The race model of the Stop task posits that responses undergo inhibition when a stopping process, triggered by the onset of the Stop signal, completes before the ''going'' processes triggered by the onset of the 2AFC stimulus. The race between stopping and going processes is the model's namesake, and is supported by the monotonically-decreasing relationship of interstimulus interval (ISI) to successful inhibition: larger ISIs give the ''going'' process an increasing advantage in the race, and thus leads to less successful inhibition. We observed a similar phenomenon in our Double Go task, such that increasing ISIs led to less slowing of first responses; this effect was visible at the group level (Figure S7C) but also even at the level of individual subjects (Fig. 6B), who showed substantial variability in the earliest ISI to yield zero observable slowing.

We utilized this behavioral variability to estimate individual differences in Double Go task performance. First, we estimated the

probability that each trial belonged to either the ''slowed'' or ''unslowed'' distributions of reaction times. This categorization was accomplished by fitting a mixture model to the difference between reaction times of Double $Go_{Signal}^{1st\ RT}$ and Double $Go_{NoSignal}^{Only\ RT}$ trials of corresponding percent rank. To the extent these reaction times come from the same (i.e., unslowed) distribution, these equipercentile residuals should be centered on zero; however, there was pronounced positive skew (Fig. 6C), indicating that a substantial proportion of trials did undergo slowing. We considered as ''slowed'' those trials that were marginally less likely to come from a Gaussian distribution centered on zero, relative to an alternative distribution with a positive mean (see overlaid curves on Fig. 6C, and Text S1). This method clearly separated ''slowed'' from ''unslowed'' trials on the basis of the first RT on Double Go_{Signal} trials: ''unslowed'' trials showed approximately zero slowing relative to corresponding trials within the No Signal distribution, whereas ''slowed'' trials were significantly longer than corresponding trials within the No Signal distribution (Fig. 6D).

Next, we estimated for each subject the amount of time that must elapse after signal presentation until responses are categorized as ''slowed'' (yielding the time of signal detection [TOSD], our measure of context-monitoring), and the difference between that subjects' ''slowed'' and ''unslowed'' reaction times (yielding the duration of slowing [DoS], our measure of stopping from the Double Go task). If motoric stopping (or, equivalently, motor plan replacement) is controlled, and initiated in this controlled fashion in the Double Go task, then the process of motoric stopping or motor plan replacement should should cease (as estimated by DoS, in the Double Go task) in proportion to how quickly competing motor plans can be stopped, as assessed by SSRT in the Stop task. That is, the ''controlled motoric stopping'' and ''controlled motor plan replacement'' accounts both predict that DoS and SSRT should be positively correlated.

However, DoS and SSRT were not positively correlated – they instead showed a weak negative correlation (Pearson $R = -.188$, $p = .048$; Fig. 6E), in direct contradiction to the prediction motivated by these alternative accounts. SSRT was instead positively correlated only with TOSD – i.e., the efficiency with which signals could be detected (Fig. 6E; $R = .418$, $p < .0005$) – as predicted by accounts which posit that context-monitoring underlies the commonalities of the Double Go and Stop Signal tasks. This positive relationship persisted when controlling for DoS ($R = .410$, $p < .0005$), indicating that the overlapping variance in TOSD and SSRT does not reflect motoric stopping or motor plan replacement. Strikingly, this relationship of context-monitoring to SSRT was also regionally-specific: SSRT and TOSD overlapped in their relationship to hemodynamics only within the rVLPFC (Fig. 6F).

A second, independent assessment of the origin of the observed commonalities across our tasks is also enabled by our formal model. Specifically, the model identifies exactly which trials undergo motoric stopping/slowing within the Double Go task, and thus permits these trials to be excluded from analysis. To the extent that similar hemodynamic, electroencephalographic, and pupillometric patterns are observed when these ''slowed'' trials are excluded, it would suggest that the commonalities across our tasks do not reflect a motoric stopping process common to these tasks.

Consistent with the claim that a common and cognitively-controlled process of context-monitoring – and not a common process of motoric stopping – underlies the commonalities of our tasks, a complete re-analysis of the data without such ''slowed'' trials replicated all of our primary results: the increased transient hemodynamic response in the rVLPFC during the Double Go task, the prominent sustained hemodynamic activity observed in that task, the multivariate hemodynamic commonalities across tasks, the increased Stop P3 response in the Double Go task, the

hemodynamics [8–9,11,19–21] (but see [10] for an exception). By

analysis). Similarly, in the Stop Task, the signal turned red if subjects failed to successfully stop their response on that trial (in all experiments). Additional cross-experiment differences in our tasks suggest the generality of our results across minor variations in experimental procedure (see Figure S1 & Table S1).

Statistical Analysis of fMRI

Data were acquired with a 3T GE Signal whole-body MRI scanner at the University of Colorado Health Sciences Center, using T2-weighted echo-planar imaging (EPI) (TR = 2000 ms, TE = 32 ms, flip angle = 70°). Additional acquisition details are available in Text S1.

Image pre-processing and analyses were conducted with FSL (FMRIB's Software Library). The first six volumes of each run were discarded to allow the MR signal to reach steady state, the remaining images in each participant's time series were motion corrected using MCFLIRT, and non-brain voxels were removed

calculated as:

$$BIC = -2 \cdot \sum_{n=1}^{N} \ln\left(\sum_{d=1}^{D} \Phi_d \, L_d RT_n\right) + D_p \cdot \ln(N)$$

Where N is the total number of observations, D is the total number of distributions fit, D_p is the total number of free parameters used in fitting those distributions, Φ_d is the weight of the d^{th} distribution, and $L_d(RT_n)$ is the likelihood of the n^{th} RT given the best fit parameters for the d^{th} distribution (μ and σ for Gaussian and k and Θ for Gamma).

We next categorized individual trials as slowed or unslowed using the likelihood of observing each RT under either of the two fitted distributions. RTs were categorized as slowed if there was even weak evidence in favor of the RT belonging to that distribution (as quantified by a difference in BIC of $\geq$2.35); otherwise RTs were categorized as unslowed. Other standards of evidence lead to similar results as those presented here, but do not as cleanly separate the slowed and unslowed trials (c.f. Fig. 6D).

To calculate TOSD, we subtracted the signal delay from the n^{th} percentile of no signal trial RTs, where n corresponds to the

as the difference between slowed 1st responses on Double Go_{Signal}trials and responses of the same percent rank on Double $Go_{No\text{-}Signal}$ trials. The time of signal detection can be estimated as the amount of time that must elapse following a signal before responses are slowed. (C) The process model of the Double Go Task predicts that slowing should be larger when signals are presented earlier; this prediction was confirmed.
(TIF)

Table S1 Differences between Experiments 1–3.
(DOCX)

Table S2 Mixture model estimates.
(DOCX)

Table S3 Descriptive statistics for model-based analyses across Experiments.
(DOCX)

Table S4 Statistical results from all primary analyses both with and without those trials designated by the mixture model as ‘‘slowed’’ in the Go task.

43. Aisa B, Mingus B, O'Reilly RC (2008) The emergent neural modeling system.