american physiological society

JNP JOURNAL OF NEUROPHYSIOLOGY

J Neurophysiol 131: 757–767, 2024.
First published March 13, 2024; doi:10.1152/jn.00434.2023

RESEARCH ARTICLE

Control of Movement

Action initiation and action inhibition follow the same time course when compared under matched experimental conditions

Yue Du, Alexander D. Forrence, Delaney M. Metcalf, and Adrian M. Haith

Department of Neurology, Johns Hopkins University, Baltimore, Maryland, United States

Abstract

The ability to initiate an action quickly when needed and the ability to cancel an impending action are both fundamental to action control. It is often presumed that they are qualitatively distinct processes, yet they have largely been studied in isolation and little is known about how they relate to one another. Comparing previous experimental results shows a similar time course for response initiation and response inhibition. However, the exact time course varies widely depending on experimental conditions, including the frequency of different trial types and the urgency to respond. For example, in the stop-signal task, where both action initiation and action inhibition are involved and could be compared, action inhibition is typically found to be much faster. However, this apparent difference is likely due to there being much greater urgency to inhibit an action than to initiate one in order to avoid failing at the task. This asymmetry in the urgency between action initiation and action inhibition makes it impossible to compare their relative time courses in a single task. Here, we demonstrate that when action initiation and action inhibition are measured separately under conditions that are matched as closely as possible, their speeds are not distinguishable and are positively correlated across participants. Our results raise the possibility that action initiation and action inhibition may not necessarily be qualitatively distinct processes but may instead reflect complementary outcomes of a single decision process determining whether or not to act.

NEW & NOTEWORTHY The time courses of initiating an action and canceling an action have largely been studied in isolation, and little is known about their relationship. Here, we show that when measured under comparable conditions the speeds of action initiation and action inhibition are the same. This finding raises the possibility that these two functions may be more closely related than previously assumed, with potentially important implications for their underlying neural basis.

action inhibition; action initiation; reaction time; stop-signal task; timed response

INTRODUCTION

In some situations, external events require us to rapidly generate an action. For instance, if a pedestrian steps into the street, a driver may need to quickly step on the brake pedal. At other times, we may need to prevent a planned action from being initiated. For instance, a pedestrian about to step into the street may need to prevent themselves from doing so if a speeding car suddenly approaches. These reactive processes of action initiation and action inhibition in response to external events are complementary to one another in dictating whether or not we perform actions in the world. Action initiation and action inhibition are generally considered to be qualitatively distinct, however, and possibly supported by distinct neural circuitry. In previous work, these two processes have mostly been studied separately with distinct experimental approaches. Consequently, it has not been possible to directly compare their properties and, more generally, understand exactly how they relate to one another.

The speed of action initiation has been studied through simple reaction time (RT) tasks. In this approach, participants generate a predetermined response (e.g., a button press) as rapidly as they can after a cue is presented. People are typically able to initiate a response in this way in ~180–220 ms (1, 2). Action inhibition has mostly been studied with the stop-signal task. In this task, people are asked to choose and generate a response to an imperative "go" stimulus as quickly as possible (e.g., pressing a specific button) but must cancel this response in the event of a "stop" signal. It is

Correspondence: Y. Du (ydu27@jhmi.edu).
Submitted 24 November 2023 / Revised 15 February 2024 / Accepted 7 March 2024

typically found that people can successfully cancel an action if the cue arrives ~200–250 ms before the action would be initiated (3–6). However, the exact timings of action initiation and action inhibition vary widely and seem to be strongly dependent on experimental conditions, such as the number of choice responses involved and the frequency of trials in which a response is required (7–10). This makes it challenging to compare action inhibition and action initiation by comparing previous studies.

In principle, it should be possible to compare the speed of action initiation and action inhibition within the stop-signal task, since it engages both processes. In general, it takes a shorter time to cancel a response than to initiate one in the stop-signal task (10, 11). However, this apparent speed advantage of action inhibition over action initiation likely does not reflect a general difference in their properties. Because stop-signal tasks were originally designed to examine action inhibition, they contain an inherent asymmetry between the urgency to generate a response to the go cue and the urgency to cancel a response after the stop cue; cancellation of a response has to occur before the movement is generated, but, in contrast, a response can be voluntarily produced at any time after the go cue. In fact, despite being asked to respond as quickly as they can, participants have been found to systematically delay the time at which an action is generated to accommodate the stopping behavior (12–15), leading to much longer response times than would normally be required to initiate a response. This suggests that the time at which the response is seen to be generated may not truly reflect how fast an action could be initiated (16, 17). In addition, because of their primary focus on action inhibition, stop-signal tasks also contain an asymmetry in the frequency of trials that require a response or not (e.g., 25% "stop" trials vs. 75% "respond" trials). These asymmetries make it impossible to directly compare the properties of action initiation and action inhibition within a single task.

To more fairly compare the relative speeds of action initiation and action inhibition, we performed an experiment to measure them in two separate tasks in which participants had to either generate an action or cancel an intended action in a reactive manner. By measuring action initiation and inhibition in separate tasks, we were able to match the conditions as closely as possible. Key to our experimental design was that participants were trained to always respond at a prescribed time in each trial (i.e., a timed-response approach) (17) or to withhold their prepared responses at the prescribed time (i.e., anticipated response inhibition approach) (3, 11, 18). By occasionally and unexpectedly switching the required behavior, either from requiring no response to requiring a response or from requiring a response at the prescribed time to requiring no response, we were able to measure the time course of initiation and inhibition in symmetric conditions, allowing us to directly compare them.

MATERIALS AND METHODS

Experimental Design

Participants.

Thirty-six right-handed participants (15 female; 1 nonbinary) between 18 and 41 yr of age took part in the study. The experimental procedure was approved by the Johns Hopkins School of Medicine Institutional Review Board. All participants gave written informed consent and received $15 per hour for their participation. Data from one participant who did not follow the task instructions well were excluded from analyses (Supplemental Figs. S1 and S2).

General procedures.

Participants sat in front of a laptop with a gray screen and with a keypad next to it. The keypad was positioned so that participants could comfortably rest the index finger of their right hand on a mechanical key mounted on the keypad. On every trial, a white target line was placed with the same distance from the bottom of the screen and a circle was displayed at the top center of the screen (Fig. 1). Once the trial started, the circle moved downward vertically and participants were asked to either press the key or do nothing when the moving circle reached the target line. The circle stopped moving once a response was registered, or it kept moving toward the bottom of the screen.

In each trial, whether or not a response was required in a given trial depended on the color of the circle (white or black; more details in *R-to-NR condition* and *NR-to-R condition*). Since perceptual processing (potentially influenced by color and illuminance contrast of visual cues) plays a critical role in action initiation (1) and action inhibition (19), we

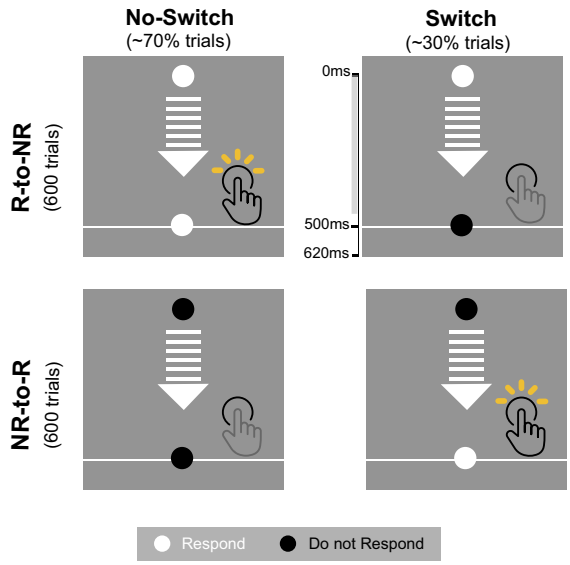

Figure 1. Experiment procedure. Participants were asked to either press a key or do nothing when a moving circle reached the target line. Whether or not a response was required in a given trial depended on the color of the circle (e.g., white = respond; black = do not respond). The actual circle color was counterbalanced across participants, controlling for the potential perceptual differences of black and white colors. The circle always started with the same color within each block of 100 trials. In ~70% of trials, the circle remained the same color throughout. However, in another ~30% of trials, the circle changed color before it hit the target line, forcing participants to cancel a preplanned response or to initiate a response when the circle crossed the line. By manipulating the time at which the circle color changed in each condition, we were able to compare people's ability to stop themselves from generating a planned response ("R-to-NR" condition; *top*) to the ability to rapidly generate a response ("NR-to-R" condition; *bottom*). Participants completed 12 blocks of 100 trials, generating 204 switch trials in each condition.

counterbalanced the association between response and the color of the moving stimulus so that for half of the participants white color cued a response (i.e., press the key) and black color indicated that no response was needed (i.e., do not press the key) and this association was reversed for the other half of the participants, so as to control for the potential perceptual differences between black and white colors.

Feedback was given in each trial. For trials where a response was required, a red cross mark was shown if the circle did not intersect the line when it was stopped by a response or if the circle left the screen with no response having been generated and a green check mark appeared if the circle intersected the target line when it was stopped by a response. The feedback was used to encourage participants to respond with accurate timing and minimize tendencies to delay their response to gain more time to make decisions. In trials in which no response was required, a green check mark was displayed if the circle left the screen without a response having been made and a red cross was displayed if any response was generated at all.

Criterion task.

Before the experimental trials began, participants completed two criterion blocks. In these two blocks, all trials required a response, for participants to become familiar with the timing requirement of the response. The meaning of the color used in this task was consistent with that used in subsequent tasks for each individual. In the first and easier criterion block, the moving circle started from the top center of the screen and dropped toward the bottom of the screen with a constant speed, which took 900 ms in total. A white target line was placed 750 ms from the top and thus 150 ms from the bottom. The circle diameter was sized such that it took 120 ms for the circle to move across the target line. The block ended with five consecutive correct responses (i.e., any part of the circle stopped on the line). Participants then performed the second and more difficult criterion block, which matched the conditions of the main experiment, i.e., the diameter of the circle was reduced to 60 ms and it took 500 ms from the trial onset to the center of the circle intersecting the target line and another 120 ms to the bottom of the screen. Similarly, five consecutive correct trials were required to end this block.

After successfully completing these two criterion blocks, participants then performed the main task with a response-to-no response (R-to-NR) condition and a no response-to-response (NR-to-R) condition. Half of the participants performed the R-to-NR condition for six blocks of 100 trials followed by performing the NR-to-R condition for six blocks of 100 trials. The order of these two conditions was reversed for the other half of the participants. The speed estimated for each condition did not depend on the order of which condition was first performed (Supplemental Fig. S3).

R-to-NR condition.

This task is also known as the adaptive or anticipated stop-signal task (3, 10, 11, 18) and has been used to examine how fast participants can decide to cancel a prepared response that was originally planned to be executed. The moving circle always started with the color that cued a response (white for half of participants and black for the other half). In a random ~30% of trials (204 out of 600 trials), the circle turned to the not-responding color while it was moving toward the target line. The time of color switch before the center of the circle intersected the target line was randomly drawn from a uniform distribution between 50 and 500 ms with a step size of 16.7 ms. The choice of this step size was constrained by the refresh rate of the monitor, which was 60 Hz. Thus, there were 28 possible time points at which the circle color changed. The closer the time point was to the target line, the shorter time available to make a decision.

NR-to-R condition.

This task, conceptually similar to the timed-response task commonly used in motor reaching tasks (17, 20), was used to examine how fast participants can initiate a response. Trials started with the circle defaulted to the not-responding color (white for half of participants and black for the other half) and switched to the responding color in a random subset of trials. Consistent with the R-to-NR condition, the proportion of color-switch trials was ~30% of trials (204 out of 600 trials) and the time of color change ranged from 50 to 500 ms. These switch trials and their corresponding color change times were matched between these two conditions on a trial-by-trial basis. One participant showed clear evidence of guessing the required response in the NR-to-R condition (Supplemental Figs. S1 and S2), and we therefore excluded this participant from subsequent analysis.

Data and Statistical Analysis

Speed-accuracy trade-off.

Our primary analysis was focused on the trials in which the color switched in both conditions. We assessed the time course over which participants were able to abort an impending action in the R-to-NR condition by constructing a speed-accuracy trade-off relating the time available to cancel a response and the probability of the response being successfully canceled. For visualization purposes, we estimated this speed-accuracy trade-off with a 50-ms sliding window on the time available to cancel a response (i.e., allowed RT; see below for details). Similarly, the speed-accuracy trade-off can also reveal how rapidly a response can be initiated in the NR-to-R condition.

Response correctness.

In our original speed-accuracy trade-off analysis (Fig. 2), a trial in which a response was required was considered to be successful only if a response was made while the circle was not below the target line (i.e., within a 30-ms time window below the target line). In practice, participants often did generate a response but did so after the circle was no longer overlapping the line. We designated these trials as failures and considered them to be equivalent to not generating a response at all. However, we designated a trial that does not require a response as correct if participants did not press the button before the trial ended. This nonproduced response could be >30 ms after the target line if it were generated.

This designation of correctness included a strong requirement of timing accuracy when initiating a response in the NR-to-R condition but not when canceling a response in the R-to-NR condition (Fig. 2*C*). This asymmetry in the analysis may have affected our estimation of the relative timing of

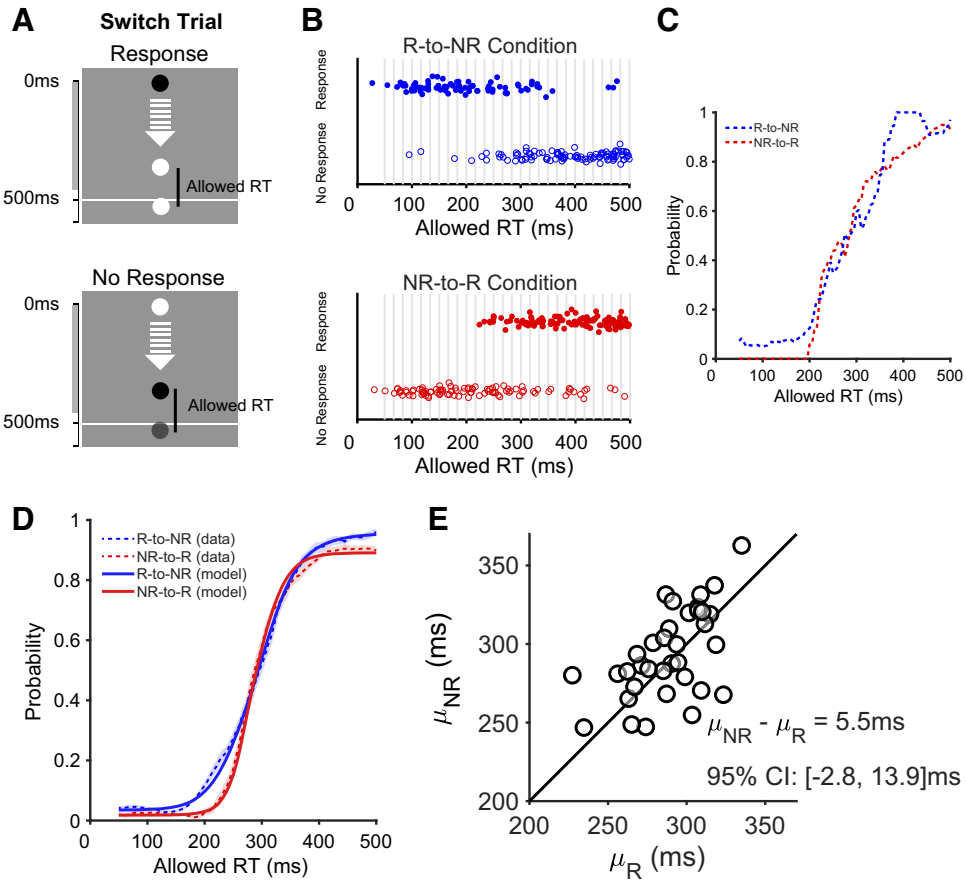

Figure 2. Canceling an impending response is not faster than initiating one. *A*: the allowed reaction time (RT) was quantified as the time elapsed from color change to the time of the button press when participants generated a response. When no response occurred, the allowed RT was approximated as the time interval between the color change to the typical time of button presses in comparable trials (see MATERIALS AND METHODS for more details). *B*: behavior of 1 exemplar participant. In trials in which only a very short allowed RT was allowed, this participant consistently made the wrong choice as to whether to respond or not. When a longer allowed RT was allowed, this participant was able to consistently make the correct choice to respond or not. Vertical jitter was added to allow individual data points to be seen more easily. NR-to-R, no response-to-response; R-to-NR, response-to-no response. *C*: the raw data were used to construct speed-accuracy trade-offs, showing the probability of a correct choice as a function of allowed RT. *D*: mean speed-accuracy trade-offs for each condition across all participants (dashed lines) were well captured by a computational model (solid lines) in which we assumed that the decision to respond or not could be thought of as a discrete event occurring at a random time $T \sim \mathcal{N}(\mu, \sigma^2)$ after the circle changes color. *E*: we used this model to estimate the average time needed to cancel an intended response μ_{NR} in the R-to-NR condition and the average time to initiate a response μ_R in the NR-to-R condition. Across all participants, μ_{NR} (294.5 ± 28.3 ms; mean ± SD) was comparable to μ_R (289.0 ± 24.4 ms; mean ± SD; equivalence and noninferiority test). 95% CI, 95% confidence interval.

going and stopping. To match the timing requirement, we further used a 35-ms time window below the target line. Results from these two different criteria for correctness are consistent with one another (Supplemental Fig. S4 and RESULTS).

Allowed RT.

By manipulating the time at which the circle changed its color (between 50 and 500 ms before the targeted line), we forced participants to cancel an intended response in the R-to-NR condition or press a button in the NR-to-R condition within a particular amount of time, referred to as allowed RT. In the R-to-NR condition, when participants failed to cancel an impending response the allowed RT was quantified as the time elapsed from color change to the time of the button press. When no response was generated, the actual allowed RT was not observable and so instead the allowed RT was approximated by the intended allowed RT, i.e., the time interval between the center of the circle and the target line at the moment of color change. In the NR-to-R condition, the allowed RT was calculated as the time interval between color change and the time of response if participants pressed a button before or when the circle reached the target line, whereas it was approximated as the interval between color change and the target line (i.e., intended allowed RT) if participants did not generate a response or the response was made later than the target line (i.e., the circle was no longer overlapping the target line).

When using the intended RT described above, we inherently assumed that the not-produced response would have had accurate timing (i.e., exactly 500 ms after the trial started) if it was produced, which, however, is not true in reality. All participants had an idiosyncratic tendency to respond consistently earlier or later than the target line (Supplemental Fig. S5). To better approximate the true RT,

we first calculated how much later or earlier each participant responded to trials in nonswitch trials in which the circle was in the responding color throughout the trial (i.e., 396 out of 600 trials, ~70% in the R-to-NR condition). From these measurements, we randomly drew a sample and added it to an intended RT to approximate the unobservable true allowed RT. We repeated this bootstrapping process 1,000 times for each individual and modeled the mean speed-accuracy trade-off.

Speed-accuracy trade-off model.

To quantify the speed-accuracy trade-off, we assumed that canceling an intended response in the R-to-NR condition occurred at a random time $T_{NR} \sim \mathcal{N}(\mu_{NR}, \sigma^2_{NR})$. A response would be correctly aborted with a probability α_{NR} (close to 0) if the available allowed RT was shorter than T_{NR}. Similarly, a response would be correctly aborted with probability β_{NR} (close to 1) if allowed RT was longer than T_{NR}. A value of β_{NR} that is less than 1 indicates that a response may not be correctly stopped even given long allowed RT, reflecting "trigger failure" where the stop process is never initiated because of participants being inattentive to the signal. The presence of trigger failure could yield overestimations of the action inhibition speed (6, 21, 22), because the speed was often calculated with respect to the point where the probability of responding (or not responding) was at 0.5 (23), with an implicit assumption that β_{NR} was 1. Instead, our model estimated the center of the speed-accuracy trade-off (i.e., not necessarily at 0.5 accuracy level if $\beta_{NR} < 1$), so it took the trigger failure effect into account to estimate the stopping speed μ_{NR}, confirmed by parameter recovery analysis (Supplemental Fig. S6).

According to our model, the probability, in trial i, of observing a correct response cancellation ($c = 1$), given the preparation time (t^i) is given by

$$\begin{aligned} P_{NR}(c^i|t^i) &= \alpha_{NR} P(t^i \le T_{NR}) + \beta_{NR} P(t^i > T_{NR}) \\ &= \alpha_{NR}\left[1 - \Phi_{NR}\left(t^i|\mu_{NR}, \sigma^2_{NR}\right)\right] + \beta_{NR}\Phi_{NR}\left(t^i|\mu_{NR}, \sigma^2_{NR}\right) \end{aligned}$$

where $\Phi_{NR}\left(t^i \,|\, \mu_{NR}, \sigma^2_{NR}\right)$ is the cumulative normal distribution of T_{NR}.

Similarly, in the NR-to-R condition, the probability of correctly initiating a response given the preparation time (t^i) is

$$\begin{aligned} P_R(c^i|t^i) &= \alpha_R P\left(t^i \le T_R\right) + \beta_R P\left(t^i > T_R\right) \\ &= \alpha_R\left[1 - \Phi_R(t^i)\right] + \beta_R \Phi_R(t^i) \end{aligned}$$

where $\Phi_R(t^i) = \Phi_R\left(t^i|\mu_R, \sigma^2_R\right)$ is the cumulative normal distribution of T_R.

We estimated the parameters μ, σ, α, β for both conditions, using maximum likelihood estimation with the MATLAB function fmincon. Parameter recovery analysis showed that the parameter of interest, μ, could be reliably estimated from our model (Supplemental Fig. S7).

Modeling the time of response.

Time of response was the time elapsed from the trial onset to the time at which the response was made, if any. In our speed-accuracy trade-off analysis, we relied on a proxy allowed RT for trials in which a response was not produced. To avoid this reliance, we also estimated the speed of making a response in the NR-to-R condition by fitting the time of response, y, with two linear functions, which intersected at an intended allowed RT of μ_T:

$$y = \begin{cases} \mu_0 + \beta_1(t - \mu_T) + \mathcal{M}(0, \sigma^2_{rt}, \delta), & t < \mu_T \\ \mu_0 + \beta_2(t - \mu_T) + \mathcal{N}(0, \sigma^2_r), & t \ge \mu_T \end{cases}$$

In this model, we took the time of transition between these components, μ_T, as the minimum time at which participants could initiate a response as required after the color change. For intended allowed RT $t \ge \mu_T$, we assumed that participants could successfully time the response so that the button press would occur as the circle crossed the target line. In this case, the observed time of response would follow a Gaussian distribution centered on the required time of response, μ_0 (~500 ms). When intended allowed RT $t < \mu_T$, we assumed that participants behaved in a reactive manner to the appearance of stimulus and that therefore the time of response would increase as a function of $t - \mu_T$. We observed that the timing of these late responses followed a typical reaction time distribution, with participants occasionally exhibiting unusually late responses. We therefore assumed that the residual term in the upper equation followed an exponentially modified Gaussian distribution $\mathcal{M}(0, \sigma^2_{rt}, \delta)$, a long-tailed distribution that is commonly used to describe reaction time distributions. This choice could still accommodate participants who did not generate occasional very late responses since $\mathcal{M}$ approaches a Gaussian distribution when δ is close to zero.

Although this model resembled participants' behavior in the NR-to-R condition quite closely, we also observed that participants' behavior would deviate from the pattern described by the model on a small subset of outlying trials. We assumed that these trials either reflected *1*) lapses in which participants may have been inattentive to the stimulus so that even when the intended allowed RT was quite long they responded only very late (i.e., after the stimulus disappeared from the screen; the time of response > 620 ms) or *2*) anticipatory response in which participants correctly guessed that they would need to press the button despite the stimulus changing color very late. Including these trials in our analysis could significantly influence the model fit, as demonstrated by the parameter recovery analysis (Supplemental Fig. S10). We therefore sought to exclude such trials from our analysis by excluding trials with very late responses despite a very long intended RT (lapses) or correctly timed trials with a very low intended RT (anticipatory guesses). We determine the cutoffs of the long enough and short enough intended allowed RT on an individual basis based on the time point at which the performance accuracy reached within 5% of the upper asymptotic of the speed-accuracy trade-off and the time point at which the performance accuracy increased above the lower chance level of the speed-accuracy trade-off by 5%, respectively (see *Modeling speed-accuracy trade-off*). There were 10 out of 7,770 trials where participants made a late response given long allowed RT (lapses) and 59 out of 7,770 trials where participant anticipated to respond (anticipatory guess). The anticipation trials primarily occurred in four participants.

Parameters μ_0, β_1, β_2, μ_T, σ_{rt}, σ_r, δ were estimated by maximum likelihood estimation with MATLAB function fmincon.

To avoid a local minimum estimation, we ran the maximum likelihood estimation with 100 random starting values. A parameter recovery analysis indicated that our model fitting yielded unreliable estimation of true parameters (Supplemental Fig. S8). We found, based on parameter recovery, that it was better to constrain σ_{rt} to be >0.005, to avoid poor-quality fits. For the same reason, we also regularized the fits by penalizing the log-likelihood with

$$LL* = LL - \gamma(\sigma_{rt} - 0.03)^2 - \gamma(\delta - 0.03)^2 - \gamma(\mu_0 - 0.5)^2$$

σ_{rt} and δ were included to avoid unrealistic estimation of $\sigma_{rt} \approx 0$ and $\delta \approx 0$. We chose 0.03 for σ_{rt} and δ, as it was the mean value of initial estimation across participants. We also regularized μ_0 and set it to 0.5 s because our data showed that participants tended to respond around the target line given a long enough RT (Supplemental Fig. S5). We set γ = 2,000, which avoided overfitting these three parameters to the particular value we selected. Parameter recovery demonstrated that this regularized fitting procedure led to reliable estimation of the true parameters when applied to synthetic data (Supplemental Fig. S9).

Statistical analysis.

Data (e.g., μ_{NR} vs. μ_R) were analyzed by paired *t* test at the significance level of $\alpha = 0.05$ after examining the normality of samples. For this test, we reported the 95% $(1 - \alpha)$ confidence interval (CI). Because nonsignificant outcomes from hypothesis testing do not necessarily support that two samples are not different from one another, we further conducted the equivalence and noninferiority test (24, 25), which allows us to statistically test the equivalence between μ_{NR} and μ_R. Under this test, the null hypothesis is that μ_{NR} and μ_R and are different, whereas they are assumed to be equivalent under the alternative hypothesis. Power analysis indicates that the sample size n = 35 (out of 36) had 80% power to detect an effect size of 0.7 between conditions (24). Therefore, we set the upper and lower bounds of the equivalence and noninferiority test as 0.7 and −0.7, corresponding to the equivalence bound between 16.5 ms and −16.5 ms in the scale of reaction time. Importantly, the significant result from the equivalence test does not suggest that the difference between μ_{NR} and μ_R is exactly zero. Instead, their equivalence was considered within the equivalence margin (i.e., −16.5 ms to 16.5 ms). As such, $\mu_{NR} - \mu_R$ could be significantly different from zero based on a paired *t* test (e.g., difference = 8 ms, which might not be biologically meaningful regarding human reaction time), while this difference is still bounded within the equivalence margin. For the equivalence test, we reported the 90% confidence interval, which yields a α = 0.05 significance level for testing equivalence (25), matching the same significance level with the paired *t* test we used.

RESULTS

Participants performed two tasks designed to assess how quickly they could generate or cancel a movement in response to an unexpected cue. Participants viewed a circle moving vertically downward to cross a horizonal line (Fig. 1) and were instructed to either press a button when the circle overlapped the target line or do nothing, depending on the color of the circle as it crossed the line. By varying the initial color of the circle and unexpectedly changing the color during the trial (switch trials; 30% of all trials), we created two separate conditions to assess the speed of action initiation and action inhibition, respectively: In the Response-to-No-Response (R-to-NR) condition, participants needed to rapidly abort an initially prepared response [similar to anticipated stopping paradigms (3, 6, 11, 18)]; in the No-Response-to-Response (NR-to-R) condition, participants needed to rapidly initiate a response that they had not originally intended to.

In switch trials of the R-to-NR condition, participants had to abort an impending response. The amount of time available to participants to inhibit their response, which we term the allowed RT (Fig. 2*A*; MATERIALS AND METHODS), was determined by the time at which the circle changed color, which varied from 50 to 500 ms before crossing the target line. The performance of one exemplar participant is shown in Fig. 2*B*, *top*. When the allowed RT was very short (<100 ms), this participant almost always failed to abort their prepared response, suggesting that making a response is prepotent (see also Supplemental Fig. S2). However, at longer allowed RTs (>300 ms), this participant was able to correctly cancel the response in almost all trials. At intermediate allowed RTs, the participant was sometimes successful and sometimes unsuccessful in canceling their response. This behavior was broadly consistent with previous findings in the stop-signal task (6, 8, 10, 11, 18).

In switch trials of the NR-to-R condition, participants had to rapidly generate a response. Behavior in this condition was similar but complementary to that in the R-to-NR switch trials (Fig. 2*B*, *bottom*). The same participant failed to initiate a timely response if the circle changed color shortly before it crossed the line (<100 ms allowed RT), but when allowed a longer RT to react to the color change (>300 ms allowed RT) they always correctly generated a response as the circle crossed the target line.

From this raw response data, we constructed a speed-accuracy trade-off for each condition (R-to-NR and NR-to-R), based on a 50-ms sliding window on the allowed RT. This trade-off function describes the probability of correctly aborting a response (R-to-NR) or the probability of generating a well-timed response (NR-to-R) as a function of allowed RT (Fig. 2*C*). For the exemplar participant in Fig. 2*B*, the centers of the two speed-accuracy trade-off functions were both located at ~280 ms, indicating that the average times required to cancel a response (R-to-NR condition) or initiate a response (NR-to-R) were comparable. Averaged behavior across all participants (n = 35 out of 36; MATERIALS AND METHODS; Supplemental Figs. S1 and S2) showed the same pattern (Fig. 2*D*, dashed lines), with comparable RTs required to initiate a response and to cancel the initiation of a response.

To more precisely quantify how quickly participants could cancel a response in the R-to-NR condition, we considered a simple model, similar to the classic race model of stopping behavior (4) and related to prior work on movement initiation (17), in which we assumed that the cancellation of a response was a discrete event occurring at a random time $T_{NR} \sim \mathcal{N}(\mu_{NR}, \sigma^2_{NR})$ after the circle changed color. On a given trial, if the time needed to cancel a response (T_{NR}) was shorter than the allowed RT, participants would successfully avoid generating a response. But if the required time, T_{NR}, was longer than the allowed RT, participants would fail to

cancel the impending response. This led to a predicted probability of being correct that increased smoothly as a function of allowed RT, as observed in the data. An analogous model was applied to the NR-to-R condition. In this case, the commitment to initiating a response was assumed to be made at a random time $T_R \sim \mathcal{N}(\mu_R, \sigma_R^2)$ after the circle changed color. Note that this event did not determine the timing of the response, which we assume was determined by a separate process and was synchronized with the motion of the circle. On a given trial, the participant would succeed at correctly generating the response at the right time only if T_{NR} was shorter than the allowed RT on that trial. Thus, μ_{NR} and μ_R represented the average times required either to cancel a movement or initiate a movement, respectively.

We fitted these models to each participant's data via maximum likelihood estimation, yielding model fits that closely matched the empirical data (Fig. 2*D*, solid lines; see Supplemental Fig. S1 for individual fitting). By comparing the estimates of μ_{NR} to μ_R across participants, we found that the time required to cancel the initiation of an impending response (μ_{NR}: 294.5 ± 28.3 ms; mean ± SD) and the time required to initiate a response (μ_R: 289.0 ± 24.4 ms; mean ± SD) were not significantly different from one another [$\mu_{NR} - \mu_R$: 5.5 ± 24.4 ms, Cohen's $d = 0.207$, $t(34) = 1.34$, paired t test, $P = 0.19$, 95% CI: [−2.8 ms, 13.9 ms]; Fig. 2*E*] and were also highly correlated across participants ($\rho = 0.58$, $P < 0.001$). The equivalence and noninferiority test demonstrated that these two speeds were statistically equivalent ($P_{lower} < 0.001$, $P_{upper} = 0.004$; 90% CI: [−1.5 ms, 12.5 ms]). These speeds were not depending on whether the NR-to-R or R-to-NR blocks were completed first and did not change differentially with practice (Supplemental Fig. S3). Thus, when we compared action initiation and action inhibition under experimental conditions that were as closely matched as possible, we found that generating a response takes the same amount of time as canceling one.

Because μ_{NR} and μ_R represent the RT at which performance accuracy reaches the center of the speed-accuracy trade-off function, their estimates were sensitive to exactly how we defined whether a trial was performed correctly or not. In our initial analysis, a late response (if the circle was below the target line without overlapping it at the time the response was made) of an NR-to-R trial was considered as incorrect, whereas in R-to-NR trials it was impossible to distinguish between trials in which a response was canceled at the correct time and trials in which a response was delayed before being canceled. This unavoidably imposed a stronger requirement of timing precision on initiating a response than canceling a response [asymptotic accuracy: $\beta = 0.89 \pm 0.06$ vs. $\beta = 0.96 \pm 0.05$; $t(34) = -5.26$, paired t test, $P < 0.0001$, 95% CI: [−0.09, −0.04]]. However, this aspect of the analysis did not greatly bias our estimate of the time course of action initiation. When we relaxed the timing requirement for correct responding in the NR-to-R condition to match the asymptotic accuracy in the R-to-NR condition [0.94 ± 0.06 vs. 0.96 ± 0.05; $t(34) = -1.9$, paired t test, $P = 0.07$, 95% CI: [−0.04 ms, 0.01 ms]], the resulting estimate of μ_R (286.3 ± 23.6 ms) was only 2.7 ms different from the original analysis and was only 8.2 ms shorter than μ_{NR} [$\mu_{NR} - \mu_R$: 8.2 ± 23.5 ms, Cohen's $d = 0.31$, $t(34) = 2.06$, paired t test, $P = 0.047$, 95% CI: [0.0 ms, 16.3 ms], power = 0.25; Supplemental Fig. S3]. Although there was a statistical difference between μ_{NR} and μ_R after the upper asymptotic accuracies were matched, the difference of only 8.2 ms may not be large enough to be meaningful at the scale of typical human reaction times. Moreover, the statistical power to detect this difference based on our sample size was only 0.25. Indeed, the difference between μ_{NR} and the updated μ_R fell within the equivalence bound between −16.5 ms and 16.5 ms ($P_{lower} < 0.001$; $P_{upper} = 0.009$, 90% CI: [1.5 ms, 14.9 ms]). Therefore, our conclusion about action initiation and action inhibition having a similar time course was not strongly biased by the presence of late responses.

Similar to other analyses of behavior in stop-signal tasks, the preceding analyses relied on designating a proxy RT for trials where a response was never generated, and it is possible that this might have influenced the results. We considered an alternative analysis of behavior in the NR-to-R condition that only included trials in which a response was generated and was based on the actual time of those responses as a function of the intended allowed RT (i.e., time interval between color change and the target line). We observed that when the intended allowed RT was long (e.g., >350 ms), participants were often able to make a response around the target line. However, when the intended allowed RT was short (e.g., <250 ms) participants often did generate a response, but later than required by the task: The time of response kept increasing with shorter intended allowed RT (Fig. 3*A*). This suggested that the time of response should exhibit two distinct patterns that were joined at a particular intended RT. We assumed that the timing of the transition between these two distinct patterns coincided with the minimum time at which participants could successfully initiate a movement in response to the color change. To estimate the threshold allowed RT at which this switch occurred, μ_T, we fit a simple model of the timing of participants' responses as a function of intended allowed RT (MATERIALS AND METHODS). The model comprised two linear components: one to represent accurately timed responses at longer intended allowed RTs and one to represent delayed, reactively generated responses at very short intended allowed RTs (Fig. 3*A*, solid light blue line), with the transition between these occurring at μ_T. The estimates of μ_T (301.1 ± 37.1 ms; mean ± SD) obtained with this approach were in close agreement with the original estimate of μ_R based on designative proxy RTs ($\rho = 0.68$; $P < 0.0001$). These alternative estimates of the time needed to initiate a movement also did not differ significantly from the time that participants needed to cancel a previously intended response, μ_{NR}, estimated in the R-to-NR condition [$\mu_{NR} - \mu_T$: −6.57 ± 35.23 ms, Cohen's $d = -0.20$, $t(34) = -1.11$, paired t test, $P = 0.28$, 95% CI: [−18.73 ms, 5.53 ms]; Fig. 3*B*]. Results from the equivalence and noninferiority tests confirmed that these two speeds were equivalent to each other ($P_{lower} = 0.002$; $P_{upper} < 0.001$, 90% CI: [−16.6 ms, 3.6 ms]). This result accords with our initial analysis that action initiation and action inhibition follow a similar time course.

DISCUSSION

By matching experimental conditions as closely as possible between initiating an action and inhibiting an impending action, we found that the time requirements for these two

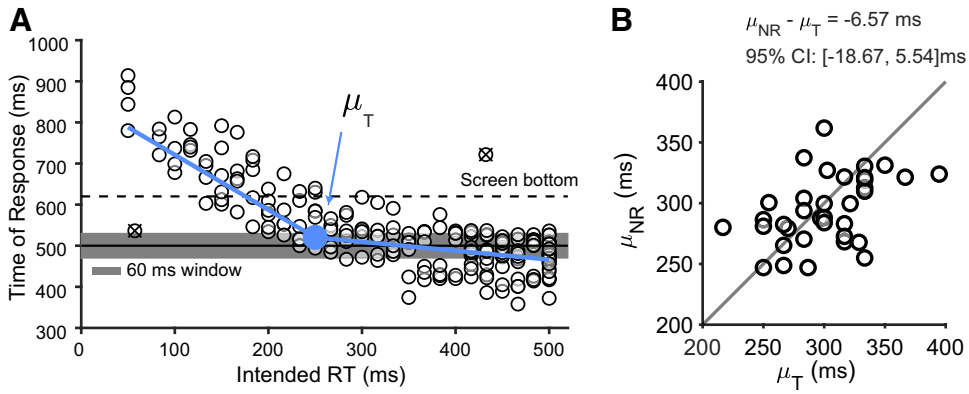

Figure 3. Alternative analysis based on the pattern of delayed responses. *A*: the time of response data as a function of intended reaction time (RT) from the same exemplar participant as in Fig. 2*B*. The time of response was measured as the time interval between trial start and the time at which a response is made. Intended allowed RT was the time interval between color change and the target line. Since this analysis focused only on trials in which a response was actually generated, it did not rely on approximating unobserved RT. We fitted the time of response data with 2 linear functions of intended allowed RT, which intersected at μ_T, a parameter that represents the minimum allowed RT at which an accurately timed response could still be made. *B*: across all participants, the estimates of μ_T (301.1 ± 37.1 ms; mean ± SD), the alternative estimate of time required to initiate a response, were comparable to μ_{NR} (294.5 ± 28.3 ms; mean ± SD), the original estimate of the time required to cancel a response. The crossed circles are illustrative examples of outlier trials, attributed to attentional lapses and anticipatory guesses, which were excluded from the model fits (see MATERIALS AND METHODS). 95% CI, 95% confidence interval.

processes are comparable, both at ~290 ms (note that this included an ~33 ms delay caused by the visual display, which was not accounted for in our calculations).

We used the anticipated stop-signal task (i.e., the R-to-NR condition) to examine action inhibition. Although less widely used than the classic and orthodox stop-signal task (6), the anticipated stop-signal task has been proven to provide reliable estimates of how quickly one can cancel a prepared response (3, 10, 18). Using the anticipated stop-signal task allowed us to compare estimates of the timing of action inhibition to an analogous measure of the timing of action initiation using a timed-response approach. Measuring these two timings in two separate conditions also allowed us to match the design of the experiment as closely as possible across conditions (e.g., trial frequency and trial sequence).

By focusing our analysis on the speed-accuracy trade-off (i.e., the probability of trial success as a function of allowed RT), our estimate of the speed of action inhibition in the anticipated stop-signal task was not influenced by the possibility of trigger failures (i.e., trials in which the stop process is never initiated because of participants being inattentive to the signal), which can bias estimates of the speed of action inhibition in the classic stop-signal task (6). In particular, trigger failures influence only the asymptote but not the slope of the speed-accuracy trade-off in our results (Supplemental Fig. S6) and consequently do not impact our estimates of the speed of action inhibition. Although it may be possible to mitigate the effects of trigger failure through Bayesian analysis approaches (22), the anticipated stop-signal task resolves this problem in a more straightforward manner.

It could still be argued that the results from the anticipated stop-signal task might not be generalizable to action inhibition measured through the classic stop-signal task. However, the time needed to inhibit an action found in the classic stop-signal task (3–5, 10) falls in a range similar to the typical response times identified in simple reaction time studies (1, 2), with both typically ~180–300 ms. Both times can be reduced to ~150 ms if triggered by an unexpected event (17, 26, 27). Moreover, the speed of action initiation and action inhibition are both sensitive to the proportion of trials that require a response (or no response). Action initiation speed becomes faster when there are a greater fraction of trials requiring a response (7, 9), whereas action inhibition becomes more proactive (28) and its speed appears to become faster when more trials require no response (Ref. 29, but see Ref. 30).

It is arguable that distinct preparation processes may be involved between conditions. In particular, preparing the required action in advance in the R-to-NR condition may invoke a preparatory suppression process that prevents the action from being generated too early, whereas in the NR-to-R condition a response might reflect simple reaction without the movement being prepared in advance (31). However, because participants only performed a single action in the experiment, we consider it unlikely that the action would not be prepared beforehand. Nevertheless, it is impossible for us to exactly know the preparatory state of participants in our task. Future investigations that probe the neurophysiological mechanisms involved may help to determine whether or not preparation is indeed the same in these two experimental conditions.

To our knowledge, only one previous study experimentally compared the speeds of action initiation and action inhibition (32). This study had several methodological differences from the work presented here. There was a 50% frequency of switch trials in each condition (as opposed to 30% in our study), which induced greater proactive action initiation and action inhibition. This study also used already well-learned color-response mappings (green = respond vs. red = don't respond) that are likely to lead to faster behavior compared with the arbitrary color-response mappings that we used (33). In addition, a more continuous interceptive movement was required in that task, which may invoke distinct response inhibition mechanisms compared with the simple key press in our task (34–36). Nevertheless, that study found similar speeds

between action initiation and action inhibition, consistent with our results. This consistent result suggests that the similar time course between action initiation and action inhibition might be robust across different types of task variations, including how familiar the task requirements are, the type of action generated, and the degree to which the behavior may be proactive.

Our finding of comparable speeds between reactive action initiation and action inhibition was based on the equivalence and noninferiority test instead of Bayesian hypothesis testing. Bayesian hypothesis testing is considered the "standard" test to detect the absence of an effect. The Bayes factor, however, depends on selected priors that are often unknown and subjective. In our data, the Bayes factor indeed was not stable; how strongly the null hypothesis appeared to be favored depended strongly on the selected priors, although the Bayes factor robustly supported the null hypothesis in all cases. Despite being less widely employed, the equivalence and noninferiority test has long been established (24, 25). Its result does not suggest that the speeds of action initiation and action inhibition are exactly the same. Instead, there was no systematic speed advantage of one process over the other. As an advantage of the equivalence test, these two speeds were concluded to be significantly equivalent within the [−16.5 ms, 16.5 ms] boundary based on our sample size. Future studies with larger sample sizes may be needed to test their equivalence within smaller equivalence intervals. However, it is questionable whether response time differences at the magnitude of a few milliseconds are meaningful with respect to human behaviors and nontrivial noise in measurements.

The fact that action inhibition is not faster than action initiation suggests that to successfully cancel responses on a regular basis, the timing of generating a response would have to be systematically delayed. That is, successfully canceling an action depends not only on the inhibitory processes associated with preventing an action from being generated, the "emergency brake" (27, 37–43), but also on being able to delay the timing at which any response would occur, buying additional time for permitting the possibility of use of the emergency brake. This capacity to decouple our responses to external stimuli from the appearance of the stimulus itself is known as "freedom from immediacy" (Refs. 16, 17; see also Refs. 44, 45). Indeed, slowing of responses has been commonly observed in previous stop-signal tasks (12–14, 46–48). As a result, response times in stop-signal tasks are typically hundreds of milliseconds longer than what we observed in the present study (e.g., Ref. 10) and are accompanied by delayed motor cortex excitability (49). Intriguingly, it has been found that for participants to successfully cancel the response, the response only needs to be delayed by an amount of time slightly greater (~5 to 20 ms) than the time lag between seeing the primary go signal and seeing the stop signal (12). The fact that this adjustment can be so small provides further evidence that action initiation and action inhibition occur over similar timescales.

By measuring action initiation and action inhibition under two separate and symmetric conditions, while tightly controlling the timing of responses, we observed comparable and highly correlated speeds between these two processes. Although it is premature to draw a conclusion solely based on the behavioral data, we speculate that action initiation and action inhibition might not reflect two distinct processes, as is often assumed. Instead, action initiation and action inhibition may be two complementary outcomes of a single process determining the decision about whether to act or not. Indeed, in the context of voluntary movement, the decision of whether or not to act is thought to be an important function that is independent of decisions about how and when to act (16).

If it really is the case that action initiation and action inhibition are more closely related than previously thought, they would share similar neural substrates (e.g., Refs. 50, 51), but this may not be apparent when they are assessed under very different experimental conditions. It has been well established that action selection in humans is interactively controlled through multiple cortico-basal ganglia pathways. The direct and indirect pathways work together to modulate ongoing movement and steer the responding or no-responding decision in a proactive manner (8, 40, 52, 53). Reactive inhibition, as studied in our task and most other stop-signal tasks, is believed to be controlled by the hyperdirect pathway that serves as an emergency brake that can abruptly abort a response (27, 37–43). However, this understanding has largely been established through the lens of conventional stop-signal tasks, which prioritize the importance of stopping a response over that of initiating a response. One possible hypothesis based on our behavioral data is that the hyperdirect pathway may work like a switch, not only serving as an emergency brake but also potentially serving as an emergency accelerator, allowing a response to be initiated reactively even it is not initially planned to be. Beyond the basal ganglia, there also remains a question of whether action initiation follows a similar temporal cascade from the frontal and motor cortices to muscles and behavior that has been identified in action stopping (54–57). We suggest that to better understand the similarity and differences in neural basis of action initiation and action inhibition, it is critical for future research to examine them in tandem, under closely matched experimental conditions, as we have shown here.

Another approach to examining the relationship between action initiation and action inhibition is by comparing them in patient populations. Under closely matched experimental conditions (e.g., NR-to-R vs. R-to-NR), correlated deficits in both action initiation and action inhibition would suggest a general mechanism for deciding whether to act or not to act. It is also possible, however, that an impairment in this process could manifest as an asymmetry between acting and not acting, if the underlying decision becomes biased in a particular direction. In particular, being predisposed toward acting rather than not acting may lead to faster and yet premature action initiation and deficit performance of action inhibition, or vice versa. For example, a recent study using the classic stop-signal task reported impairment of action inhibition in patients with Parkinson's disease, whereas these patients made response ~250 ms faster than healthy peers (58). We hope that our findings and framework can help inform future studies to understand the possible mechanisms underlying how these two facets of action control may be related to one another in both clinical and healthy populations.

DATA AVAILABILITY

All data generated from this study and the code for reproducing the experiment and the results are available at the corresponding author's personal GitHub page: https://github.com/YueDu-Science/Inhibition-Initiation.

SUPPLEMENTAL MATERIALS

Supplemental Figs. S1–S10: https://doi.org/10.6084/m9.figshare.25222553.

ACKNOWLEDGMENTS

Present addresses: A. D. Forrence, Dept. of Psychology, Yale University, New Haven, CT 06511; D. M. Metcalf, Medical College of Georgia, Augusta, GA 30912.

GRANTS

This work was supported by funding from Sunburst Entertainment Group (Y. Du, A. M. Haith).

DISCLOSURES

No conflicts of interest, financial or otherwise, are declared by the authors.

AUTHOR CONTRIBUTIONS

Y.D., A.D.F., and A.M.H. conceived and designed research; Y.D., A.D.F., and D.M.M. performed experiments; Y.D. analyzed data; Y.D. and A.M.H. interpreted results of experiments; Y.D. prepared figures; Y.D. drafted manuscript; Y.D., A.D.F., D.M.M., and A.M.H. edited and revised manuscript; Y.D., A.D.F., D.M.M., and A.M.H. approved final version of manuscript.

REFERENCES

1. **Luce RD.** *Response Times: Their Role in Inferring Elementary Mental Organization*. New York: Oxford University Press, 1991.
2. **Welford A.** Choice reaction time: basic concepts. In: *Reaction Times*. New York: Academic Press, 1980, p. 73–128.
3. **He JL**, **Hirst RJ**, **Puri R**, **Coxon J**, **Byblow W**, **Hinder M**, **Skippen P**, **Matzke D**, **Heathcote A**, **Wadsley CG**, **Silk T**, **Hyde C**, **Parmar D**, **Pedapati E**, **Gilbert DL**, **Huddleston DA**, **Mostofsky S**, **Leunissen I**, **MacDonald HJ**, **Chowdhury NS**, **Gretton M**, **Nikitenko T**, **Zandbelt B**, **Strickland L**, **Puts NAJ.** OSARI, an open-source anticipated response inhibition task. *Behav Res Methods* 54: 1530–1540, 2022. doi:10.3758/s13428-021-01680-9.
4. **Logan GD**, **Cowan WB.** On the ability to inhibit thought and action: a theory of an act of control. *Psychol Rev* 91: 295–327, 1984. doi:10.1037/0033-295X.91.3.295.
5. **Matzke D**, **Strickland LJ**, **Sripada C**, **Weigard AS**, **Puri R**, **He J**, **Hirst R**, **Heathcote A.** Stopping timed actions (Preprint). *PsyArXiv:* 1–45, 2021. doi:10.31234/osf.io/9h3v7.
6. **Verbruggen F**, **Aron AR**, **Band GP**, **Beste C**, **Bissett PG**, **Brockett AT**, **et al.** A consensus guide to capturing the ability to inhibit actions and impulsive behaviors in the stop-signal task. *eLife* 8: e46323, 2019. doi:10.7554/eLife.46323.
7. **Buckolz E**, **Rodgers R.** The influence of catch trial frequency on simple reaction time. *Acta Psychol* 44: 191–200, 1980. doi:10.1016/0001-6918(80)90067-0.
8. **Dunovan K**, **Lynch B**, **Molesworth T**, **Verstynen T.** Competing basal ganglia pathways determine the difference between stopping and deciding not to go. *eLife* 4: e08723, 2015. doi:10.7554/eLife.08723.
9. **Gordon IE.** Stimulus probability and simple reaction time. *Nature* 215: 895–896, 1967. doi:10.1038/215895a0.
10. **Leunissen I**, **Zandbelt BB**, **Potocanac Z**, **Swinnen SP**, **Coxon JP.** Reliable estimation of inhibitory efficiency: to anticipate, choose or simply react? *Eur J Neurosci* 45: 1512–1523, 2017. doi:10.1111/ejn.13590.
11. **Slater-Hammel AT.** Reliability, accuracy, and refractoriness of a transit reaction. *Res Q Am Assoc Health Phys Educ Recreat* 31: 217–228, 1960. doi:10.1080/10671188.1960.10613098.
12. **Lappin JS**, **Eriksen CW.** Use of a delayed signal to stop a visual reaction-time response. *J Exp Psychol* 72: 805–811, 1966. doi:10.1037/h0021266.
13. **Leotti LA**, **Wager TD.** Motivational influences on response inhibition measures. *J Exp Psychol Hum Percept Perform* 36: 430–447, 2010. doi:10.1037/a0016802.
14. **Verbruggen F**, **Logan GD.** Response inhibition in the stop-signal paradigm. *Trends Cogn Sci* 12: 418–424, 2008. doi:10.1016/j.tics.2008.07.005.
15. **Verbruggen F**, **Logan GD.** Proactive adjustments of response strategies in the stop-signal paradigm. *J Exp Psychol Hum Percept Perform* 35: 835–854, 2009. doi:10.1037/a0012726.
16. **Brass M**, **Haggard P.** The what, when, whether model of intentional action. *Neuroscientist* 14: 319–325, 2008. doi:10.1177/1073858408317417.
17. **Haith AM**, **Pakpoor J**, **Krakauer JW.** Independence of movement preparation and movement initiation. *J Neurosci* 36: 3007–3015, 2016. doi:10.1523/JNEUROSCI.3245-15.2016.
18. **Coxon JP**, **Stinear CM**, **Byblow WD.** Intracortical inhibition during volitional inhibition of prepared action. *J Neurophysiol* 95: 3371–3383, 2006. doi:10.1152/jn.01334.2005.
19. **Salinas E**, **Stanford TR.** The countermanding task revisited: fast stimulus detection is a key determinant of psychophysical performance. *J Neurosci* 33: 5668–5685, 2013. doi:10.1523/JNEUROSCI.3977-12.2013.
20. **Ghez C**, **Favilla M**, **Ghilardi MF**, **Gordon J**, **Bermejo R**, **Pullman S.** Discrete and continuous planning of hand movements and isometric force trajectories. *Exp Brain Res* 115: 217–233, 1997. doi:10.1007/pl00005692.
21. **Band GP**, **Van Der Molen MW**, **Logan GD.** Horse-race model simulations of the stop-signal procedure. *Acta Psychol (Amst)* 112: 105–142, 2003. doi:10.1016/s0001-6918(02)00079-3.
22. **Matzke D**, **Love J**, **Heathcote A.** A Bayesian approach for estimating the probability of trigger failures in the stop-signal paradigm. *Behav Res Methods* 49: 267–281, 2017. doi:10.3758/s13428-015-0695-8.
23. **Verbruggen F**, **Chambers CD**, **Logan GD.** Fictitious inhibitory differences: how skewness and slowing distort the estimation of stopping latencies. *Psychol Sci* 24: 352–362, 2013. doi:10.1177/0956797612457390.
24. **Lakens D.** Equivalence tests: a practical primer for t tests, correlations, and meta-analyses. *Soc Psychol Personal Sci* 8: 355–362, 2017. doi:10.1177/1948550617697177.
25. **Walker E**, **Nowacki AS.** Understanding equivalence and noninferiority testing. *J Gen Intern Med* 26: 192–196, 2011. doi:10.1007/s11606-010-1513-8.
26. **Carlsen AN**, **Chua R**, **Inglis JT**, **Sanderson DJ**, **Franks IM.** Prepared movements are elicited early by startle. *J Mot Behav* 36: 253–264, 2004. doi:10.3200/JMBR.36.3.253-264.
27. **Wessel JR**, **Aron AR.** On the globality of motor suppression: unexpected events and their influence on behavior and cognition. *Neuron* 93: 259–280, 2017. doi:10.1016/j.neuron.2016.12.013.
28. **Aron AR.** From reactive to proactive and selective control: developing a richer model for stopping inappropriate responses. *Biol Psychiatry* 69: e55–e68, 2011. doi:10.1016/j.biopsych.2010.07.024.
29. **Castro-Meneses LJ**, **Johnson BW**, **Sowman PF.** The effects of impulsivity and proactive inhibition on reactive inhibition and the go process: insights from vocal and manual stop signal tasks. *Front Hum Neurosci* 9: 529, 2015. doi:10.3389/fnhum.2015.00529.
30. **Zandbelt BB**, **Bloemendaal M**, **Neggers SF**, **Kahn RS**, **Vink M.** Expectations and violations: delineating the neural network of proactive inhibitory control. *Hum Brain Mapp* 34: 2015–2024, 2013. doi:10.1002/hbm.22047.
31. **McInnes AN**, **Lipp OV**, **Tresilian JR**, **Vallence AM**, **Marinovic W.** Premovement inhibition can protect motor actions from interference by response-irrelevant sensory stimulation. *J Physiol* 599: 4389–4406, 2021. doi:10.1113/JP281849.
32. **Marinovic W**, **Plooy AM**, **Tresilian JR.** Preparation and inhibition of interceptive actions. *Exp Brain Res* 197: 311–319, 2009. doi:10.1007/s00221-009-1916-0.

33. **Ceceli AO**, **Myers CE**, **Tricomi E.** Demonstrating and disrupting well-learned habits. *PloS One* 15: e0234424, 2020. doi:10.1371/journal.pone.0234424.
34. **Hervault M**, **Huys R**, **Farrer C**, **Buisson JC**, **Zanone PG.** Cancelling discrete and stopping ongoing rhythmic movements: do they involve the same process of motor inhibition? *Hum Mov Sci* 64: 296–306, 2019. doi:10.1016/j.humov.2019.02.010.
35. **Schultz KE**, **Denning D**, **Hufnagel V**, **Swann N.** Stopping a continuous movement: a novel approach to investigating inhibitory control. *J Cogn Neurosci* 35: 1108–1132, 2023. doi:10.1162/jocn_a_01998.
36. **Morein-Zamir S**, **Nagelkerke P**, **Chua R**, **Franks I**, **Kingstone A.** Inhibiting prepared and ongoing responses: is there more than one kind of stopping? *Psychon Bull Rev* 11: 1034–1040, 2004. doi:10.3758/BF03196733.
37. **Aron AR.** The neural basis of inhibition in cognitive control. *Neuroscientist* 13: 214–228, 2007. doi:10.1177/1073858407299288.
38. **Aron AR**, **Robbins TW**, **Poldrack RA.** Inhibition and the right inferior frontal cortex: one decade on. *Trends Cogn Sci* 18: 177–185, 2014. doi:10.1016/j.tics.2013.12.003.
39. **Chen W**, **de Hemptinne C**, **Miller AM**, **Leibbrand M**, **Little SJ**, **Lim DA**, **Larson PS**, **Starr PA.** Prefrontal-subthalamic hyperdirect pathway modulates movement inhibition in humans. *Neuron* 106: 579–588.e3, 2020. doi:10.1016/j.neuron.2020.02.012.
40. **Dunovan K**, **Verstynen T.** Believer-skeptic meets actor-critic: rethinking the role of basal ganglia pathways during decision-making and reinforcement learning. *Front Neurosci* 10: 106, 2016. doi:10.3389/fnins.2016.00106.
41. **Hannah R**, **Aron AR.** Towards real-world generalizability of a circuit for action-stopping. *Nat Rev Neurosci* 22: 538–552, 2021. doi:10.1038/s41583-021-00485-1.
42. **Jahanshahi M**, **Obeso I**, **Rothwell JC**, **Obeso JA.** A fronto-striato-subthalamic-pallidal network for goal-directed and habitual inhibition. *Nat Rev Neurosci* 16: 719–732, 2015. doi:10.1038/nrn4038.
43. **Wiecki TV**, **Frank MJ.** A computational model of inhibitory control in frontal cortex and basal ganglia. *Psychol Rev* 120: 329–355, 2013. doi:10.1037/a0031542.
44. **Kornblum S.** Simple reaction time as a race between signal detection and time estimation: a paradigm and model. *Percept Psychophys* 13: 108–112, 1973. doi:10.3758/BF03207243.
45. **Ollman RT**, **Billington MJ.** The deadline model for simple reaction times. *Cogn Psychol* 3: 311–336, 1972. doi:10.1016/0010-0285(72)90010-2.
46. **Corneil BD**, **Cheng JC**, **Goonetilleke SC.** Dynamic and opposing adjustment of movement cancellation and generation in an oculomotor countermanding task. *J Neurosci* 33: 9975–9984, 2013. doi:10.1523/JNEUROSCI.2543-12.2013.
47. **Gulberti A**, **Arndt P**, **Colonius H.** Stopping eyes and hands: evidence for non-independence of stop and go processes and for a separation of central and peripheral inhibition. *Front Hum Neurosci* 8: 61, 2014. doi:10.3389/fnhum.2014.00061.
48. **Özyurt J**, **Colonius H**, **Arndt PA.** Countermanding saccades: evidence against independent processing of go and stop signals. *Percept Psychophys* 65: 420–428, 2003. doi:10.3758/bf03194573.
49. **Rawji V**, **Modi S**, **Rocchi L**, **Jahanshahi M**, **Rothwel J.** Proactive inhibition is marked by differences in the pattern of motor cortex activity during movement preparation and execution. *J Neurophysiol* 127: 819–828, 2022. doi:10.1152/jn.00359.2021.
50. **Filevich E**, **Kühn S**, **Haggard P.** Negative motor phenomena in cortical stimulation: implications for inhibitory control of human action. *Cortex* 48: 1251–1261, 2012. doi:10.1016/j.cortex.2012.04.014.
51. **Mostofsky SH**, **Simmonds DJ.** Response inhibition and response selection: two sides of the same coin. *J Cogn Neurosci* 20: 751–761, 2008. doi:10.1162/jocn.2008.20500.
52. **Bariselli S**, **Fobbs WC**, **Creed MC**, **Kravitz AV.** A competitive model for striatal action selection. *Brain Res* 1713: 70–79, 2019. doi:10.1016/j.brainres.2018.10.009.
53. **Mikhael JG**, **Bogacz R.** Learning reward uncertainty in the basal ganglia. *PLoS Comput Biol* 12: e1005062, 2016. doi:10.1371/journal.pcbi.1005062.
54. **Jana S**, **Hannah R**, **Muralidharan V**, **Aron AR.** Temporal cascade of frontal, motor and muscle processes underlying human action-stopping. *eLife* 9: e50371, 2020. doi:10.7554/eLife.50371.
55. **Raud L**, **Thunberg C**, **Huster RJ.** Partial response electromyography as a marker of action stopping. *eLife* 11: e70332, 2022. doi:10.7554/eLife.70332.
56. **Raud L**, **Huster RJ.** The temporal dynamics of response inhibition and their modulation by cognitive control. *Brain Topogr* 30: 486–501, 2017. doi:10.1007/s10548-017-0566-y.
57. **Tatz JR**, **Soh C**, **Wessel JR.** Common and unique inhibitory control signatures of action-stopping and attentional capture suggest that actions are stopped in two stages. *J Neurosci* 41: 8826–8838, 2021. doi:10.1523/JNEUROSCI.1105-21.2021.
58. **Zhong S**, **Choi JW**, **Hashoush NG**, **Babayan D**, **Malekmohammadi M**, **Pouratian N**, **Christopoulos V.** A neurocomputational theory of action regulation predicts motor behavior in neurotypical individuals and patients with Parkinson's disease. *PLOS Comput Biol* 18: e1010111, 2022. doi:10.1371/journal.pcbi.1010111.