



UvA-DARE (Digital Academic Repository)

A unified account of simple and response-selective inhibition

Gronau, Q.F.; Hinder, M.R.; Salomoni, S.E.; Matzke, D.; Heathcote, A.

DOI

[10.31234/osf.io/aqk8d](https://doi.org/10.31234/osf.io/aqk8d)

[10.1016/j.cogpsych.2023.101628](https://doi.org/10.1016/j.cogpsych.2023.101628)

Publication date

2024

Document Version

Final published version

Published in

Cognitive Psychology

License

Article 25fa Dutch Copyright Act (<https://www.openaccess.nl/en/policies/open-access-in-dutch-copyright-law-taverne-amendment>)

[Link to publication](#)

Citation for published version (APA):

Gronau, Q. F., Hinder, M. R., Salomoni, S. E., Matzke, D., & Heathcote, A. (2024). A unified account of simple and response-selective inhibition. *Cognitive Psychology*, 149, Article 101628. <https://doi.org/10.31234/osf.io/aqk8d>, <https://doi.org/10.1016/j.cogpsych.2023.101628>

General rights

It is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), other than for strictly personal, individual use, unless the work is under an open content license (like Creative Commons).

Disclaimer/Complaints regulations

If you believe that digital publication of certain material infringes any of your rights or (privacy) interests, please let the Library know, stating your reasons. In case of a legitimate complaint, the Library will make the material inaccessible and/or remove it from the website. Please Ask the Library: <https://uba.uva.nl/en/contact>, or a letter to: Library of the University of Amsterdam, Secretariat, P.O. Box 19185, 1000 GD Amsterdam, The Netherlands. You will be contacted as soon as possible.

UvA-DARE is a service provided by the library of the University of Amsterdam (<https://dare.uva.nl>)



A unified account of simple and response-selective inhibition[☆]

Quentin F. Gronau^{a,*}, Mark R. Hinder^{b,1}, Sauro E. Salomoni^b, Dora Matzke^c, Andrew Heathcote^c

^a University of Newcastle, Australia

^b University of Tasmania, Australia

^c University of Amsterdam, The Netherlands

ARTICLE INFO

Dataset link: <https://osf.io/5z9vk/>

Keywords:

Selective stopping

ABSTRACT

Response inhibition is a key attribute of human executive control. Standard stop-signal tasks require countermanding a single response; the speed at which that response can be inhibited indexes the efficacy of the inhibitory control networks. However, more complex stopping tasks, where one or more components of a multi-component action are cancelled (i.e., response-selective stopping) cannot be explained by the independent-race model appropriate for the simple task (Logan and Cowan 1984). Healthy human participants ($n = 28$; 10 male; 19–40 years) completed a response-selective stopping task where a ‘go’ stimulus required simultaneous (bimanual) button presses in response to left and right pointing green arrows. On a subset of trials (30%) one, or both, arrows turned red (constituting the stop signal) requiring that only the button-press(es) associated with red arrows be cancelled. Electromyographic recordings from both index fingers (first dorsal interosseous) permitted the assessment of both voluntary motor responses that resulted in overt button presses, and activity that was cancelled prior to an overt response (i.e., partial, or covert, responses). We propose a simultaneously inhibit and start (SIS) model that extends the independent race model and provides a highly accurate account of response-selective stopping data. Together with fine-grained EMG analysis, our model-based analysis offers converging evidence that the selective-stop signal simultaneously triggers a process that stops the bimanual response and triggers a new unimanual response corresponding to the green arrow. Our results require a reconceptualisation of response-selective stopping and offer a tractable framework for assessing such tasks in healthy and patient populations.

Significance Statement

Response inhibition is a key attribute of human executive control, frequently investigated using the stop-signal task. After initiating a motor response to a go signal, a stop signal occasionally appears at a delay, requiring cancellation of the response. This has been conceptualised as a ‘race’ between the go and stop processes, with the successful (or failed) cancellation determined by which process wins the race. Here we provide a novel computational model for a complex variation of the stop-signal task, where only one component of a multicomponent action needs to be cancelled. We provide compelling muscle activation data that support our model, providing a robust and plausible framework for studying these complex inhibition tasks in both healthy and pathological cohorts.

[☆] This work was supported by the Australian Research Council (DP200101696). DM is supported by a Vidi grant (VI.Vidi.191.091) from the Dutch Research Council (NWO).

* Corresponding author.

E-mail addresses: quentin.f.gronau@gmail.com (Q.F. Gronau), mark.hinder@utas.edu.au (M.R. Hinder), sauro.salomoni@utas.edu.au (S.E. Salomoni), d.matzke@uva.nl (D. Matzke), ajheathcote@gmail.com (A. Heathcote).

¹ These authors contributed equally.

<https://doi.org/10.1016/j.cogpsych.2023.101628>

Received 4 May 2023; Received in revised form 17 October 2023; Accepted 10 December 2023

Available online 9 January 2024

0010-0285/© 2023 Published by Elsevier Inc.

1. Introduction

Response inhibition is a key feature of executive function most commonly assessed using the stop-signal task (Verbruggen et al., 2019). The standard stop-signal task requires rapid choice responses (e.g., a left or right button press) to a “go” stimulus (e.g., a left or right pointing arrow), that on some trials must be cancelled when a second stimulus occurs some time later (the stop-signal delay, SSD; Table 1 lists the acronyms used in this article). Assuming a simple independent race between processes or “runners” with variable speeds triggered by the go and stop signals (Logan & Cowan, 1984), the standard task affords a key index of executive control, the running time of the inhibitory process. Go response time (RT) is much slower than the running time of the inhibitory process, which is called stop-signal reaction time (SSRT), consistent with stopping being mediated by a specialised fronto-basal ganglia network (Sebastian, Forstmann, & Matzke, 2018).

Although the standard task is widely used (Matzke, Verbruggen, & Logan, 2018), more complex response-selective inhibition tasks—where only some components of a multi-component action are cancelled, with the remaining components being executed as quickly as possible—may be more relevant to inhibition in the real-world (Aron & Verbruggen, 2008), and to executive-function disorders (Aron, 2011). A common example occurs when initiating a manual gear change, braking with one foot and pressing the clutch with the other. When a dangerous change in road conditions occurs, clutching is aborted while still braking. In the laboratory analogue, participants simultaneously press a button with each hand in response to the go signal (e.g., a pair of left and right pointing green arrows). If a selective stop signal occurs (e.g., one arrow turns red) the corresponding button press is withheld but the other still made.

Here we show that the traditional explanation of selective stopping, as instantiated in the Activation-Threshold Model (ATM, see Fig. 1A: MacDonald, Coxon, Stinear, & Byblow, 2014; MacDonald, McMorland, Stinear, Coxon, & Byblow, 2017), is inconsistent with electromyographic (EMG) activity in a task mixing two types of stopping trials, selective stopping and standard stopping (i.e., both arrows turn red and no response should be produced). We then propose a new model, called simultaneously inhibit and start (SIS, see Fig. 1B), which provides a unified explanation of both standard selective stopping trials as well as go trials (i.e., trials with no stop signal). We show that SIS accurately characterises both EMG and all aspects of behaviour, including the probability of stopping and the full distribution of RT in any case where responses do occur.

The traditional approach assumes a global inhibitory process is triggered by both standard and selective stop trials, with a subsequent “restart” process required to produce a unimanual response on selective-stop trials. The need to globally stop and then subsequently restart has been used to explain why, when selective stopping is successful, the unimanual response is delayed relative to the bimanual response on trials without a stop signal, typically by 100–200 ms (Coxon, Stinear, & Byblow, 2007). The ATM explains the common finding that unimanual EMG amplitudes are larger than bimanual EMG amplitudes by the restart needing to exceed a higher activation threshold caused by the global inhibition. More specifically, the ATM was developed based on variations in cortico-motor excitability observed during unimanual stop trials that were taken to suggest a series of three distinct processing stages, (i) a go stage: the go signal triggers preparation and activation of a synchronous two-component (“bimanual”) response via neuronal coupling of effector representations; (ii) a stopping stage: if the stop signal arrives before EMG amplitude is sufficient to cause execution of a bimanual response, the threshold activation required to do so is raised; (iii) a restart stage: a new one-component (unimanual) response is then prepared and activation boosted sufficiently to overcome the higher threshold and trigger a unimanual response. Simulations performed by MacDonald et al. (2017) show that the delay between the onset of the second stage, which causes the activation threshold to increase exponentially towards a higher level, and third stage, which causes activation to increase exponentially towards an even higher level, is usually quite substantial. Their Fig. 5F shows the delay is at the very least 60 milliseconds (ms) and usually much longer.

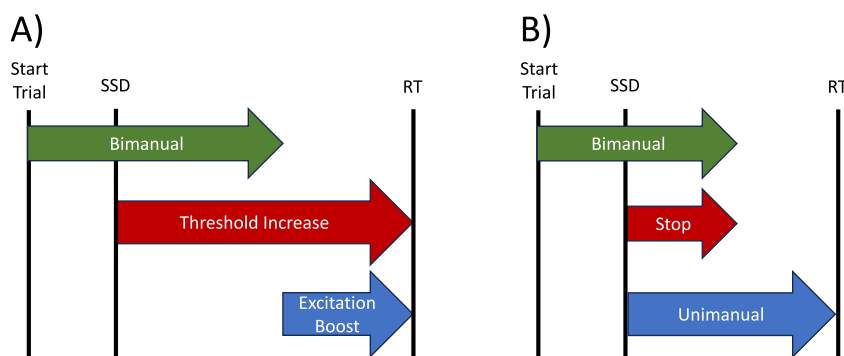


Fig. 1. Models of a successful selective-stop trial. (A) The Activation-Threshold Model (ATM): Presentation of go stimulus starts a unitary bimanual-response production process (green arrow), then at a time SSD later, before the bimanual process has reached the response thresholds a selective stop signal causes a threshold increase (red arrow), and after a further delay an excitation boost is added to response activation (blue arrow) so that eventually the threshold is reached at time RT and a unimanual response is produced. (B) The simultaneous inhibit and start (SIS) model: Presentation of go stimulus starts a unitary bimanual-response runner (green arrow), then at SSD the stop signal triggers both a stop runner (red arrow) and a unimanual response runner. The stop runner catches and inhibits the bimanual runner and some time later the unimanual runner triggers a corresponding response. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

Table 1
Acronyms and their meanings.

Type	Acronym	Meaning
Stimulus	DS	dual stimulus (no following stimuli)
	SS	stop stimulus (DS followed by bimanual stop signal)
	LS	left stimulus (DS followed by unimanual stop-right signal)
	RS	right stimulus (DS followed by unimanual stop-left signal)
Response	DR	dual response, bimanual (correct action for DS)
	NR	non-response (correct action for SS)
	LR	left response, unimanual (correct action for LS)
	RR	right response, unimanual (correct action for RS)
Measure	EMG	electromyographic activity
	RT	response time (onset of DS to button press)
	SSD	stop-signal delay (onset of DS to stop signal)
	SSRT	stop-signal RT (running time for the stop process)
Theory	ATM	Activation-Threshold Model
	SIS	Simultaneously Inhibit and Start

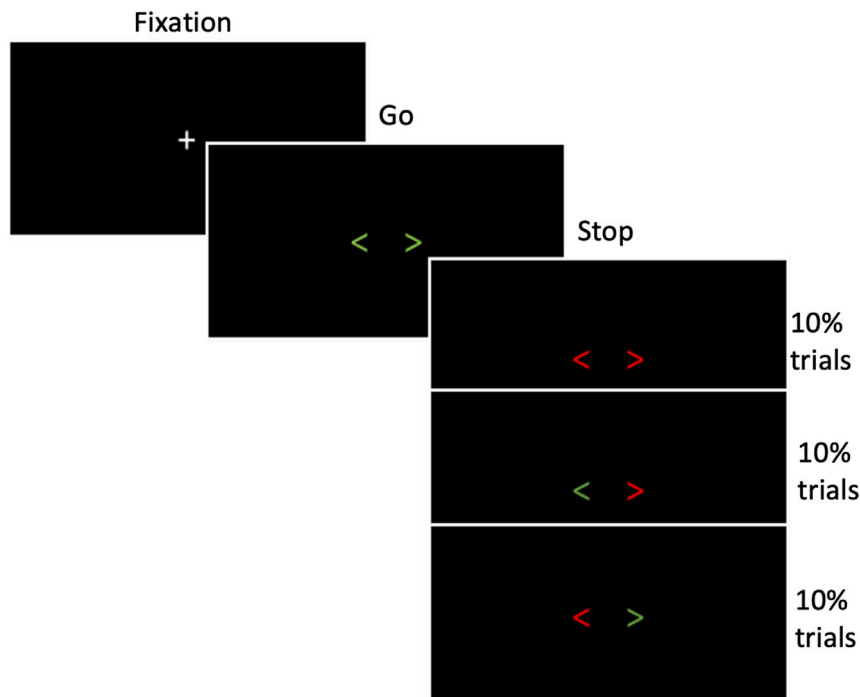


Fig. 2. The presentation of the stop-signal task visual stimuli in the current study. A centrally located fixation cross was presented for between 500–700 ms. At the offset of the fixation cross, the go stimulus would then appear for 500 ms. On 30% of trials, the imperative go signal changed to the stop signal at SSD ms after presentation of the go stimulus. The 30% of stop trials were comprised of 10% of each stop type (bimanual stop, left stop, right stop).

SIS instead explains selective stopping by an extension of the race-model architecture applied to the standard stop-signal task, where the selective-stop signal simultaneously inhibits the bimanual response and starts a new unimanual response. SIS and ATM differ in three key ways, (i) rather than the second and third stages occurring in series, they occur in parallel, (ii) rather than the third stage being a restart, it is constituted of an entirely new selective-response process; (iii) rather than being global, the action of the stopping process is specific to the bimanual runner. If the same type of inhibitory mechanism is used in the standard and selective stop-signal tasks, SIS must assume its stopping process will be much faster than any go process. This predicts that successful stopping will typically rely on the stop runner being quick enough to overtake the bimanual runner. The selective process then triggers a corresponding response some time later. Unimanual responses are delayed relative to the bimanual responses simply because the onset of the selective process is delayed by the SSD. This predicts that the unimanual vs. manual delay is a function of the SSD, although it could be shorter than the SSD if the unimanual runner is faster than the bimanual runner as indicated by EMG amplitudes. However, under the unified account proposed by SIS, the go runners should rely on the same neural mechanisms, and so the unimanual runner should not be much faster. Similarly, if standard and selective stopping rely on the same neural mechanisms, the stop runner should be much faster than either type of go runner in selective stopping.

Importantly the two theories make clearly contrasting predictions about the time course of “partial” responses—EMG bursts insufficient to cause a response—that sometimes occur in both hands on successful selective-stop trials which later have a single response-generating burst in one hand. ATM predicts a gap between partial responses and the unimanual response due to the need to restart. A partial response occurs when EMG activity increases as activation approaches but does not reach the response threshold before the stopping process elevates the threshold, causing EMG activity to decline before it is sufficient to trigger a response. There is then necessarily a delay before the restart process can boost activation sufficiently to cross the new higher threshold. SIS, in contrast, predicts that partial bursts can occur any time up to just before the unimanual response depending on the margin between the finishing times of inhibitory, bimanual, and unimanual runners. That is, if the stop runner arrives before a sufficient amplitude is generated to trigger a bimanual response and causes EMG in both hands to start to decrease it is still possible that a fast unimanual runner might arrive very soon after, causing EMG to increase in the corresponding hand while it continues to decrease in the other hand.

If running times are as variable in selective stopping as they are known to be in the standard paradigm (Matzke et al., 2013), such “blending” of partial and response-generating bursts should be sufficiently common to be detectable. To enhance measurement of partial bursts and our ability to detect any gap we used both stiff and conventional compliant response buttons in our experiment. To foreshadow our results, for both button types, our EMG analysis found no evidence for a gap, leading us to propose the SIS model. The simple structure of SIS allowed us to implement it as a quantitative cognitive model and extend the Bayesian hierarchical methods that have been successfully used to estimate the distribution of SSRT in the standard paradigm (Matzke, Love, & Heathcote, 2017; Matzke et al., 2013) to estimate both SSRT and unimanual runner finishing time distributions in selective stopping. We simultaneously fit the SIS model to response probabilities and the distributions of RTs for bimanual and unimanual responses to determine whether it provides an accurate quantitative account of the data. We also test the predictions that the unimanual runner is faster than the bimanual runner but the stop runner is faster than both.

2. Materials and methods

2.1. Experimental design

2.1.1. Participants

Twenty-eight participants (ten male) aged between 19 and 40 years ($M = 24.1$, $SD = 4.6$) undertook a single experimental session after giving informed consent; the procedures had received ethical approval from the Tasmanian Human Research Ethics Committee. Twenty-six participants were self-reportedly right-handed, whereas the other two were left-handed. Participants were screened for neurological impairments, muscular conditions that limited their ability to comfortably undertake the task, and colour blindness (Ishihara, 1972).

2.1.2. Procedure

Participants were seated at a desk approximately 80 cm away from a computer monitor with forearms pronated and resting comfortably on the desk. Each index finger rested on a button, with the buttons positioned approximately shoulder width apart. Following initial on-screen instructions (together with verbal explanations to provide any clarification), participants were required to respond to visual stimuli created in PsychoPy (Peirce, 2009) presented on a black background by pressing the buttons. They completed 1260 trials (60 practice trials and 1200 experimental trials) in a single experimental session in the laboratory, which lasted no more than two hours. Breaks were administered between each 120-trial block to prevent fatigue, with a longer break after 630 trials, if desired.

We manipulated the force required to register a response to the visual stimuli by using two sets of response buttons—either stiff or compliant, which respectively required 7.4 or 1.0 Newton of force to generate a response. The buttons were swapped half way during the experiment (the order of button use was counterbalanced across participants), but the procedure with each button type was otherwise identical.

On each trial, a fixation cross was presented for 500–700 ms (uniformly distributed) in the centre of the screen at the start of each trial; the jittering of presentation period was designed to prevent temporal prediction of the upcoming go stimulus. Following the fixation cross, two green arrows appeared (the ‘go’ stimulus), requiring fast-as-possible simultaneous responses on both left and right buttons. Successful responses, defined as a bimanual response occurring in a window of 150–1500 ms after the go stimulus, were followed by visual feedback of the response time for 500 ms prior to commencement of the next trial. Failed responses received the feedback message ‘incorrect’, and button presses after the response window received the feedback message ‘missed’. Any responses faster than 150 ms were met with the feedback message ‘too fast’.

On 30% of trials, one or both green arrows turned red, indicating that the participant must try to cancel the response on the same side as the red arrow. Thus, some stop trials required both left and right responses to be cancelled (both arrows turn red: ‘bimanual stop’), whereas others required only one response to be cancelled while the other response continued (either the left or the right arrow turns red: left or right stop—‘selective stops’). Each of the three stop types was equally likely (10% of all trials) and presented in a pseudo-randomised order. The delay between the onset of the ‘go’ and ‘stop’ stimuli (stop-signal delay, SSD), was staircased in 50 ms increments based on whether or not participants successfully cancelled their response on the previous stop trial: Successful inhibition resulted in longer SSD on the subsequent stop trial, whereas failed inhibition resulted in shorter SSD on the subsequent trial. SSD was incremented independently for each stop type, and re-set to 150 ms at the start of each 120-trial block.

Successfully inhibited trials were followed by the message ‘good’ presented on the screen, whereas failed inhibitions were provided with the feedback message ‘incorrect’.

Each participant undertook a total of 630 trials with each button type: 30 practice trials with three stop trials of each variant, followed by 5 blocks of 120 trials, each constituted of 90 go trials, 10 bimanual stop, 10 left stop and 10 right stop trials. Practice trials were not used in the analyses.

2.1.3. Electromyographic (EMG) recordings

EMG signals were recorded using adhesive electrodes (Ag/AgCl) positioned in a belly-tendon montage on the first dorsal interosseous (FDI) muscle of each hand, with a ground electrode positioned on the ulnar bone on each wrist. The analogue signals were band-pass filtered at 20–1000 Hz, amplified 1000 times, sampled at 2000 Hz (CED Power 1401 and CED 1902, Cambridge, UK) and saved into a PC for offline analysis. Prior to commencing the task, participants were provided with initial visual feedback of the EMG signal (biofeedback) to ensure they completely relaxed their muscles. If signals became noisy during the experiment (i.e., unintended activation of the muscles, not related prior to the presentation of stimuli), the experimenter instructed the participant to relax their hands.

2.2. Statistical analysis

2.2.1. Behavioural data

Raw behavioural data, and a script in the R language that prepares it for further analysis (prepareData.R) are provided in supplementary materials at <https://osf.io/5z9vk/>. Consistent with the standard assumption that neuronal coupling of effector representation typically produces bimanual responses, RTs for left and right button presses were typically very close to synchronous. Coupling was reflected in the standard deviation of right–left RT (38 ms) being much less than predicted assuming independence (median 180 ms, bootstrap 95% confidence interval 179–182 ms). However, occasionally, bimanual responses were asynchronous with difference intervals ranging up to 1000 ms. Based on inspection of the histogram of differences (see prepareData.R at <https://osf.io/5z9vk/>) we removed all bimanual responses with asynchrony greater than 60 ms (1.9% of trials overall, 0.1% to 5.8% on an individual basis) from further analysis. We treated the remaining differences due to motor noise by averaging the left and right button press times to produce a single RT measure for bimanual response trials. We also removed trials with RTs less than 200 ms (0.13%), as they are very likely to be anticipatory responses which are too fast to have been triggered by the go stimulus.

For all trial types, participants occasionally made unusual errors. On dual-stimulus (DS) trials (i.e., trials without a stop signal), they constituted either no response or a unimanual left or right response (0.7%, 0.5% or 0.4% of this trial type). On dual stop (SS) trials (i.e., trials with two red arrows), they are unimanual left or right responses (2% each). On unimanual left-stimulus (LS) trials (i.e., red arrow on the right), they are unimanual right responses or no response (1.2% and 0.3%) and on unimanual right-stimulus (RS) trials, they are unimanual left responses or no response (0.9% and 0.5%).

Here we report results for the basic SIS model that does not account for these unusual responses, and so they were removed from the modelling analyses (1.4% of trials overall, 0.2% to 5.1% on an individual basis). We also developed an extended model that does account for these rare unusual responses in a descriptive manner on the assumption that they are produced by motor errors. We do not make any strong claims about the motor-error account given the rarity of the associated responses makes inference about them very uncertain. However, in supplementary materials we provide detailed results for the extended model to demonstrate they are consistent with the results for the basic model reported here, showing that omitting the unusual responses has no material impact on substantive conclusions in this case. On the OSF website, we provide R code to fit both basic and extended models using hierarchical Bayesian methods implemented through the Dynamic Models of Choice software (DMC; Heathcote et al., 2019). Supplementary materials also provide a formal specification of the basic and extended models.

We fit the models in a Bayesian hierarchical manner (Matzke, Dolan, Batchelder, & Wagenmakers, 2015; Shiffrin, Lee, Kim, & Wagenmakers, 2008) using the default settings in the DMC software, which implements differential evolution Markov chain Monte Carlo sampling (Turner, Sederberg, Brown, & Steyvers, 2013) to obtain samples from the posterior distribution. Specifically, we started by fitting each participant’s data separately. These individual estimates were used as the start points for the hierarchical fitting. Then we ran the hierarchical fitting in an automatic manner until the chains had converged to the posterior using the DMC function `h.run.converge.dmc`. Finally, we ran the chains for 250 more iterations using the DMC function `h.run.dmc`. These final iterations were used for all analyses. Convergence of the MCMC chains was assessed by visual inspection and the \hat{R} statistic (Brooks & Gelman, 1998) which was below 1.1 for all parameters.

The prior distributions that we used for the group-level parameters are displayed in Table 2 along with their definitions. The times for bimanual and unimanual runners to trigger a response, and for the stop runner to inhibit the go runner, are described by separate ex-Gaussian distributions, which are the convolution of a normal distribution (with mean μ and standard deviation σ) and an exponential distribution (with mean τ). The ex-Gaussian distribution, is positively skewed, with mean $\mu + \tau$, variance $\sigma^2 + \tau^2$, and third central moment given by $2 \times \tau^3$ relating to the skewness of the distribution. Differences between responses made with stiff and compliant buttons were accounted for by allowing different normal means for both bimanual and unimanual runners. With probability p_{TF} the model also allowed for the possibility that the stop signal may fail to trigger the stop runner (Matzke, Love et al., 2017), indicating a lapse of attention to task goals.

Table 2

The prior distributions for the group-level mean parameters were all truncated normal distributions with the below specified location, scale, lower bound, and upper bound. The priors on the group-level standard deviations were exponential distributions with the rate parameter set to 1. Note that the prior for the trigger failure parameter P_{TF} was specified on the probit scale (Φ^{-1} denotes the probit transformation).

Type	Parameter	Location	Scale	Lower bound	Upper bound
Bimanual normal mean	$\mu^{\text{compliant}}$	0.4	1	0	2
Bimanual normal mean	μ^{stiff}	0.4	1	0	2
Bimanual normal SD	σ	0.05	1	0	0.5
Bimanual exponential mean	τ	0.15	1	0	0.5
Stop normal mean	$\mu_{\text{stop}}^{\text{compliant}}$	0.2	1	0	2
Stop normal mean	$\mu_{\text{stop}}^{\text{stiff}}$	0.2	1	0	2
Stop normal SD	σ_{stop}	0.025	1	0	0.5
Stop exponential mean	τ_{stop}	0.075	1	0	0.5
Unimanual normal mean	$\mu_U^{\text{compliant}}$	0.4	1	0	2
Unimanual normal mean	μ_U^{stiff}	0.4	1	0	2
Unimanual normal SD	σ_U	0.05	1	0	0.5
Unimanual exponential mean	τ_U	0.15	1	0	0.5
Probability of trigger failure	P_{TF}	$\Phi^{-1}(0.1)$	1	-6	6

2.2.2. EMG analysis

The EMG signals were digitally filtered using a fourth-order band-pass Butterworth filter at 20–500 Hz. EMG envelopes were obtained by full-wave rectification and low-pass filtering at 10 Hz, and used to extract all EMG parameters, except for onset and offset times.

We used a single-threshold algorithm (Hodges & Bui, 1996) to detect the precise times of onset and offset of EMG bursts. Signals were rectified and smoothed by low-pass filtering at 50 Hz. A moving average with a window of 500 ms was used to find the segment with the lowest RMS amplitude (baseline), and EMG bursts were identified when the amplitude was above 3 SD from baseline. For robustness, EMG bursts separated by less than 20 ms were merged together.

From the EMG onsets and offsets detected, we set time constraints to identify up to two bursts of interest: First, the RT-generating burst was identified as the last burst where $\text{GO-signal} < \text{onset} < [\text{RT} - 100 \text{ ms}]$. In addition, some trials exhibited partial bursts, i.e., EMG bursts that were not strong enough to generate a button press. In successful stop trials, partial bursts represent motor commands that, although initiated, were cancelled in response to the stop signal. Based on theoretical considerations that match our model assumptions, partial bursts were identified in each hand as the earliest burst where (i) EMG offset $>$ SSD (i.e., the initial bimanual response is cancelled after presentation of the stop signal); (ii) time of peak EMG $<$ onset of RT-generating burst from the responding hand (i.e., the EMG associated with the new unimanual response must commence [i.e., an onset is detected] after the EMG associated with the initial bimanual/partial response begins to decrease [i.e., after its peak is detected]); and (iii) peak EMG amplitude $>$ 20% of the average peak from successful bimanual go trials (avoid inclusion of spurious muscle activity unrelated to the task). In bimanual stop trials, in the absence of an RT-generating burst, we estimated an upper time constraint for the time of peak partial burst, i.e., condition (ii), that would be comparable to the one used in selective-stop trials: The 90th percentile from the distribution of EMG onsets of the RT-generating burst from successful selective-stop trials. These time and amplitude constraints were critical to avoid inclusion of spurious bursts, in particular those initiated together with or after the RT-generating burst, which are likely related to mirror activity or other spurious muscle activity unrelated to the stop-signal task.

From each burst detected (RT-generating and partial), we extracted the times of onset, peak and offset, as well as the peak amplitude and peak slope (also called rate of EMG onset) from the EMG envelopes. The amplitude of EMG envelopes from each hand was normalised by the average peak EMG from successful bimanual go trials, independently for stiff and compliant buttons. We used generalised linear mixed models (GLMMs) to compare the parameters extracted across trial types.

3. Results

We denote the four trial types occurring in the experiment in terms of the stimulus that defines the behaviour desired of the participant: (1) dual stimulus (DS, two green arrows), requiring a bimanual go response; (2) stop stimulus (SS, two red arrows), requiring both left and right responses be withheld, i.e., bimanual stop; (3) left stimulus (LS, green left arrow and red right arrow), requiring a unimanual left response; and (4) right stimulus (RS, red left arrow and green right arrow), requiring a unimanual right response. The four corresponding trial outcomes are: (1) a dual response (DR, i.e., bimanual button press); (2) a non-response (NR, i.e., no button press); (3) a left response (LR, i.e., unimanual left button press); and (4) a right response (RR, i.e., unimanual right button press).

The results reported here focus on the seven most common stimulus–response combinations, which are predicted by both theories that we test: SS-NR, SS-DR, DS-DR, LS-LR, LS-DR, RS-RR and RS-DR (98.5% of the data). In supplementary materials (<https://osf.io/5z9vk/>), we show that an extension of SIS is able to accurately account for all stimulus–response combinations by allowing for motor errors (i.e., slips of the finger causing the wrong button to be pressed or responses too weak to cause a button press to be recorded).

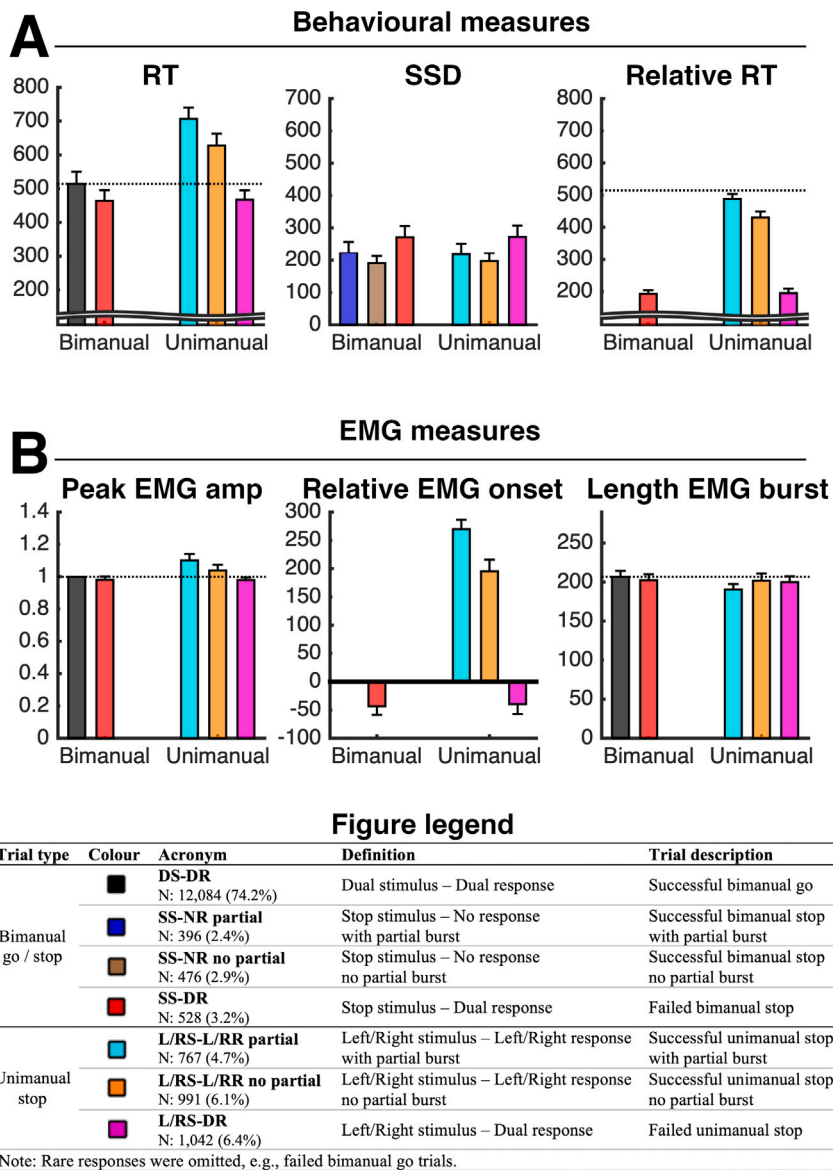


Fig. 3. Average results with 95% confidence intervals (CI) for responses for each trial type, using the stiff button. Successful trials were broken down by whether or not a partial response occurred, and with LS and RS results collapsed. The horizontal dotted lines represent the average for DS-DR as a reference. Compared with bimaneal trials with no stop signal (DS-DR), unimaneal responses showed faster relative RT and higher peak EMG amplitude, suggesting that the unimaneal runner is faster and produces stronger responses than the bimaneal runner. Among unimaneal responses (LS-LR, RS-RR), trials with partial burst have slightly later relative EMG onset and higher EMG amplitude than those without partial bursts.

3.1. Behavioural and EMG results

Fig. 3 shows the behavioural and EMG profiles obtained using stiff buttons; similar results for compliant buttons can be seen in the supplementary material. The left panel on the top row of Fig. 3 shows the fundamental signature of the standard race model: RT on failed bimaneal stop trials (i.e., SS-DR) is faster than on bimaneal go trials (i.e., DS-DR), because slow dual-response runners tend to loose to inhibitory runners more often than fast dual-response runners (GLMM $p < 0.001$, all post hoc $p < 0.001$). This is also true for failed trials requiring selective stopping (i.e., L/RS-DR). In contrast, a substantial stopping delay relative to trials without a stop signal is evident for successful unimaneal responses (i.e., L/RS-L/RR), with the delay being larger when there is a partial response.

The second panel on the top row of Fig. 3 shows that, as expected, trials on which stopping failed (either bimaneal or unimaneal) have longer SSDs (which was set by a staircase algorithm, see Materials and Methods) than trials where stopping succeeded (GLMM $p < 0.001$, post hoc $p < 0.001$). On successful bimaneal stop trials, there is little difference in average SSD between trials with and

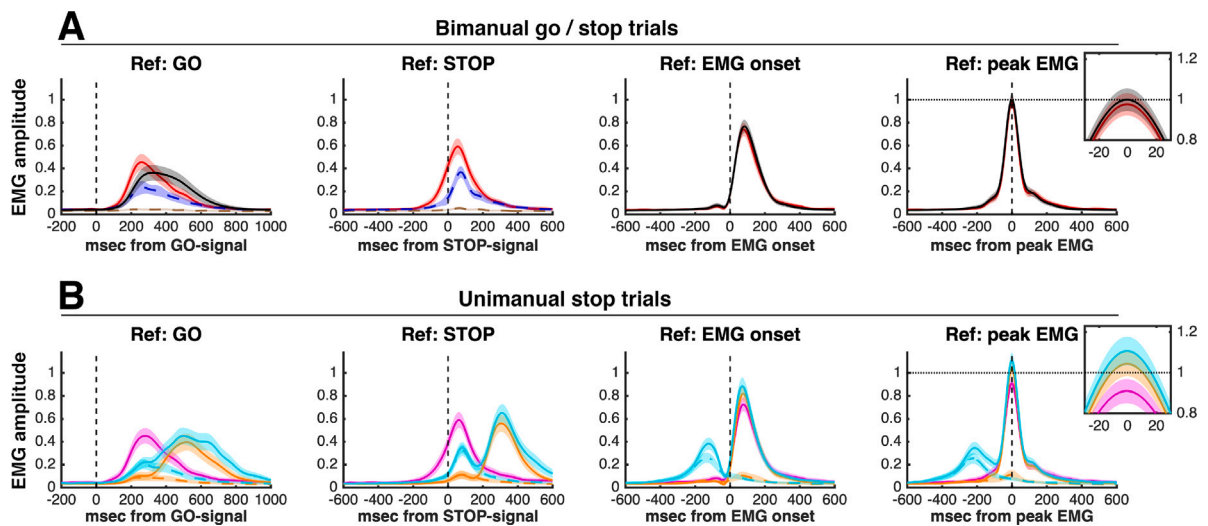


Fig. 4. Average EMG profiles and 95% CI for each trial type, using the stiff buttons. Bimanual and unimanual responses are displayed in the top and bottom rows, respectively. Each column represents the same EMG profiles aligned relative to a different time reference, represented by a vertical dashed line. For selective-stopping trials, EMG profiles of responding and stopping hands are represented by solid and dashed lines, respectively.

without partial responses. However, unimanual stop trials with partial responses have longer SSDs (post hoc $p = 0.01$), consistent with there being more time for the initiation of a (partial) dual response to begin before it is terminated by inhibition. The third panel on the top row of Fig. 3 plots RT with SSD subtracted (i.e., “relative” RT). Relative RT is very quick on failed stop trials, consistent with dual responses being triggered by the go stimulus (GLMM $p < 0.001$). Consistent with the unimanual runner being triggered by the selective stop signal, and being faster than the bimanual runner, relative RT from unimanual responses is slightly faster than RT from bimanual trials with no stop signal (post hoc $p < 0.001$). The difference between unimanual trials with and without a partial response is also attenuated by removing the systematic difference between them in SSD (post hoc $p < 0.001$).

The bottom row of Fig. 3 compares the peak EMG amplitude, the relative onset of RT-generating burst (relative to SSD), and its length. There is no evidence of a difference among these statistics between successful bimanual go trials or failed (bimanual or unimanual) stop trials (GLMM $p < 0.001$, post hoc both $p > 0.6$). In contrast, EMG amplitude is elevated on successful selective-stop trials (GLMM $p < 0.001$, post hoc both $p < 0.001$), consistent with relative RT being faster for these trials than for trials with no stop signal. Moreover, unimanual trials with partial bursts show later relative onsets of EMG than trials without partial bursts (post hoc $p < 0.001$).

Fig. 4 shows, for each trial type, the average EMG profiles obtained using the stiff buttons collapsed over left and right hands (analogous results for compliant buttons are given in supplementary materials). The top row shows responses to bimanual go and stop trials and the second row shows responses to unimanual stop trials, where the responding and stopping hands are represented by solid and dashed lines, respectively. In each column, the profiles were referenced to a different time event in order to highlight distinct aspects of the EMG responses. Note that, because EMG profiles were normalised to the average peak of the RT-generating burst, the amplitude of successful bimanual go trials is 1 when all peaks are aligned (i.e., “Ref: peak EMG”, right-most column). As the reference is moved away from the peak EMG, the shape and timing of the average EMG profiles blur as a result of time-averaging individual profiles whose peaks are not perfectly aligned.

When the EMG profiles are referenced to the go signal (“Ref: GO”), it is clear that responses on failed stop trials (both SS-DR and L/RS-DR) happened earlier than on successful bimanual go trials, i.e., they had an earlier EMG onset and peak latency. Also, note the similarity of responses between bimanual and unimanual failed stop trials, as both correspond to imperative bimanual responses that were not successfully cancelled.

Profiles referenced to the stop signal (“Ref: STOP”) highlight the strong dependence between SSD and responses during selective-stop trials. Consistent with (Jana, Hannah, Muralidharan, & Aron, 2020), failed stop trials show the largest SSDs, followed by successful trials with partial responses then trials without partial responses (see Fig. 3). This reflects that, on failed stop trials, the stop signal was presented too late for the already large EMG response to be cancelled. On successful stop trials with partial bursts, although the bimanual response had already been initiated when the stop signal was presented, it was still small enough to allow for cancellation. Finally, on successful stop trials with no partial bursts, the stop signal was presented early enough so no significant bimanual EMG activity had been initiated. Furthermore, note that, during successful unimanual stop trials, the EMG profiles of the responding and stopping hands (solid and dashed lines) are nearly identical before the onset of the RT-generation burst. We applied statistical parametric mapping to compare the EMG profiles between the responding and stopping hands, which confirmed that they differ significantly only 180 ms after the stop signal, during initiation of the unimanual response.

Profiles referenced to EMG onset show (i) strong similarity between responses to successful bimanual and failed stop trials (top row), in addition to (ii) a clear separation between partial and RT-generating bursts during successful selective-stop trials (second

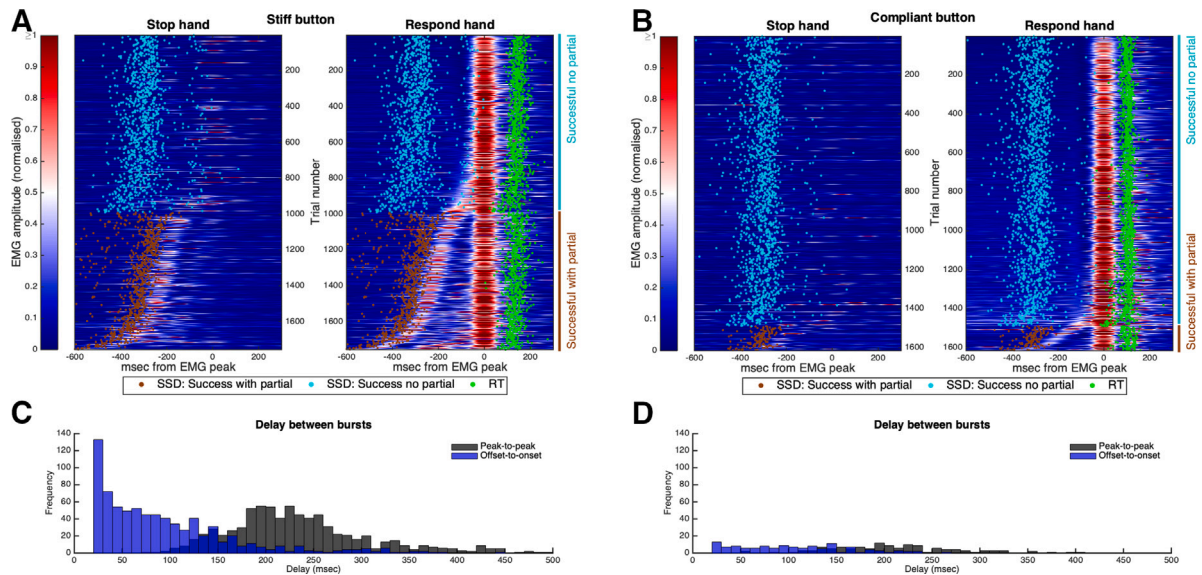


Fig. 5. A and B. Single-trial EMG profiles from all subjects and trials from stiff and compliant buttons, respectively, colour-coded by normalised EMG amplitude. Signals from each trial were referenced to the peak of the RT-generating burst of the responding hand. To highlight the relative timing of partial bursts, trials with partial burst are displayed at the bottom, sorted by the distance between the partial and RT-generating bursts. Other trials were sorted by the amplitude prior to RT-generating burst. C and D. Histogram of the delay between partial and RT-generating bursts, represented as peak-to-peak or offset-to-onset times.

row). Note that the algorithm used to detect EMG bursts requires that the EMG amplitude returns to baseline levels for at least 20 ms between the partial and main responses, otherwise they are identified as a single burst.

In profiles referenced to peak EMG, the dotted horizontal line marks the average (normalised) amplitude of the RT-generating burst during successful bimanual go trials. These plots indicate that the amplitude of the EMG response is similar during successful and failed bimanual trials, whereas successful unimanual responses show higher peak EMG amplitude, as highlighted on the small inset plot.

As expected, partial bursts were smaller than RT-generating bursts: Considering successful unimanual trials, the partial bursts have lower normalised amplitude (average partial vs. RT-generating: 0.51 vs. 1.12), slower rate of onset (8.8 vs. 19.5), and shorter duration (91 vs. 192 ms) than RT-generating bursts of the same (i.e., responding) hand.

3.2. Merging of partial and RT-generating bursts

We developed a robust algorithm to detect partial EMG bursts independently in the responding and stopping hands. The algorithm avoids detecting partial bursts that previous analyses have reported as occurring together with, or after, the RT-generating burst (Raud et al., 2020)—which likely correspond to mirror activity or other spurious muscle activation not related to the stop-signal task. Briefly (see Materials and Methods for full details), for each trial we find all EMG bursts using a single-threshold EMG onset detection algorithm (Hodges & Bui, 1996), and a partial burst is identified when (i) the offset happens after the stop signal; (ii) the peak latency happens before the onset of the RT-generating burst; and (iii) the normalised peak amplitude is at least 0.2. These conditions were necessary to ensure we select only bursts corresponding to activity that was cancelled in response to the stop signal.

In Fig. 5-A, we present the entire distribution of successful single-trial EMG responses to unimanual trials by pooling together data from all subjects, separated by stop and responding hands. The EMG profile from each trial is represented by a horizontal line, colour-coded by normalised amplitude, and referenced to the time of the peak RT-generating burst of the responding hand. To represent variations in the timing of partial bursts, successful trials with partial bursts are displayed at the bottom, sorted by the delay between the peaks of partial and RT-generating bursts, and thus illustrate the complete distribution of timing between the two bursts across our entire data set. The partial bursts, when present, correspond to cancelled bimanual responses, and usually (65% of the times for stiff buttons) happen on both hands simultaneously, with an average difference in the onset of partial bursts between hands of only 14 ms.

However, there were other trials where the EMG amplitude did not return to baseline after an initial response, but rather the partial burst merges with the RT-generating burst, resulting in a single burst with a long tail (i.e. skewed) to the left, prior to peak EMG. To represent the temporal relationship between partial and RT-generating bursts, we sorted the remaining trials (successful stop trials with no partial response) based on the average amplitude in the period between -160 and -80 ms from the peak EMG of the responding hand. The results provide evidence that, rather than a clear delay to the main response, the stopping delay manifests as a continuous distribution of timing between the partial and the RT-generating bursts. On some trials, the delay is large enough

to allow clear separation between bursts; on other trials, the delay is shorter until it becomes less than the precision of our EMG detection algorithm, and hence it becomes impossible to distinguish the two bursts, as they merge into a single burst.

Fig. 5-C shows the distribution of the delays between the partial and RT-generating bursts, both between the burst peaks, and the length of the silent period between bursts, i.e., the time between the offset of the partial and the onset of the RT-generating responses. The delay between peaks has a uni-modal distribution and is comprised only of positive values. Consistent with the merging between bursts evident qualitatively in Fig. 5-A, the silent period shows a distribution skewed towards small values, bordering on the 20 ms limit of our EMG onset detection algorithm for a large number of trials.

Figs. 5-B and 5-D show analogous results for compliant buttons. Compared with the stiff button, responses using the compliant button had a much smaller percentage of successful unimanual stop trials with partial bursts (14.8% vs. 51.4% of successful stop trials, or 8.2% vs. 32.1% of all unimanual stop trials), in addition to a larger percentage of failed stop trials (42.3% vs. 37.2%). These differences arise from the amount of force (and thus length of time) it takes to generate a press with each button type: The compliant button requires much less force to register a button press than the stiff button (1N vs. 7.4N). Hence participants have less time available to cancel the initial bimanual response after they perceive the stop signal. Nevertheless, the results for the compliant buttons are also consistent with a merging between partial and RT-generating bursts.

In summary, our results for both stiff and compliant buttons challenge the global-inhibition theory, which states that cancellation of the bimanual response is followed by a delay then re-initiation of a unimanual response (Coxon et al., 2007; MacDonald et al., 2014, 2017). Instead, they are consistent with an independent race between a bimanual runner, triggered by the go stimulus, and two runners that are triggered by the stop signal, namely a stop runner that can cancel the already active bimanual go runner and a new go runner that can initiate a unimanual response. If the stop runner wins, the partial response is cancelled, EMG amplitude returns to baseline levels before initiation of the RT-generating burst, and two distinct bursts are observed. In contrast, if the go runner wins, although a temporary reduction (or distortion) in EMG amplitude may still be observed prior to the peak of the RT-generating burst, the initial response is not completely cancelled, and a single burst is observed. In the next section we outline the details of the cognitive model we developed and fit it to our behavioural data.

3.3. Cognitive modelling

The SIS model is based on Matzke et al.'s (2013) BEESTS model of the standard stop-signal task. It is mathematically tractable, with straightforward solutions for the likelihoods of all stimulus–response combinations (provided in supplementary materials) enabling efficient fitting. Fig. 6 graphically depicts the ways in which it accounts for the seven common stimulus–response pairs occurring over the four trial types.

On DS trials (i.e., trials with no stop signal), the go stimulus triggers a single dual-go runner that simultaneously produces the left-hand and right-hand components of a bimanual response (DR) at time RT (drawn from an ex-Gaussian distribution). The unity of bimanual responses is supported by the coordinated way in which they occur, with each button press typically made almost simultaneously. In contrast, trial-to-trial variability in RT spans a much wider time scale (e.g., $SD \approx 100$ ms), so if each button press were to any appreciable degree independent, much less synchronicity would be expected (see MacDonald et al., 2014, for further evidence see prepareData.R at <https://osf.io/5z9vk/>).

On SS trials (i.e., trials requiring both button presses to be withheld), the stop signal can trigger a stop runner, which also has an ex-Gaussian finishing time distribution. However, with probability p_{TF} the stop signal may fail to trigger the stop runner (Matzke, Love et al., 2017), indicating a lapse of attention to task goals. Such “trigger failures” occur at elevated levels in clinical populations, where they can explain part or most of performance deficits in the standard stop-signal paradigm (Matzke, Hughes, Badcock, Michie, & Heathcote, 2017; Weigard, Heathcote, Matzke, & Huang-Pollock, 2019). They are also found in healthy participants, usually at lower levels, but often with appreciable individual differences (Matzke, Curley, Gong, & Heathcote, 2019; Matzke, Love et al., 2017), and when their effect is neglected it can lead to over-estimation of SSRT (Skippen et al., 2020, 2018). Hence, we allowed for the possibility that they also occurred in selective stopping. If a trigger failure occurs, processing follows the same course as in a DS trial. If it does not, two outcomes are possible: the stop runner wins the race and both left and right responses are withheld (i.e., a NR outcome) or the go runner wins and both responses are made (i.e., a DR outcome).

For the remaining two trial types, LS (i.e., the go signal is followed by a stop signal indicating respond left but not right) or RS (i.e., the go signal is followed by a stop signal indicating respond right but not left), trigger failures may also occur with probability p_{TF} , in which case processing again follows the same course as in a DS trial. Otherwise, the stop signal triggers both the stop runner and either a left runner (on LS trials), capable of producing a unimanual left response, or a right runner (on RS trials), capable of producing a unimanual right response. These successful unimanual responses can occur in two ways: (1) the stop runner wins the race, stopping the bimanual runner but not the unimanual runner, which then continues, or (2) the unimanual runner beats the other two runners. If, as it occurs in the standard stop-signal paradigm, a stop runner is much faster than a go runner, and this extends to unimanual as well as dual-go runners, then the latter case would be expected to be rare.

Our initial test compared versions of the model with and without trigger failures. All model fitting and assessment was performed in a Bayesian manner (see Materials and Methods for details), and we accompany model-based estimates with 95% credible intervals to indicate uncertainty (Morey, Hoekstra, Rouder, Lee, & Wagenmakers, 2016).² We selected between models using Ando's (2010)

² A 95% credible interval ranging from x to y contains 95% of the posterior mass and thus indicates that one can be 95% confident that the true value lies between x and y .

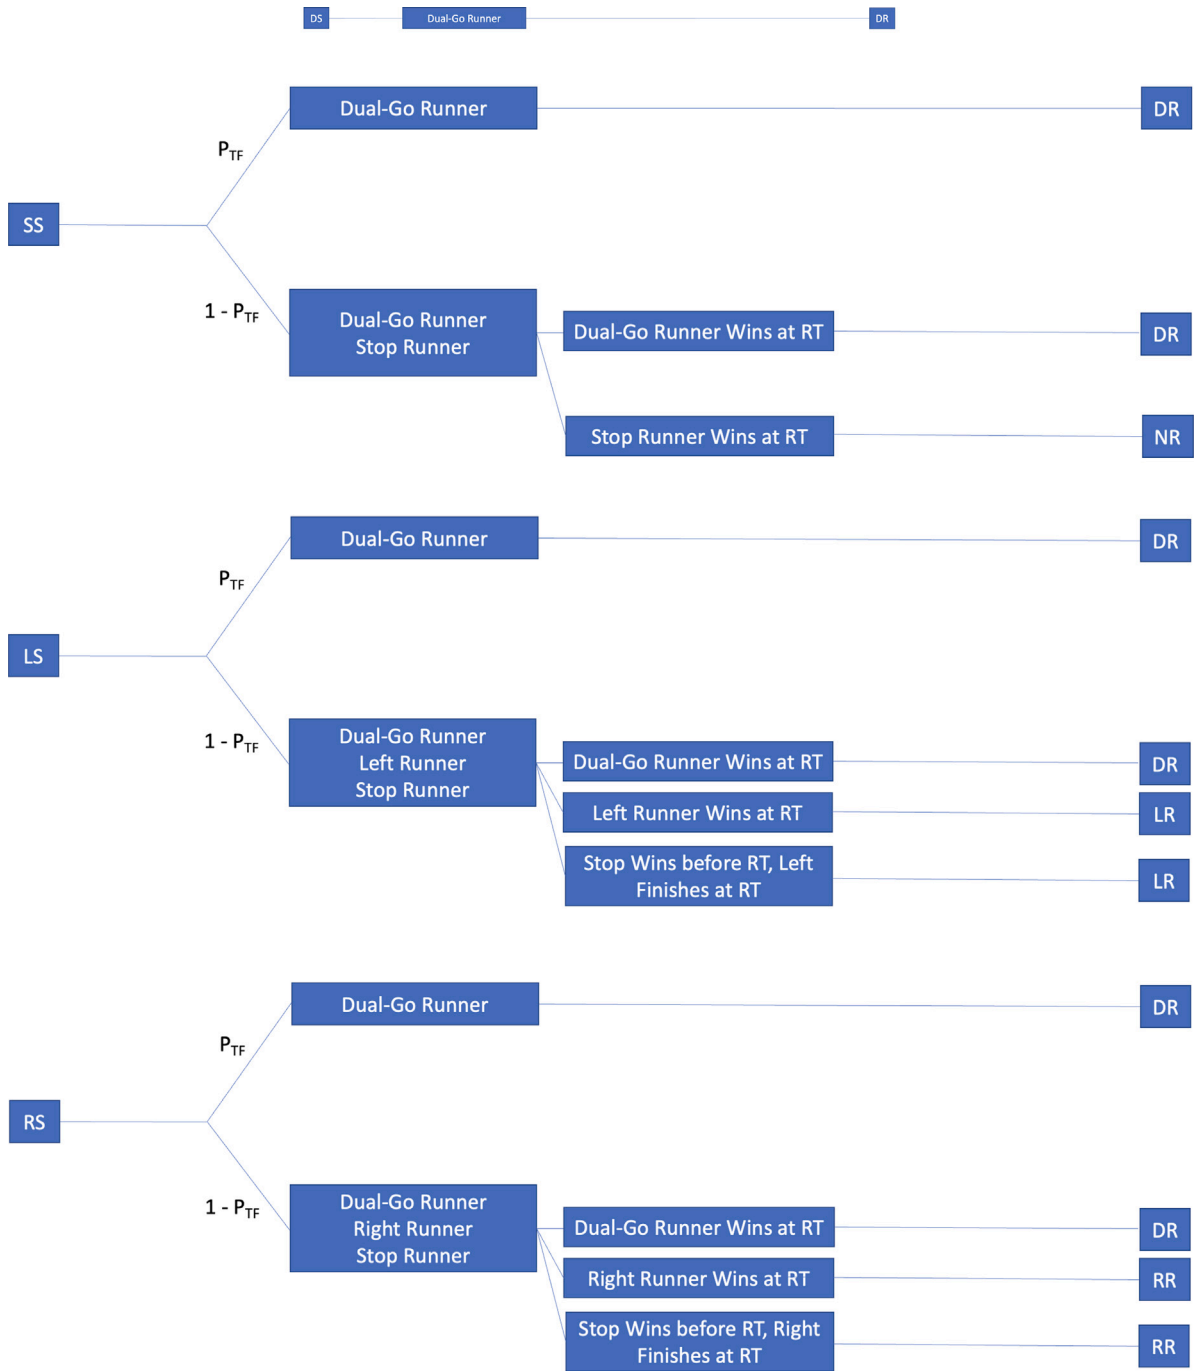


Fig. 6. Model architecture for the four trial types: DS (go stimulus with no stop signal), SS (go stimulus followed by a signal to stop both components of the bimanual response), LS (go stimulus followed by a signal to stop the right response but continue the left response), and RS (go stimulus followed by a signal to stop the right response but continue the left response).

BPIC criterion, which balances the extra model flexibility, and hence better fit to the present data, afforded by adding trigger failures against the potential for over-fitting, and hence poorer prediction of new data. BPIC favoured trigger failures by a margin of 115, providing very strong evidence that they should be included in the model. Although on average the percentage of trigger failures was small (0.8%, with 95% credible interval [0.4%–1.5%]), there was substantial individual variation, with trigger failures occurring on 10.1% [3.7%–14.5%] of trials for one participant. In light of these results, we selected the model with trigger failures for further analysis.

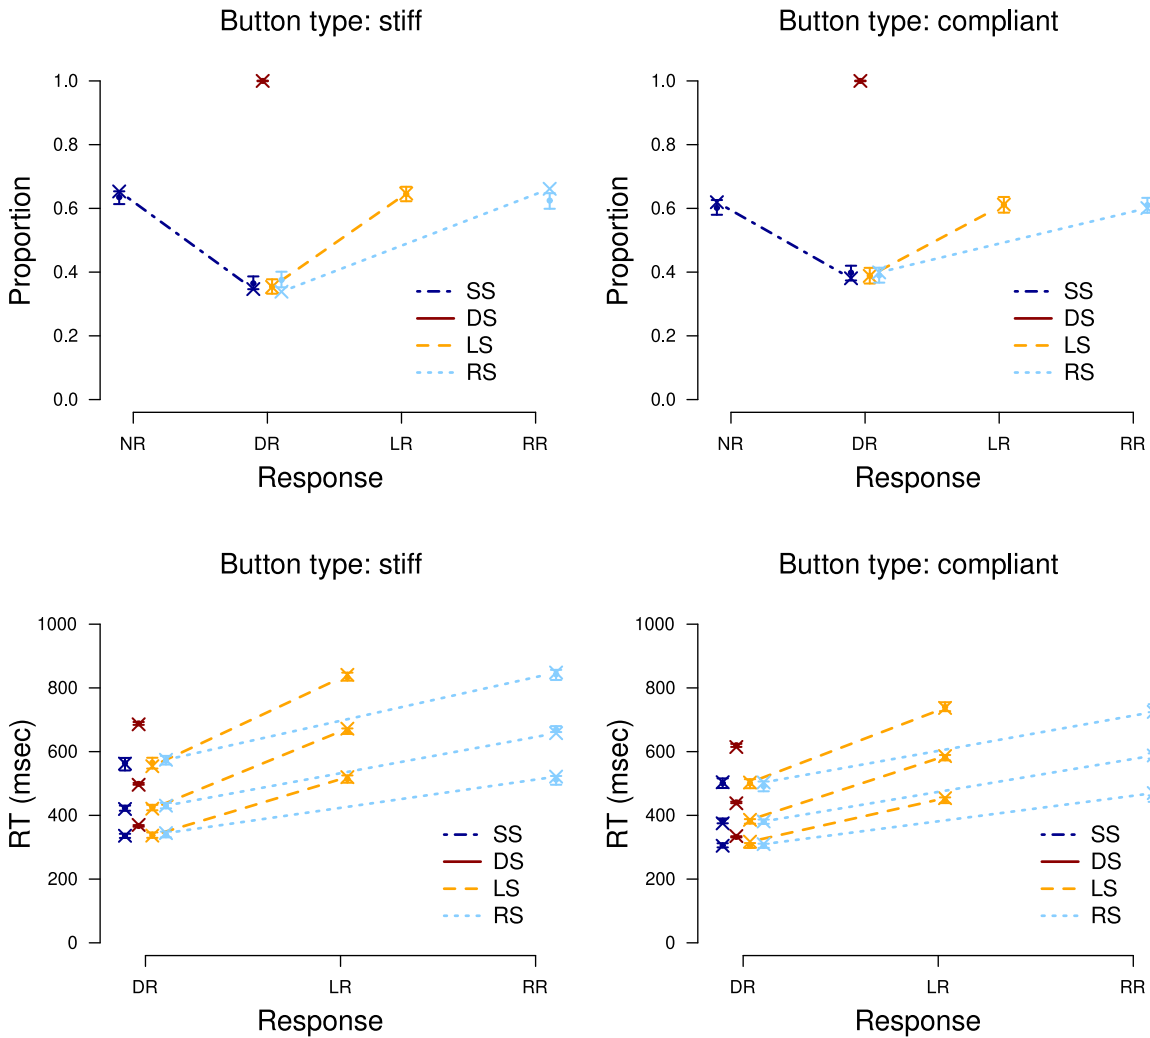


Fig. 7. Fit of SIS with trigger failures. Error bars show 95% posterior predictive credible intervals generated by simulating 100 data sets based on random draws of parameter vectors from the posterior distribution. RT distributions are represented by sets of increasing points from the same condition for the 10th, 50th, and 90th percentiles. SS = dual-stop stimulus trials; DS = dual-go stimulus trials; LS = left-go stimulus trials; and RS = right-go stimulus trials. NR = no response; DR = dual (bimanual) response; LR = left unimanual response; RR = right unimanual response.

Fig. 7 shows model fit, with a similar pattern of performance, but slightly faster responding, for compliant relative to stiff buttons. Failures of stopping (i.e., dual responses) occurred at about the same rate in all three conditions with stop signals. Substantial slowing of unimanual responses relative to dual responses is evident. The figure also shows that the race model’s signature prediction, faster responding with than without a stop signal, occurs to the same extent for selective and non-selective stopping, and increases as responses slow, so that it is quite substantial for the 90th percentile. The model captures all of these effects, including the fine-grained effects on RT distributions.

Estimates of the model’s parameters are provided in Table 3. Fig. 8 combines parameters to show the model’s estimates of the average speed of the runners. The time for the dual-stop runner, which corresponds to SSRT, is much faster than the times for the go runners, just as is the case in the standard stop-signal paradigm. However, it is also slower in an absolute sense than the SSRT from the standard paradigm for comparable populations performing the same type of easy choice task, which is typically considerably less than 200 ms when trigger failures are taken into account (Skippen et al., 2020). Table 3 also shows that slowing due to stiff buttons (accounted for in the μ parameter) was slightly larger for bimanual than unimanual responses, and that trigger failures were rare, suggesting that participants were highly engaged with the task.

Fig. 8 also shows that the unimanual runner is slower than the stop runner. As a result, it rarely beats the stop runner, and so most unimanual responses occur after the stop runner first beats the bimanual runner. The unimanual runner is, however, faster than the bimanual runner, consistent with our behavioural findings, and that of others (MacDonald et al., 2014, 2017).

Table 3

Parameter estimates for the basic SIS model. Displayed are the posterior medians (i.e., 50%) for the group-level means, accompanied by 95% credible intervals (i.e., 2.5% and 97.5%). Estimates for P_{TF} have been converted from the probit to the probability scale for easier interpretation.

Parameter	2.5%	50%	97.5%
$\mu^{\text{compliant}}$	0.37	0.40	0.43
μ^{stiff}	0.42	0.46	0.49
σ	0.05	0.06	0.06
τ	0.05	0.06	0.07
$\mu_{\text{stop}}^{\text{compliant}}$	0.20	0.21	0.22
$\mu_{\text{stop}}^{\text{stiff}}$	0.21	0.22	0.23
σ_{stop}	0.00	0.01	0.03
τ_{stop}	0.00	0.01	0.02
$\mu_{\text{U}}^{\text{compliant}}$	0.36	0.38	0.39
$\mu_{\text{U}}^{\text{stiff}}$	0.41	0.42	0.43
σ_{U}	0.05	0.06	0.07
τ_{U}	0.01	0.03	0.04
P_{TF}	0.00	0.01	0.02

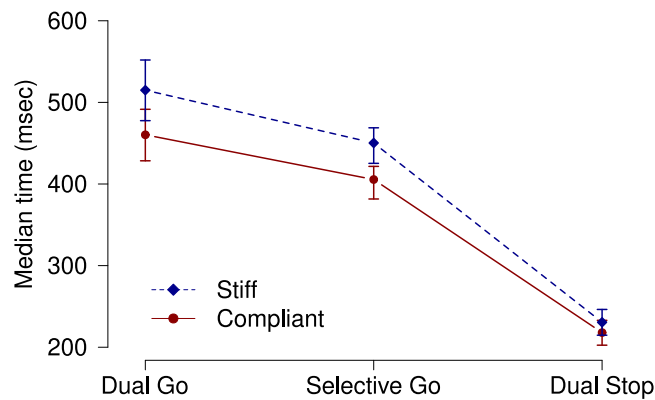


Fig. 8. Average estimated running times for SIS with trigger failures in terms of the median over each participant's mean RT (i.e., $\mu + \tau$ estimates); the plot displays the posterior median estimate for this quantity accompanied by a 95% credible interval. Note that the model assumes that the left and right unimanual runners have the same speed. The individual-level estimates (not displayed) indicate that there was a meaningful difference between the running times of the dual-go and unimanual runners for both the stiff button (median difference: 60 ms, 95% credible interval: [57 ms, 64 ms]) and the compliant button (median difference: 51 ms, 95% credible interval: [47 ms, 54 ms]). However, there was no meaningful difference in running time between the stiff and the compliant dual stop runners (median difference: 13 ms, 95% credible interval: [-2 ms, 27 ms]).

4. Discussion

Previous investigations of selective stopping suggested that a global stop mechanism, which quickly cancels both components of an initial multicomponent action, is followed by a subsequent restart of the component of the initial action that was not required to stop (MacDonald et al., 2014, 2017). Here, we present novel and converging lines of evidence that support a reconceptualisation of the mechanisms underlying selective stopping.

We utilised a bimanual response task requiring simultaneous button presses with left and right index fingers (a bimanual go response), with the subsequent stop signal necessitating cancellation of either one or both button presses. Our cognitive model, which provided a highly accurate account of the behavioural data, suggests that selective stopping is accomplished by the same type of independent race architecture that accurately characterises non-selective stopping (i.e., as assessed using a standard stop-signal task), but with the addition that a selective-stop signal simultaneously triggers not only a stop runner but also a new unimanual go runner. Specifically, the bimanual go runner commences upon presentation of the go stimulus, and races against the stop runner that commences upon presentation of the stop stimulus, irrespective of whether that stimulus requires both, or only one, component of the original action to be cancelled. If successful, the stop runner cancels the initial bimanual response. If the stop signal also indicates that selective stopping is required, a unimanual go runner that can trigger the non-stopping hand also joins the race at the time of stop-signal presentation.

There are potential performance advantages to this parallel processing architecture compared to a serial architecture where a unimanual response component is re-initiated after successful stopping has occurred (MacDonald et al., 2017). Typically, in our model fits for selective-stop trials, the fast stop process causes the initial bimanual action to be withheld only shortly before the

slower unimanual runner produces the required (left or right) response. The stopping delay, characteristic of selective-stopping tasks and replicated in our data (~ 100 ms), is thus a consequence of a new unimanual go response being triggered by the stop signal, which is delayed relative to the go signal. Interestingly, however, typical SSDs are a little longer than this stopping delay. This is consistent with the model fits that indicate unimanual go runners are slightly faster than the dual go runners for both the stiff and compliant buttons (Fig. 5), which partially mitigates the stopping delay despite the substantially later start time. Our behavioural data indicate that RTs on successful selective-stop trials, calculated relative to the stop signal, are faster than the standard bimanual response times (Fig. 2A—Relative RT panel), corroborating the conclusions drawn from model fits. Moreover, our EMG measures also suggest that unimanual responses during successful selective-stop trials are faster (earlier EMG onset) and stronger (higher peak EMG amplitude) than the initial bimanual responses. Finally, despite the model allowing flexibility for the stop process to vary between stiff and compliant buttons, the fits suggest that this stopping process was invariant over button types. Accordingly, the cancellation process was unaffected by both the nature of the stop (selective or bimanual) or the mechanics of the cancelled response (stiff or compliant buttons), reflecting its fast and relatively automatic nature.

Overall, and in contrast to the standard independent race model, our new model provides a way of estimating SSRT in selective-stopping paradigms that has a clear justification in terms of a plausible cognitive and neural process, and thus we recommend its use. To this end, we make the methods used to fit the SIS model openly available (<https://osf.io/5z9vk/>). We also note that SIS may be applicable to other types of complex stopping, such as stop-change paradigms where the stop signal indicates that an ongoing choice response must be changed to the alternative choice response (see Camalier et al., 2007, for a similar proposal).

Single trial analysis of overt and covert EMG (i.e., muscle activation above and below the threshold for generating a button press: ‘RT-generating’ and ‘partial’ bursts, respectively) fully support the conclusions drawn from the cognitive modelling. We imposed tight temporal constraints on when the overt and (preceding) partial responses could be detected, based upon expected response times and the model framework. This enabled us to make significant advances to recent work in this field (e.g., Raud et al., 2020; Raud, Thunberg, & Huster, 2022) and ensured that all partial bursts detected were related to cancellation of initiated (bimanual) actions. On selective-stop trials, we were able to reliably capture both partial and RT-generating bursts in the same hand (the non-stopping effector). These partial bursts were observed when responding with the stiff buttons on approximately half of successfully inhibited trials, irrespective of whether the required stopping was bimanual or selective, validating the model-based conclusion that the stopping mechanism was invariant across stop types. The partial burst, when observed, was commonly ($\sim 65\%$ of such trials) detected simultaneously in the stopping as well as the responding hand, indicative of cancellation of the initial bimanual response.

Critically, in successful selective-stop trials with partial bursts, the temporal separation between the partial and subsequent RT-generating bursts followed a continuous distribution, with any estimated separation only constrained by the limit of our algorithm to detect EMG onsets (i.e., we required a nominal 20 ms separation to characterise the partial and RT bursts as being distinct). This continuum from distinct, well-defined bursts, to a merging of bursts (Fig. 5) is wholly consistent with the cognitive model, which predicts that—based on the specific parameters of the stop and unimanual go distributions—the unimanual response time approaches the time at which the bimanual response is cancelled. A framework that predicts a nominal restart of one component of the initial bimanual response following successful cancellation (like the ATM: MacDonald et al., 2017) cannot account for this merging of the partial and subsequent RT generating bursts as a non-negligible restart time is predicted.

Previous EMG algorithmic approaches to assess response cancellation in selective-stopping tasks (for a recent review, see Raud et al., 2022) have searched for partial and RT generating bursts in separate effectors (i.e., the stopping and responding hands, respectively) without imposing tight temporal constraints on their identification. This approach yielded an average delay between the peaks of the partial and the subsequent RT burst of ~ 145 ms. However, in some cases the range of delays (“delta peak” in Raud et al., 2020, Fig 3E) extends to 459 ms, with a significant proportion of trials yielding negative delays up to -219 ms, where partial bursts occur after the RT burst. Such instances cannot represent cancellation of the bimanual response, but rather represent mirror activation during the voluntary selective response or later activity not time-locked to visual responses (e.g., postural adjustments or muscle ‘twitches’). Including partial EMG bursts occurring after the overt burst will result in overestimation of the key measure of ‘CancelTime’, i.e. the time between presentation of the stop signal and the ensuing reduction in amplitude of the muscle activity (Jana et al., 2020). We therefore propose that the current measure of CancelTime derived using our EMG algorithm (~ 110 ms averaged across selective-stop and bimanual-stop trials) is a more accurate reflection of the minimum time required to cancel initiated responses, suggesting the process is faster than suggested by previous EMG measures (e.g. 148–174 ms Raud et al., 2020). The short stopping latencies in our EMG data are also consistent with reports of neural braking of the bimanual response 117 ms after the stop signal (MacDonald et al., 2017). Note that CancelTime is necessarily a downward-biased measure of SSRT because it does not take account of cases where slow stop runners lose the race. However, although SSRT is conventionally thought to be greater than 200 ms (Verbruggen et al., 2019), estimates are shorter when trigger failures are taken into account (e.g., 130 ms in Skippen et al., 2020).

A key finding is that our EMG estimates of CancelTime did not vary reliably between selective-stop and bimanual-stop trials (Fig. 1B, Ref: STOP; 2nd and 3rd row; also Results section ‘Basic Behavioral and EMG Results’). This finding further supports that the stopping mechanism does not vary between selective-stop and bimanual-stop trials, and corroborates the model framework which designates a universal stop process that is invariant between bimanual and selective trials. Furthermore, neither the model’s SSRT estimates (Fig. 5) nor the EMG estimates of CancelTime vary between stiff and compliant buttons. In contrast, Raud et al. (2020) reported substantially longer CancelTimes for bimanual compared to selective-stop trials, likely because of inclusion of partial responses not associated with action stopping in the calculations.

In future work SIS could be extended, with the addition of extra assumptions, to incorporate predictions about EMG in general and partial responses in particular. For example, EMG could be assumed to begin to increase at some time t_{onset} before a response

racer finishes, and the finishing time could be mapped to a “point of no return” (t_{pnr} , likely at or just after the EMG peak) after which a response cannot be inhibited. If the stop runner finishes before the point of no return but after t_{onset} then a partial response results. However, if the stop runner finishes at an earlier time then a “cortical stop” occurs without an accompanying partial response. It is also possible that when the stop runner finishes it takes some short time, $t_{inhibit}$, to start inhibiting a rising bimanual EMG burst, in which case a partial response occur when it finishes between $t_{onset} - t_{inhibit}$ and $t_{pnr} - t_{inhibit}$. Further, a relationship between the speed of runners and the amplitude of the EMG bursts they induce would need to be specified, with our results here suggesting an inverse relationship (i.e., faster runners are associated with larger magnitudes). When fitting such a joint model the onset and other times might be modelled as either constants or random variables, and an account would have to be taken of the resolution with which partial bursts can be detected (e.g., up to 20 ms before the onset of the response generating burst with our algorithm).

In conclusion, we present a new cognitive model that provides a parsimonious explanation of the key action and inhibition processes in a response-selective inhibition paradigm, which is wholly consistent with our rigorous EMG analyses. Together, the model and EMG data offer additional insights into the ecologically relevant task of selective stopping of planned actions, and provide robust methods to accurately measure the inhibitory processes it depends on.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

Data and code are available at <https://osf.io/5z9vk/>.

Acknowledgements

This work was supported by the Australian Research Council (DP200101696). DM is supported by a Vidi grant (VI.Vidi.191.091) from the Dutch Research Council (NWO).

References

- Ando, T. (2010). *Bayesian model selection and statistical modeling*. Boca Raton, FL: CRC Press.
- Aron, A. R. (2011). From reactive to proactive and selective control: Developing a richer model for stopping inappropriate responses. *Biological Psychiatry*, 69(12), e55–e68. <http://dx.doi.org/10.1016/j.biopsych.2010.07.024>, URL <http://dx.doi.org/10.1016/j.biopsych.2010.07.024>.
- Aron, A. R., & Verbruggen, F. (2008). Stop the presses. *Psychological Science*, 19(11), 1146–1153. <http://dx.doi.org/10.1111/j.1467-9280.2008.02216.x>, URL <http://journals.sagepub.com/doi/10.1111/j.1467-9280.2008.02216.x>.
- Brooks, S. B., & Gelman, A. (1998). General methods for monitoring convergence of iterative simulations. *Journal of Computational and Graphical Statistics*, 7, 434–455.
- Camalier, C., Gotler, A., Murthy, A., Thompson, K., Logan, G., Palmeri, T., et al. (2007). Dynamics of saccade target selection: Race model analysis of double step and search step saccade production in human and macaque. *Vision Research*, 47(16), 2187–2211. <http://dx.doi.org/10.1016/j.visres.2007.04.021>, URL <http://linkinghub.elsevier.com/retrieve/pii/S0042698907001782>.
- Coxon, J. P., Stinear, C. M., & Byblow, W. D. (2007). Selective inhibition of movement. *Journal of Neurophysiology*, 97(3), 2480–2489. <http://dx.doi.org/10.1152/jn.01284.2006>, URL <http://jn.physiology.org/cgi/doi/10.1152/jn.01284.2006>.
- Heathcote, A., Lin, Y.-S., Reynolds, A., Strickland, L., Gretton, M., & Matzke, D. (2019). Dynamic models of choice. *Behavior Research Methods*, 5, 961–985.
- Hodges, P. W., & Bui, B. H. (1996). A comparison of computer-based methods for the determination of onset of muscle contraction using electromyography. *Electroencephalography and Clinical Neurophysiology/Electromyography and Motor Control*, 101(6), 511–519.
- Ishihara, S. (1972). *The series of plates designed as a test for colour deficiency: 38 plates edition*. Kanehara Trading Inc.
- Jana, S., Hannah, R., Muralidharan, V., & Aron, A. R. (2020). Temporal cascade of frontal, motor and muscle processes underlying human action-stopping. *Elife*, 9, Article e50371.
- 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. <http://dx.doi.org/10.1037//0033-295x.91.3.295>, URL <http://content.apa.org/journals/rev/91/3/295>.
- MacDonald, H. J., Coxon, J. P., Stinear, C. M., & Byblow, W. D. (2014). The fall and rise of corticomotor excitability with cancellation and reinitiation of prepared action. *Journal of Neurophysiology*, 112(11), 2707–2717. <http://dx.doi.org/10.1152/jn.00366.2014>, URL <http://www.physiology.org/doi/10.1152/jn.00366.2014>.
- MacDonald, H. J., McMorland, A. J. C., Stinear, C. M., Coxon, J. P., & Byblow, W. D. (2017). An activation threshold model for response inhibition. In G. L. Woodhall (Ed.), *PLoS ONE*, 12(1), Article e0169320. <http://dx.doi.org/10.1371/journal.pone.0169320>, URL <http://dx.plos.org/10.1371/journal.pone.0169320>.
- Matzke, D., Curley, S., Gong, C. Q., & Heathcote, A. (2019). Inhibiting responses to difficult choices. *Journal of Experimental Psychology: General*, 148(1), 124–142. <http://dx.doi.org/10.1037/xge0000525>, URL <http://eutils.ncbi.nlm.nih.gov/entrez/eutils/elink.fcgi?dbfrom=pubmed&id=30596441&retmode=ref&cmd=prlinks>.
- Matzke, D., Dolan, C. V., Batchelder, W. H., & Wagenmakers, E.-J. (2015). Bayesian estimation of multinomial processing tree models with heterogeneity in participants and items. *Psychometrika*, 80, 205–235.
- Matzke, D., Hughes, M., Badcock, J. C., Michie, P., & Heathcote, A. (2017). Failures of cognitive control or attention? The case of stop-signal deficits in schizophrenia. *Attention, Perception & Psychophysics*, 79(4), 1078–1086. <http://dx.doi.org/10.3758/s13414-017-1287-8>.
- Matzke, D., Love, J., & Heathcote, A. (2017). A Bayesian approach for estimating the probability of trigger failures in the stop-signal paradigm. *Behavior Research Methods*, 49, 267–281.
- Matzke, D., Love, J., Wiekci, T., Brown, S. D., Logan, G. D., & Wagenmakers, E.-J. (2013). Releasing the BEESTS: Bayesian estimation of stop-signal reaction time distributions. *Frontiers in Quantitative Psychology and Measurement*, 4:918.

- Matzke, D., Verbruggen, F., & Logan, G. (2018). The stop-signal paradigm. In E.-J. Wagenmakers, & J. T. Wixted (Eds.), *Stevens' handbook of experimental psychology and cognitive neuroscience, volume 5: methodology* (4th Ed.). (pp. 383–427). John Wiley & Sons, Inc.
- Morey, R. D., Hoekstra, R., Rouder, J. N., Lee, M. D., & Wagenmakers, E.-J. (2016). The fallacy of placing confidence in confidence intervals. *Psychonomic Bulletin & Review*, 23(1), 103–123. <http://dx.doi.org/10.3758/s13423-015-0947-8>.
- Peirce, J. W. (2009). Generating stimuli for neuroscience using psychopy. *Frontiers in Neuroinformatics*, 2, 10.
- Raud, L., Huster, R. J., Ivry, R. B., Labruna, L., Messel, M. S., & Greenhouse, I. (2020). A single mechanism for global and selective response inhibition under the influence of motor preparation. *Journal of Neuroscience*, 40(41), 7921–7935.
- Raud, L., Thunberg, C., & Huster, R. J. (2022). Partial response electromyography as a marker of action stopping. *Elife*, 11, Article e70332.
- Sebastian, A., Forstmann, B. U., & Matzke, D. (2018). Towards a model-based cognitive neuroscience of stopping – a neuroimaging perspective. *Neuroscience & Biobehavioral Reviews*, 90, 130–136. <http://dx.doi.org/10.1016/j.neubiorev.2018.04.011>.
- Shiffrin, R. M., Lee, M. D., Kim, W., & Wagenmakers, E.-J. (2008). A survey of model evaluation approaches with a tutorial on hierarchical Bayesian methods. *Cognitive Science*, 32, 1248–1284.
- Skippen, P., Fulham, W. R., Michie, P. T., Matzke, D., Heathcote, A., & Karayanidis, F. (2020). Reconsidering electrophysiological markers of response inhibition in light of trigger failures in the stop-signal task. *Psychophysiology*, 57(10), Article e13619. <http://dx.doi.org/10.1111/psyp.13619>, URL <https://onlinelibrary.wiley.com/doi/abs/10.1111/psyp.13619>.
- Skippen, P., Matzke, D., Heathcote, A., Fulham, W. R., Michie, P., & Karayanidis, F. (2018). Reliability of triggering inhibitory process is a better predictor of impulsivity than SSRT. *Acta Psychologica*, 192, 104–117. <http://dx.doi.org/10.1016/j.actpsy.2018.10.016>.
- Turner, B. M., Sederberg, P. B., Brown, S. D., & Steyvers, M. (2013). A method for efficiently sampling from distributions with correlated dimensions. *Psychological Methods*, 18, 368–384.
- Verbruggen, F., Aron, A. R., Band, G. P., Beste, C., Bissett, P. G., Brockett, A. T., et al. (2019). A consensus guide to capturing the ability to inhibit actions and impulsive behaviors in the stop-signal task. *eLife*, 8, Article e46323. <http://dx.doi.org/10.7554/elife.46323>, URL <https://elifesciences.org/articles/46323>.
- Weigard, A., Heathcote, A., Matzke, D., & Huang-Pollock, C. (2019). Cognitive modeling suggests that attentional failures drive longer stop-signal reaction time estimates in attention deficit/hyperactivity disorder. *Clinical Psychological Science*, 7(4), 856–872. <http://dx.doi.org/10.1177/2167702619838466>, URL <http://journals.sagepub.com/doi/10.1177/2167702619838466>.