

The Countermanding Task Revisited: Fast Stimulus Detection Is a Key Determinant of Psychophysical Performance

Emilio Salinas and Terrence R. Stanford

Department of Neurobiology and Anatomy, Wake Forest School of Medicine, Winston-Salem, North Carolina 27157-1010

The countermanding task is a standard method for assessing cognitive/inhibitory control over action and for investigating its neural correlates. In it, the subject plans a movement and either executes it, if no further instruction is given, or attempts to prevent it, if a stop signal is shown. Through various experimental manipulations, many studies have sought to characterize the inhibitory mechanisms thought to be at work in the task, typically using an inferred, model-dependent metric called the stop-signal reaction time. This approach has consistently overlooked the impact of perceptual evaluation on performance. Through analytical work and computer simulations, here we show that psychophysical performance in the task can be easily understood as the result of an ongoing motor plan that is modified (decelerated) by the outcome of a rapid sensory detection process. Notably, no specific assumptions about hypothetical inhibitory mechanisms are needed. This modeling framework achieves four things: (1) it replicates and reconciles behavioral results in numerous variants of the countermanding task; (2) it provides a new, objective metric for characterizing task performance that is more effective than the stop-signal reaction time; (3) it shows that the time window over which detection of a high-visibility stimulus effectively occurs is extremely short (~ 20 ms); and (4) it indicates that modulating neuronal latencies and the buildup rates of developing motor plans are two key neural mechanisms for controlling action. The results suggest that manipulations of the countermanding task often cause changes in perceptual detection processes, and not necessarily in inhibition.

Introduction

In the countermanding or stop-signal task, a subject is cued to make a movement to a single target. In some trials, the movement is simply executed, but in others a stop signal is given shortly after the go cue and the planned movement must be cancelled (Fig. 1). This paradigm has been widely used to study how voluntary actions are generated and controlled (Colonius et al., 2001; Kornyló et al., 2003; Akerfelt et al., 2006; Jacobson et al., 2011) and to investigate the neural correlates of such processes (Hanes et al., 1998; Ito et al., 2003; Paré and Hanes, 2003; Stuphorn and Schall, 2006). It is also commonly used to characterize abnormal mental conditions that may be related to action inhibition (Armstrong and Munoz, 2003; DeHaan et al., 2007; Lipszyc and Schachar, 2010; Thakkar et al., 2011). For concreteness, here we generally consider the version of the task in which the responses are eye movements.

The conceptual framework adopted by virtually all studies based on this task, or variants thereof, is one in which a motor and a stopping process race against each other toward a threshold (Logan and Cowan, 1984; Boucher et al., 2007; Verbruggen and Logan, 2009). The buildup of motor activity starts first, after the go signal, and the stop-related activity starts later, after the stop signal. If the motor activity reaches threshold first and wins the race, an eye movement is triggered, whereas if the stop-related activity wins, no action is executed; fixation is successfully maintained.

Here we take a different approach. We still consider the programming of a saccade as a buildup in oculomotor activity toward a threshold because there is substantial evidence for this (Hanes and Schall, 1996; Dorris et al., 1997; Lo and Wang, 2006; Brown et al., 2008; Heitz and Schall, 2012), but we make no such assumption about the stopping process. Instead, we simply consider two events: (1) the stop signal is detected with a particular timing and reliability that depend on perceptual mechanisms; and (2) then the motor plan is decelerated such that it slows down and reverses its course (Fig. 2). The idea is that, because a successfully cancelled saccade is one for which the preparatory activity simply does not reach threshold, only two quantities matter: how soon the ramping activity can start slowing down once the stop signal has been given (the detection latency), and how abrupt is the change of course toward baseline (the magnitude of the deceleration). As such, in describing the phenomenology of the task, the perceptual process is crucial because the detection of the stop signal must

Received Aug. 20, 2012; revised Jan. 24, 2013; accepted Jan. 31, 2013.

Author contributions: E.S. and T.R.S. designed research; E.S. performed research; E.S. analyzed data; E.S. and T.R.S. wrote the paper.

This work was supported by the National Institutes of Health (Grants R01EY12389 and R01EY12389-51) from the National Eye Institute, and the National Institute of Drug Abuse as part of the National Science Foundation/National Institutes of Health Collaborative Research in Computational Neuroscience Program (Grant R01DA030750).

The authors declare no competing financial interests.

Correspondence should be addressed to Dr. Emilio Salinas, Department of Neurobiology and Anatomy, Wake Forest School of Medicine, 1 Medical Center Blvd, Winston-Salem, NC 27157-1010. E-mail: esalinas@wakehealth.edu.

DOI:10.1523/JNEUROSCI.3977-12.2013

Copyright © 2013 the authors 0270-6474/13/335668-18\$15.00/0




Figure 1. The countermanding task. *a*, A no-stop trial. The instruction to make a saccade (Go) is the simultaneous disappearance of the fixation spot and appearance of the target. *b*, A correct stop trial (cancelled). When the fixation point reappears (Stop), the subject must withhold the planned eye movement. *c*, An incorrect stop trial (noncancelled). The saccade could not be stopped. The RT is measured from the onset of the go signal to the onset of the saccade. Task difficulty is controlled according to the delay between the go and stop signals (SSD). The rPT is the amount of time available for viewing the stop signal.

occur before any hypothetical inhibitory mechanism can conceivably have an impact.

Two observations motivated us to pursue this approach. First, in trying to explain the results of various task manipulations in terms of their direct effect on an inhibitory plan, previous models have overlooked the impact of the perceptual judgment; and again, any inhibitory surge must obligatorily follow the detection step. Second, the countermanding task shares key characteristics with the compelled-saccade task, a task designed to investigate fast perceptual choices (Stanford et al., 2010). The framework developed here, which derives from that work, is parsimonious, produces a psychophysical metric that has many advantages over the stop-signal reaction time (SSRT), and demonstrates that consideration of perceptual detection mechanisms is both sufficient and necessary for correctly interpreting the results of most experimental manipulations in the countermanding task.

Materials and Methods

The cancellable rise-to-threshold model. All model simulations were performed using Matlab (The Mathworks). The code is available from the authors on request.

The model consists of a variable M that represents the population activity of movement-related neurons (Bruce and Goldberg, 1985) with response fields at the saccadic target (Fig. 2). In each trial, M starts at a value of 0 and then increases linearly toward a threshold of 1000 units.

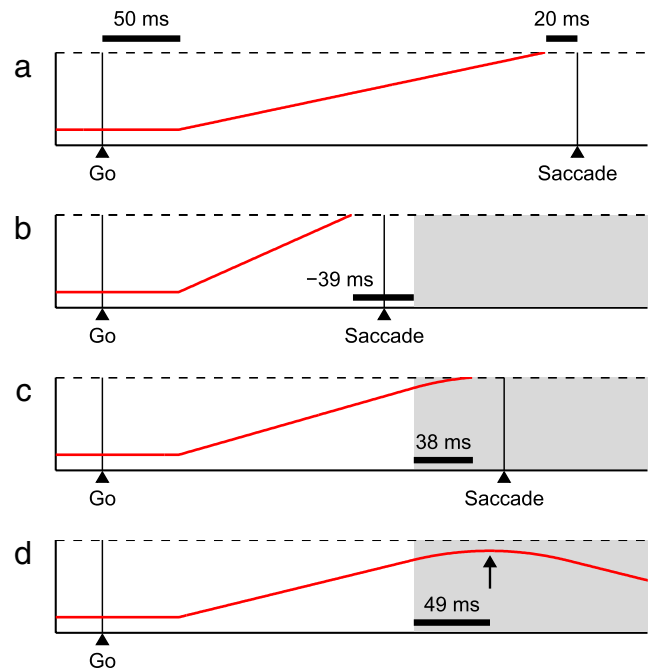


Figure 2. Trial outcomes in the countermanding task, as reproduced by the cancellable rise-to-threshold model. Each panel shows the variable M , which represents oculomotor activity, as a function of time (red traces). In these examples, each motor plan starts 50 ms (afferent delay) after the go and a saccade is produced 20 ms (efferent delay) after the threshold (dotted line) is crossed. The interval between the go and the saccade is the RT. Gray shadows represent the period (after the SSD and the afferent delay) during which the stop signal is known to the model circuit and the motor plan may decelerate. *a*, A no-stop trial. A correct saccade is produced. *b*, A noncancelled trial. The stop signal is given, but a saccade is produced before the signal has any effect on the model circuit. *c*, Another noncancelled trial. The motor plan decelerates slightly, but the incorrect saccade cannot be prevented. *d*, A cancelled trial. The motor plan decelerates and turns back toward baseline. The arrow indicates the point of maximum buildup, T_{max} . *b–d*, The SSD is 150 ms; horizontal bars and numbers indicate the ePT (Eq. 7).

When M crosses the threshold, the circuit commits to a motor action and a saccade is inevitable. The onset of the eye movement is considered to occur an efferent delay Δ_E after threshold crossing. Δ_E was set to 20 ms (Brown et al., 2008); the actual number has essentially no impact on the results.

Each simulated trial proceeds as follows: The go signal (target onset) is given at $t = 0$; and as soon as this signal reaches the model oculomotor circuit, which happens after an afferent delay Δ_T , M starts increasing (i.e., the motor plan starts developing). The rate at which M increases, the buildup rate r_B , is constant and random; in each trial, it is drawn from a Gaussian distribution with mean μ_B and SD σ_B . During this constant buildup stage, M evolves according to the following:

$$\frac{dM}{dt} = r_B \quad (1)$$

$$r_B = r_B^0 \quad (2)$$

where r_B^0 is the buildup rate drawn initially. In trials in which M reaches threshold during this stage (all no-stop trials and some noncancelled trials), a saccade is produced and a reaction time (RT) is recorded (Fig. 2*a,b*). Otherwise, the oculomotor plan changes as prescribed by the above equations until the information about the stop signal arrives.

The stop signal is presented SSD (stop-signal delay) milliseconds after the go signal, at $t = SSD$, and it reaches the model circuit at $t = SSD + \Delta_S$, where Δ_S is the afferent delay of the stop signal. At this time, the activity variable M starts decelerating: the buildup rate itself decreases at a con-

stant rate, approaching a negative value r_{DN} , which is the maximum build-down rate. The corresponding equations are as follows:

$$\frac{dM}{dt} = r_B \quad (3)$$

$$\frac{dr_B}{dt} = \frac{r_{DN} - r_B^0}{\tau} \quad (4)$$

with the added rule that, once the buildup rate reaches its new target value (i.e., once r_B is equal to r_{DN}), it stops changing, so the last derivative becomes zero. In this way, the deceleration is constant but lasts a finite maximum amount of time, which is precisely equal to τ . Two things can happen during this deceleration stage: either M turns around and never reaches threshold (the trial is cancelled; Fig. 2*d*), or even though it slows down, M still reaches threshold (the trial is noncancelled; Fig. 2*c*). So, two parameters determine the dynamics of the turnaround in activity: r_{DN} , which is the most negative slope of the firing rate when it ramps downward, and the time constant τ , which determines how fast that negative rate is reached.

For simplicity, the maximum build-down rate r_{DN} was set to -8.0 U/ms throughout; in this way, it was equal in magnitude to the average buildup rate within the standard set of parameters ($\mu_B = 8.0$ U/ms) but had the opposite sign. The actual value had only a minimal impact on the results. Indeed, we tested various additional functional forms for Equation 4 that did not place a limit on the maximum build-down rate (dr_B/dt proportional to the square of r_B^0 , proportional to r_B^0 , or equal to a constant, with variable durations) and with appropriate proportionality coefficients, the results were nearly identical to those obtained with Equation 4. Thus, the model is quite insensitive to the specific dynamics of the deceleration.

There are two afferent delays in the model: one for the target/go signal (Δ_T) and another for the stop signal (Δ_S). Allowing them to be different not only produced better fits to the data but also was crucial for experiments in which the target and the stop correspond to stimuli of different modalities. We assume that the afferent delays vary across trials and independently for the go and stop signals. The variability is Gaussian, but such that delays below a minimum amount (20 ms) are not allowed. The mean delays are μ_T and μ_S , and the SDs of the Gaussians are σ_T and σ_S .

Finally, although subjects generally do not make errors when they have ample time to detect the stop signal, under some circumstances such errors do happen. This was most evident in one particular experiment discussed below. To account for such lapses, in each simulated stop trial we consider a probability p_e of making an unconditional error (i.e., of missing the stop signal altogether, regardless of its timing). In those cases, M keeps increasing toward threshold following Equations 1 and 2 as if no stop was presented. In practice, the parameter p_e uniquely sets the maximum value of the tachometric function discussed below.

In all, the cancellable rise-to-threshold model has eight free parameters: μ_B , σ_B , μ_T , σ_T , μ_S , σ_S , τ , and p_e . The model produces different RTs and outcomes (cancelled/noncancelled) from one trial to another primarily because a different initial buildup rate (r_B^0) is drawn each time; the only other quantities that also change stochastically across trials are the afferent delays, but their variabilities (σ_T , σ_S) have a more modest influence on the results.

Model fitting The free parameters of the rise-to-threshold model were first optimized to minimize the mean absolute error between the simulated data and the monkey data published by Boucher et al. (2007). Five psychophysical curves were used to constrain the model parameters: the psychometric curve (percentage correct vs SSD) and four cumulative RT distributions, three for noncancelled trials and one for no-stop trials (Fig. 2 of Boucher et al., 2007). Thus, the error function had the form

$$E = \sum_{i,j} w_i \frac{|e_{ij} - m_{ij}|}{n_i} \quad (5)$$

where e and m are experimental and model values, respectively; the index $i = 1, 2, \dots, 5$ identifies each of the five psychophysical curves; the index j runs through each point in a curve; the factor n_i is the number of non-

Table 1. Model parameter values used

Figure/condition	μ_B	σ_B	μ_T	σ_T	μ_S	σ_S	τ	p_e
Figs. 3, 5, 12	8.0	2.4	95	40	57	12	52	0.011
Fig. 4	8.0	2.4	95	40	57	12	52	0
Fig. 7, no gap; Fig. 11	7.8	1.9	110	30	84	12	52	0.08
Fig. 7, yes gap	8.7	2.2	67	30	41	12	52	0.08
Fig. 8, 0.10 stop	5.5	0.92	117	53	56	13	34	0.016
Fig. 8, 0.69 stop	3.4	0.20	117	53	56	13	34	0.016
Fig. 9, visual stop	5.1	2.6	96	50	83	16	38	0.05
Fig. 9, auditory stop	5.1	2.6	96	50	32	16	38	0.49

All afferent delays and τ are in ms. Build-up rates are in U/ms, where the distance from baseline to threshold is 1000 U.

zero points in curve i ; and w_i weights the contribution of each curve (so $\sum_i w_i = 1$). Best-fitting parameter values were found by exhaustive search; that is, by establishing a multidimensional grid where each point is a different combination of parameters, running the model with each combination, and selecting the one that minimized the above error.

The data included in other experimental studies were typically less complete. In those cases, it was relatively easy to find ranges of values that replicated the results accurately, and the final values were chosen so that they deviated relatively little from those that fitted the Boucher et al. (2007) data. The parameter values used in each figure are shown in Table 1.

Predicted neuronal activity. Using the cancellable-rise-to-threshold model, we generated average activity traces to (1) verify that these were consistent with previous neurophysiological recordings, and (2) make predictions about the neuronal responses expected under various experimental conditions. For this, several thousand trials were run with a fixed set of parameter values and the trajectory of the variable M , which represents the oculomotor plan, was saved in all of them. Predicted neuronal responses were generated by separating the simulated trials into appropriate groups (no-stop, cancelled, etc.), synchronizing each trace either on the go signal, on the stop, or on the saccade, and averaging the M trajectories for each group. This is the standard approach for characterizing real responses recorded from oculomotor cells (e.g., Thompson et al., 1996; Port and Wurtz, 2009; Stanford et al., 2010).

Whenever an eye movement is performed, recordings from saccade-related neurons typically show a quick, stereotypical drop in activity shortly after saccade onset. This fall-off partially determines the resulting mean traces when the data are aligned on the go or the stop signals because trials with a wide range of RTs are included in the averages. So, a fixed tail was appended to each simulated trace after threshold crossing to take this into account (Ratcliff et al., 2007). The tail had no impact whatsoever on the outcomes or RTs; it simply produced more realistic average traces. It was modeled as a falling exponential function with a fixed time constant of 90 ms.

The resulting mean traces were consistent with previous single-unit studies. When aligned on the go, the buildup rate was largest for noncancelled trials, intermediate for no-stop trials, and smallest for cancelled trials (Fig. 3*a*), in accordance with differences in RT obtained across conditions (see below). Differences between cancelled and noncancelled responses were also evident when aligned on the stop signal: in the former the average activity peaked near the stop, whereas in the latter the peak occurred significantly later (Fig. 3*b*). These results agree with single-neuron recordings in the frontal eye field (FEF) and superior colliculus showing that saccade-related activity indeed turns around toward baseline during successfully cancelled trials (Hanes et al., 1998; Paré and Hanes, 2003; Boucher et al., 2007; Ray et al., 2009). A difference in buildup rate between no-stop and noncancelled trials was also seen with the data aligned on the saccade (Fig. 3*c*) and is consistent with the observation that presaccadic activity in the FEF builds up faster before short-latency saccades than before long-latency saccades (Hanes and Schall, 1996).

The ideal tachometric curve. Based on the model responses, it is possible to compute an “ideal” tachometric curve, which plots the subject’s fraction of correct responses (stopped saccades) as a function of the amount of time available to process the relevant sensory stimulus (the stop sig-

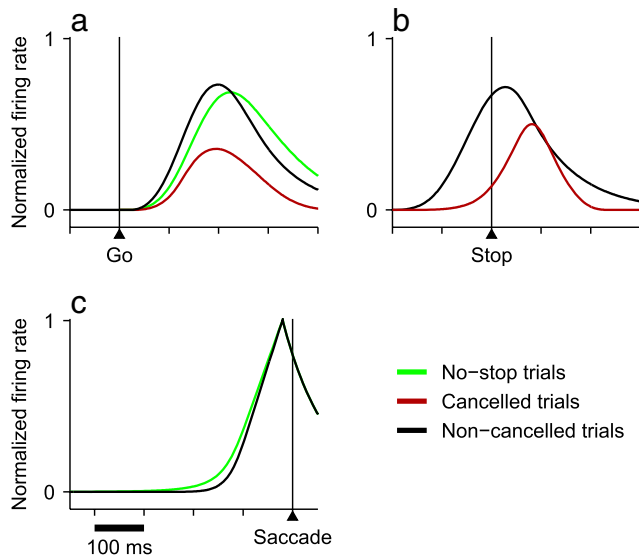


Figure 3. Neural correlates of the cancellable rise-to-threshold model. The simulated motor plans (M) were sorted according to trial type and outcome, aligned on a given event (go, stop signal, or threshold crossing), and averaged. **a**, Mean responses in no-stop (green), cancelled (red) and noncancelled trials (black) aligned on the go signal. **b**, Mean responses in cancelled (red) and noncancelled (black) trials aligned on the stop signal. **c**, Mean responses in no-stop (green) and noncancelled trials (black) aligned on threshold crossing. For stop trials, all SSDs are included in the averages.

nal). It is an ideal curve in this case because it cannot be calculated in the same way based on empirical data; for that, we have specifically developed an empirical tachometric curve (see below). However, it is instructive to consider the ideal curve as well.

The raw processing time (rPT) is the crucial quantity in the task. It is calculated as follows:

$$\text{rPT} = \text{RT} - \text{SSD} \quad (6)$$

where both the RT and the SSD are specific for each individual trial. The rPT is the theoretical limit on the amount of time available for viewing the stop signal in each trial (Fig. 1c). Negative values correspond to trials in which the saccade was initiated before the stop signal was presented, whereas positive values correspond to trials in which, in principle, the stop signal could have been seen. Not all the rPT is usable, though. Any transmission delays consume part of it and count as dead time or non-decision time. Therefore, the effective processing time (ePT) is as follows:

$$\text{ePT} = \text{rPT} - \Delta_S - \Delta_E \quad (7)$$

where Δ_S and Δ_E are the afferent and efferent delays of the stop signal, respectively. In the model, positive ePTs correspond to trials in which the motor plan was actually decelerated, whereas negative values correspond to trials in which there was no deceleration (Fig. 2). Because all results presented below are equivalent for rPT and ePT, all subsequent analyses are presented in terms of rPT only.

The above expression for rPT (Eq. 6) is well defined for noncancelled trials in which an RT is measured. In contrast, in cancelled trials, the amount of stimulus viewing time is still the quantity of interest, but no saccade and therefore no RT is produced. To overcome this problem (i.e., to calculate rPT in this case), we define a cancellation time (CT), which may be used in place of the RT in cancelled trials.

Normally, the RT is measured between the onset of the go signal and the onset of the saccade, which occurs Δ_E ms after threshold crossing. The idea is to take as the reference event, equivalent to threshold crossing, the point in time at which the neural activity reaches its maximum level, which we call T_{\max} (Fig. 4a, arrow). The CT is then as follows:

$$\text{CT} = T_{\max} + \Delta_E \quad (8)$$

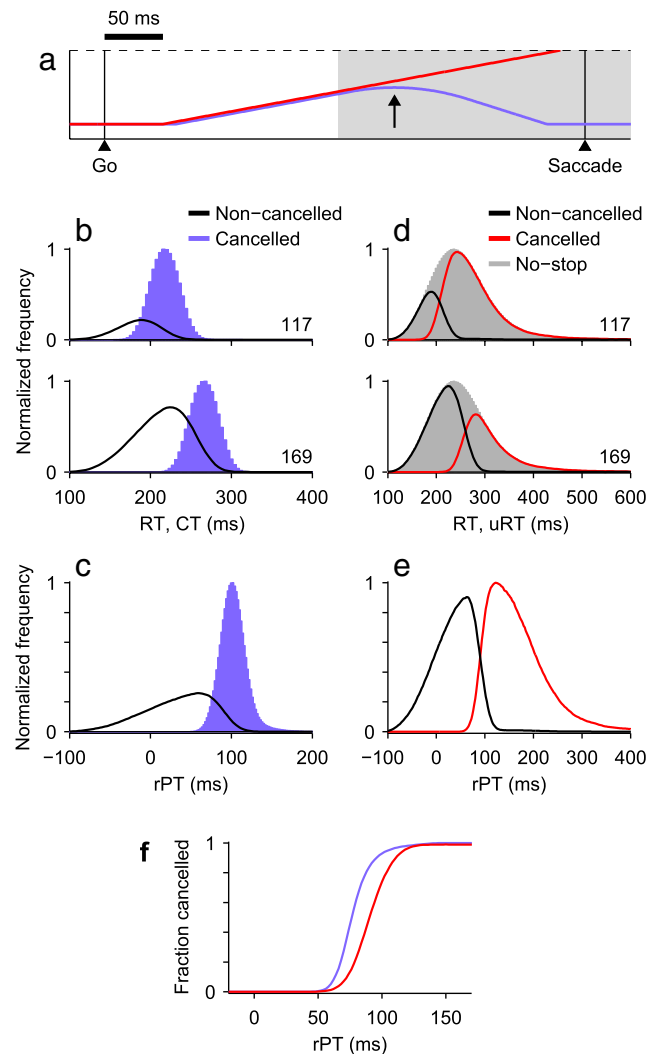


Figure 4. Calculating the ideal and empirical tachometric curves. **a**, Schematic of a successfully cancelled trial. The actual motor plan (purple trace) was decelerated and stopped; the time of maximum activity T_{\max} (arrow) defines the CT of the trial (Eq. 8). The red trace represents the point at which that same plan would have crossed threshold had it not been decelerated. The RT for the red trace is the uRT for the shown cancelled trial. **b**, Simulated distributions of RTs for noncancelled trials (black) and of CTs for cancelled trials (purple). Results are shown for SSDs of 117 and 169 ms, as indicated. **c**, Distributions of rPT values for noncancelled trials (black; based on RTs) and cancelled trials (purple; based on CTs). Results are pooled across all SSDs. **d**, Simulated RT distributions for no-stop (gray) and noncancelled (black) trials, both experimentally measurable; the difference is the inferred distribution of uRTs for cancelled trials (red). **e**, Distributions of rPT values for noncancelled trials (black; based on RTs) and cancelled trials (red; based on uRTs). Results are pooled across all SSDs. **f**, Ideal (purple) and empirical (red) tachometric curves obtained from the distributions in **c** and **e**, respectively.

Importantly, T_{\max} always marks the point in time at which the motor plan has advanced the most, whether threshold is crossed (in noncancelled trials) or not (in cancelled trials). This guarantees a continuous transition between RT and CT values for trajectories that are similar but produce different outcomes, which ultimately guarantees that the tachometric curve will be smooth. Another consequence of this is that in noncancelled trials the CT is precisely the RT. In sum, then, the rPT can be computed as follows:

$$\text{rPT} = \text{CT} - \text{SSD} \quad (9)$$

in cancelled trials (Fig. 4b,c). In this way, the rPT is always the maximum amount of time available for viewing and processing the stop signal relative to the point of maximum neural activity.

With this provision, the performance of the model can be easily plotted as a function of rPT (Fig. 4f, purple trace). For this, simulated trials are grouped according to rPT bins (Fig. 4c) and the fraction of cancelled trials is calculated for each bin. Tachometric curves were computed with bins separated by 1 ms and bin width of 20 ms.

The empirical tachometric curve. In contrast to the ideal curve, the empirical tachometric curve is derived from behavioral data alone; no simulated responses are involved. The goal is the same, though, to determine the probability of a successful cancellation as a function of processing time. In this case, the main idea is to assume that the RT associated with a cancelled saccade is simply the RT that would have been observed had the stop signal not been given (Fig. 4a). In this way, the definition of rPT stays consistent; it is always equal to the difference between the RT and the SSD for any given trial (Eq. 6). The problem here is determining the RTs that would have been observed in the trials that were, indeed, successfully cancelled. The key is that those unobserved reaction times (uRTs) can be extracted from the full distribution of RTs measured during no-stop trials. This is how.

The saccadic plans that develop quickly are precisely those that cannot be cancelled when the stop signal is presented (Fig. 3a), and hence the RTs in noncancelled trials (Fig. 4d, black curves) are simply a subset of the wider range of RTs collected in no-stop trials (Fig. 4d, gray areas). Technically, this is not quite exact because some of the RTs in noncancelled trials are slightly lengthened by the deceleration compared with those measured in no-stop trials, but in practice the difference is negligible and essentially undetectable. The RTs that remain in the full (no-stop) distribution when those fast (noncancelled) RTs are discounted or “subtracted” are precisely the uRTs (Fig. 4d, red curves).

In practice, the empirical tachometric curve is generated as follows. Suppose that a single SSD was used. The first step is to convert each measured RT into an rPT by subtracting the SSD (Eq. 6). This produces two sets of rPT values: one for no-stop trials and another for noncancelled trials. Next, create two corresponding rPT distributions (i.e., two histograms) using identical time bins for both. Say that, as a result, H_{NS}^j and H_{NC}^j give the numbers of no-stop and noncancelled rPTs in bin j , respectively. The next step is to subtract bin by bin to produce the distribution of rPTs in cancelled trials:

$$H_C^j = \alpha H_{NS}^j - H_{NC}^j \quad (10)$$

where the constant α is equal to the total number of stop trials at the given SSD divided by the total number of no-stop trials. Having done this, one now has an rPT distribution for noncancelled trials (H_{NC}^j ; Fig. 4e, black curve) and an rPT distribution for cancelled trials (H_C^j ; Fig. 4e, red curve) in the correct proportions. Using these, the empirical tachometric curve (Fig. 4f, red curve) is simply as follows:

$$f_C^j = \frac{H_C^j}{H_C^j + H_{NC}^j} \quad (11)$$

where f_C^j is the fraction of correct (cancelled) trials at bin j . The only additional complication is that, when there are multiple SSDs, the results need to be appropriately pooled before this last step. Because the CT is always shorter than the uRT, the empirical tachometric curve is shifted to the right of the ideal one (Fig. 4); their behaviors, however, are very similar.

A Matlab function that implements the above procedure is available from the authors on request.

Methods for computing the SSRT. Although the present model is conceptually unrelated to the SSRT, this quantity is nearly universally used to characterize performance in the countermanding task, so we compute it to relate our results to previous models and experiments.

Detailed procedures for computing the SSRT, and the rationale behind them, have been described before (Hanes et al., 1998; Kornyló et al., 2003; Stevenson et al., 2009; Verbruggen and Logan, 2009). Briefly, the main idea is to start with the cumulative distribution of RTs from no-stop trials, and note that in stop trials the faster RTs are precisely those that become incorrect, noncancelled saccades. The difference in time between the SSD and the longest of those fast RTs is then considered the minimum time required by the stop process (i.e., the SSRT).

In the integration method, the SSRT is calculated at each SSD by first obtaining the probability p of a noncancelled saccade for that SSD and looking up the RT at which the cumulative RT distribution reaches that same value p . The SSRT is then estimated by subtracting the SSD from the found RT value. The final result is the SSRT averaged over all SSDs for which p is between 0.1 and 0.9, to avoid the tails of the two functions.

In the mean and median methods, there is only one subtraction. In the median method, the SSRT is simply the difference between the mean RT from no-stop trials and the SSD at which $p = 0.5$, where p again is the probability of making a noncancelled saccade. The mean method is similar. It takes the mean RT from no-stop trials and subtracts a single SSD value, in this case the “mean” SSD. The mean SSD is calculated by treating the inhibition function (the probability of a noncancelled saccade as a function of SSD) as if it were a cumulative distribution of SSDs.

All three methods were implemented as described by Kornyló et al. (2003).

Comparisons with the go/stop race model. In our implementation of the race model, both a motor and a stopping plan rise linearly to threshold during stop trials. The two processes develop just like the M variable in the cancellable rise-to-threshold model; that is, once the go signal is given, the motor plan starts rising after an afferent delay (mean = 100, SD = 40 ms) and with a given buildup rate (mean = 8.3, SD = 2.6 U/ms), which vary across trials. Similarly, once the stop signal is given, the stopping plan starts rising after an afferent delay (mean = 29, SD = 0 ms) and with a given buildup rate (mean = 25.3, SD = 6.2 U/ms). Finally, again, there is a small probability that the stop signal is simply ignored ($p_e = 0.011$). In this model, the plan that reaches threshold first determines the outcome: the trial is cancelled if the stopping plan wins, and it is noncancelled if the motor plan wins. With the chosen parameter values, the resulting fits to the psychophysical data reported by Boucher et al. (2007) were just as good as those obtained with the cancellable rise-to-threshold model.

To distinguish the rise-to-threshold and race models, we simulated their responses in a different type of task trial in which the stop signal is presented only briefly, after which the go instruction again takes effect (see Results). For such transient-stop trials, in both models, we used the parameter values that fitted the Boucher et al. (2007) data, except with $p_e = 0$, we specified a stop-signal duration of 35 ms, and assumed that the motor plans resumed immediately after the end of the stop command. Including an additional delay after the stop command did not alter the expected differences between models.

Results

A simple phenomenological description of the countermanding task

There are two types of trials in the countermanding task: those in which the stop signal is presented (stop trials) and those in which it is not (no-stop trials). In no-stop trials (Fig. 1a), the subject simply has to make a saccade to a single target, and because errors are very rare in practice, we consider all such trials to be correct. In contrast, stop trials can be either correct (cancelled, Fig. 1b) or incorrect (noncancelled, Fig. 1c). The probability of success is controlled by the experimenter through the SSD, which is the time interval between the go and the stop signals; the longer the SSD in a given trial, the higher the likelihood that an incorrect saccade is produced. In no-stop and noncancelled trials, the behavioral output is the RT, which is measured between the go signal (target onset/fixation offset) and the onset of the eye movement.

We model the motor plan that develops during a no-stop trial as a gradual increase in the firing rate of a population of oculomotor neurons with motor fields at the location of the visual target (Fig. 2a). The rise in activity starts a short time after target presentation (afferent delay is Δ_T), which is the go signal, and continues until a fixed threshold level is reached. Once the threshold is crossed, the system commits to generating the saccade, which starts after a short efferent delay Δ_E . The model neu-

ronal activity builds up at a constant rate r_B^0 ; however, this initial buildup rate varies randomly from trial to trial, so the RTs produced by the model vary randomly across simulated trials.

This description based on a linear rise-to-threshold is simple, but it generates RT distributions that are close to those measured experimentally and is a reasonable approximation of the neuronal activity recorded in oculomotor areas before saccades are triggered (Carpenter and Williams, 1995; Hanes and Schall, 1996; Reddi and Carpenter, 2000; Roitman and Shadlen, 2002; Brown and Heathcote, 2007; Brown et al., 2008; Stanford et al., 2010).

In the model, stop trials start in the same way, with a gradual increase in activity toward threshold, but depending on the initial buildup rate and the time at which the stop signal is given, the developing motor plan may or may not change course and be cancelled (Fig. 2*b–d*). Intuitively, the outcome can be understood at two extremes. First, when the initial buildup rate is relatively high and/or the SSD is long, the activity level may reach threshold before the stop signal is even detected, in which case a noncancelled saccade with a short RT is produced (Fig. 2*b*). At the other end, when the initial buildup rate is relatively low and/or the SSD is short, the stop signal is detected early. This triggers a deceleration of the motor plan, which turns around and goes back to its baseline (Fig. 2*d*). In this case, fixation is successfully maintained and there is no RT; the saccade is cancelled. In between these two extremes, it is possible to have an intermediate situation in which the stop signal is detected and the motor plan slows down, but threshold is reached anyway (Fig. 2*c*). This happens because the turn-around in activity is not instantaneous, so the firing rate keeps increasing for some time as it decelerates. This results in a noncancelled trial with an intermediate RT.

In summary, there are three relevant trial/outcome combinations: no-stop trials (correct), cancelled trials (correct) and noncancelled trials (incorrect). Saccadic plans are modeled as a rise in activity toward a fixed threshold whose crossing triggers the eye movement; and after the detection of the stop signal, the rise in activity slows down and eventually reverses. We refer to this as the cancellable rise-to-threshold model.

Standard psychophysical measures of performance

Performance in the countermanding task is generally quantified in three complementary ways. First, through the psychometric curve, which gives the fraction of stop trials that are cancelled (i.e., the fraction correct) as a function of SSD (Fig. 5*a*). In general, the fraction cancelled decreases monotonically with longer SSDs. The complement of this curve (the fraction of noncancelled trials vs SSD) is often referred to as the “inhibition function,” but we avoid that terminology because, as will be shown below, such curve is unlikely to provide direct information about inhibitory dynamics.

The second way to quantify performance in the task is by plotting either the standard or the cumulative RT distributions for no-stop and noncancelled trials (Fig. 5*b*). No-stop RTs generally span a large range of values, whereas noncancelled RTs occupy a smaller range toward the faster end. In general, the shorter the SSD, the shorter the noncancelled RTs (Fig. 5*b*), but this is not always the case (see below).

Finally, the third quantity used to characterize performance in the task is the SSRT, which attempts to infer the point in time at which the motor plan is cancelled. However, no specific event in the countermanding task or in our model corresponds exactly to the SSRT; its definition depends crucially on the go/stop race model mentioned earlier in the Introduction. We defer its discussion to the sections that follow.

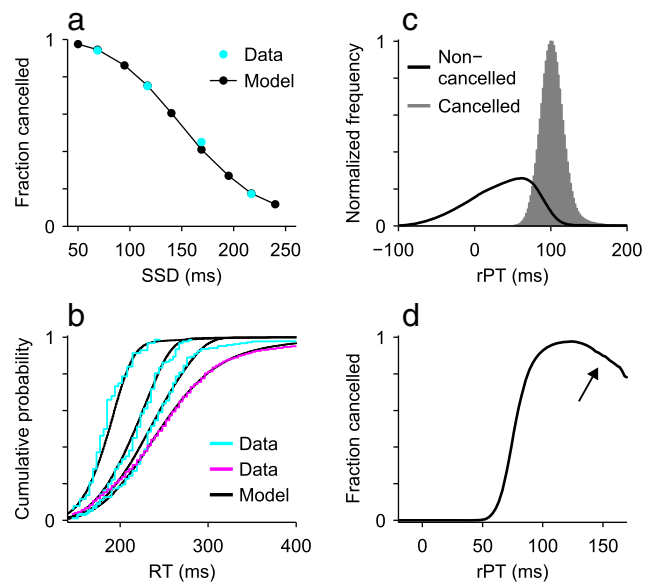


Figure 5. Quantifying performance in the countermanding task. *a*, Psychometric curve. Fraction of cancelled saccades, or successfully stopped trials, as a function of SSD. *b*, Cumulative distributions of RTs for no-stop trials (magenta trace) and noncancelled trials at three SSDs: 117, 169, and 217 ms (cyan traces, from left to right). *c*, Distributions of rPT values for cancelled (gray) and noncancelled (black) trials. Results are pooled across all SSDs. *d*, Ideal tachometric curve. Fraction of successfully cancelled stop-trials as a function of rPT. Arrow indicates curve tail, where errors occur because of lapses (parameter p_l). *a*, *b*, Blue dots indicate experimental data redrawn from Boucher et al. (2007). All other results are from model simulations with parameters optimized to fit those experimental data (see Materials and Methods).

The cancellable rise-to-threshold model replicates quite well both the psychometric curve (Fig. 5*a*) and the RT distributions (Fig. 5*b*) measured experimentally. This is true for the data shown in Figure 5, which were redrawn from the study by Boucher et al. (2007), and for numerous experiments discussed below. As will be seen, its strength is that it provides an intuitive understanding of the interplay between motor and perceptual mechanisms that determines a subject’s psychophysical performance under a wide variety of task conditions.

Fast perceptual performance revealed by the ideal tachometric curve

The present framework also allows the construction of a new, powerful metric for quantifying perceptual performance in the countermanding task. This new curve, the tachometric curve, tracks the fraction of correct choices as a function of processing time (see Materials and Methods), which is the variable that is fundamental for success in the task. We call it the tachometric curve because it reveals the effective perceptual processing speed of the subject (Stanford et al., 2010). There are two variants of this curve: ideal and empirical. We consider the ideal curve first.

As mentioned earlier, the difficulty of the task varies according to the SSD, which is a task parameter under the experimenter’s control. However, the SSD is the period of time during which the relevant sensory cue, the stop signal, is absent. What really matters is how long the cue is present and available to inform the subject’s actions, and that interval is the rPT (Fig. 1*c*). In each noncancelled trial, it is simply the difference between the RT and the SSD (Eq. 6).

The longer the rPT, the longer the subject has to view the stop signal and detect it. Consequently, one would expect a longer rPT to produce a higher probability of cancelling the eye movement.

So, the goal is to calculate the fraction of successful cancellations as a function of rPT, but there is a complication: cancelled trials do not have an overt behavioral manifestation; there is no RT to measure when a saccade is successfully stopped.

With simulated data, it is possible to work around this problem; that is, to calculate the maximum amount of cue viewing time (rPT) in a consistent way in all trials. For this, we use a quantity that is equivalent to the RT but is valid in cancelled trials, the CT (Eq. 8). Just as the RT is defined relative to the time at which threshold is crossed, the CT is defined relative to the time at which the neural activity reaches its maximum level in a trial (arrow in Fig. 4*a*). This point, T_{\max} , is a natural choice because, by definition, after T_{\max} the motor activity does not increase anymore, so it is at that point that the cancellation becomes absolutely certain. For consistency, we still define T_{\max} in noncancelled trials as the time at which threshold is crossed. This guarantees a smooth, continuous transition between RT and CT as a function of buildup rate or SSD or any other parameter that influences the trial outcome, and thus guarantees a smooth tachometric curve.

With this procedure in place, distributions of RTs for noncancelled trials (Fig. 4*b*, black curves) and distributions of CTs for cancelled trials (Fig. 4*b*, purple shades) can be obtained at each SSD. These are turned into rPT distributions simply by subtracting the corresponding SSDs (Eqs. 6 and 9). Then, by pooling the results across all SSDs, separate rPT histograms for all cancelled and all noncancelled trials are produced (Figs. 4*c* and 5*c*). Finally, based on those two distributions, it is straightforward to obtain the fraction of cancelled (i.e., correct) trials as a function of rPT (i.e., the tachometric curve; Figs. 4*f* and 5*d*).

The tachometric curve is key. First, it is much steeper (Fig. 5*d*) than the psychometric curve (Fig. 5*a*), indicating that performance in the countermanding task correlates much more tightly with rPT than with SSD. Second, the onset of the curve shows exactly when the perceptual process starts relative to the onset of the stop signal. For the parameters used in Figure 5, this occurs ~55 ms after the stop is presented (i.e., rPT ~55 ms). Third, the tachometric curve reveals the effective speed of the perceptual detection process: the fraction cancelled goes from nearly zero to 0.5 in ~20 ms. That is, 20 ms of viewing time is the difference between having virtually no chance of detecting the stop and cancelling the saccade, versus having a 50% probability of doing so. And fourth, the maximum value of the curve and the subsequent tail (Fig. 5*d*, arrow) are indicative of errors that occur at long viewing times, when the subject should have detected the stop signal but did not. These errors might, perhaps, reflect limits on cognitive, attentional, or sensory-integration processes.

The tachometric curve isolates the perceptual performance of the subject; its onset and slope reveal when the detection of the stop signal effectively begins and how fast it proceeds. It is important to stress that these properties are independent of motor execution. In particular, the curve is essentially the same regardless of the rate at which the oculomotor activity builds up (Stanford et al., 2010). The distribution of buildup rates and the afferent delay of the go signal determine the probability that each particular rPT is observed, but these quantities do not affect the relationship between rPT (stimulus viewing time) and likelihood of successful detection. Overall performance does change when the buildup rates vary, but this is because they determine which parts of the tachometric curve are sampled more heavily during an experiment; the tachometric curve itself remains almost the same (further illustration below).

We refer to this curve as “ideal” because it can only be generated with simulated data, for which the turnaround point, T_{\max} , is

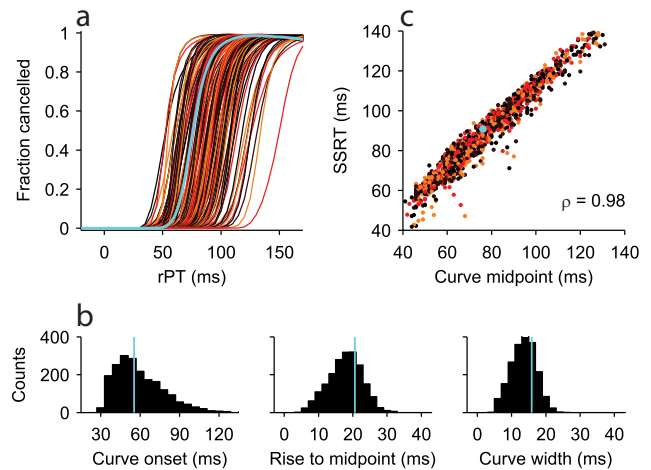


Figure 6. Sensitivity of the ideal tachometric curve to changes in model parameters. The model was run 2000 times, each one with parameters sampled from normal distributions centered on the values used for Figure 5. *a*, Resulting tachometric curves (thin lines; 125 of 2000 shown) and original curve (thick blue line). *b*, Frequency distributions for three quantities: curve onset (rPT at which the fraction cancelled exceeds 0.02), rise to midpoint (change in rPT for which the fraction cancelled goes from 0.02 to 0.5), and curve width (change in rPT for which the fraction cancelled goes from 0.25 to 0.75). Blue lines indicate values from the original curve (blue trace in *a*). *c*, SSR T computed with the integration method versus the midpoint of each tachometric curve (rPT at which the fraction cancelled equals 0.5). The blue dot corresponds to the original curve.

known. A close approximation can be constructed based on empirical data (Fig. 4), as elaborated further below, but for now it is useful to consider its robustness and implications for interpreting neural activity, as discussed next.

Robustness of the critical time window for stimulus detection

Our analysis shows that there is a critical time window during which a relatively small amount of viewing time (~20 ms) makes an enormous difference in the observed detection accuracy. But how reliable is this result? This question is important because, even though there is close agreement between the model and experimental data (Fig. 5*a,b*), the tachometric curve (Fig. 5*d*) was constructed from the model responses.

To investigate this, we set the model parameters used in Figure 5 as a reference and then performed 2000 additional runs of the model, each time with different parameter values drawn randomly from a distribution. Every extra run consisted of 100,000 simulated trials from which all the relevant psychophysical measures were computed. Each parameter was sampled over a large range, and each run had a different combination of parameter values. To get an idea of the magnitude of these variations, for the 2000 additional runs the SD of each parameter distribution was equal to 20% of the reference, whereas in the optimization procedure used for fitting the model to the experimental data, an ~10% change in a single parameter typically made a noticeable difference in the cost function (Eq. 5; see Materials and Methods). We then compared the reference tachometric curve (Fig. 6*a*, thick blue line) with the 2000 curves obtained from the additional runs (Fig. 6*a*, thin lines).

In general, the tachometric curves generated in this way varied quite widely in onset but much more moderately in rise time. The onset ranged approximately between 30 and 120 ms, whereas the rise time depended slightly on the method used to measure it but ranged approximately between 5 and 30 ms (Fig. 6*b*). These results suggest that, although the onset of the detection process

could be somewhat variable, our original estimate of ~ 20 ms required for accurate detection of the stop signal (i.e., to reach a fraction cancelled of 0.5) is no more than ~ 10 ms below the actual value. The critical time window during which the perceptual detection process effectively develops seems to be quite narrow indeed.

Relationship between the SSRT and the ideal tachometric curve

In each cancelled trial of the countermanding task, the SSRT, as conceived originally, is interpreted as the amount of time required by the inhibitory process to act; it is meant to be related to a precise point in time at which the saccadic plan is stopped. Because of this, many experiments are designed with the intention of altering the presumed inhibitory mechanisms needed to cancel the planned saccade, and attempt to resolve the resulting changes by measuring the average SSRT (see below for examples).

Further below we argue that the SSRT has inherent problems that limit its interpretability, and we propose what we consider a better alternative, the empirical tachometric curve. First, however, we want to know if there is some feature of the responses generated by the cancellable rise-to-threshold model that correlates most tightly with SSRT. In particular, what is the relationship between the SSRT, as it is calculated in practice, and the ideal tachometric curve, which we think is the key indicator of perceptual performance?

To address this, we used the same 2000 simulations discussed in the previous section, which sample a wide range of potential psychophysical behaviors in the task. For each run, we computed the mean SSRT and compared it with several other quantities, including the onset, midpoint, and width of the ideal tachometric curve, and the mean rPT in cancelled trials.

An important issue here is that there are several ways to compute the SSRT for any given empirical dataset (e.g., Hanes et al., 1998; Kornylow et al., 2003; Stevenson et al., 2009; Verbruggen and Logan, 2009). These methods do not fully agree with each other; they give different numbers regardless of the quality and quantity of the data. So, we tested the three most common alternatives: the integration method, the mean method, and the median method (see Materials and Methods). All the correlation coefficients mentioned next were significantly different from zero (permutation test, $p < 0.0001$).

The strongest association found was between the SSRT calculated via the integration method and the midpoint of the ideal tachometric curve (i.e., the rPT for which the fraction cancelled is 0.5; Fig. 6c); the correlation coefficient between them was 0.98. The second and third highest coefficients were between the same SSRT and mean rPT (0.97), and between the same SSRT and curve onset (0.93). Comparisons involving the SSRTs obtained with the mean and the median methods produced somewhat lower numbers when paired with curve midpoint (0.81 and 0.91, respectively) or curve onset (0.74 and 0.84, respectively). All these numbers depend to some extent on the parameter values chosen as a reference in the simulations, so the correlations may vary slightly. The main conclusion, however, is that changes in the SSRT, particularly when calculated with the integration method, correspond closely to shifts of the tachometric curve.

This result has a major consequence: if the tachometric curve indeed reflects the perceptual performance of the subject in the task, then it follows that the SSRT must also correlate with perceptual performance, a possibility that, to our knowledge, has not been considered in previous studies. This means that when the detection of the stop signal changes in some way that shifts the

tachometric curve, the SSRT will change too. Given that the statement “Shorter SSRT represents better inhibition abilities” (Jacobson et al., 2011) is quite representative of the dominant view, this observation has important implications for interpreting numerous experimental results: changes in the SSRT are not necessarily indicative of changes in inhibition.

The following sections discuss various experimental manipulations in the countermanding task that can be easily understood with the cancellable rise-to-threshold model. There it is further shown that variations in SSRT can indeed be attributed to variations in the perceptual detection process, a process that must take place before any inhibition comes into play.

Accounting for the “gap effect” in the countermanding task

One way to alter saccadic RTs reliably is through the “gap effect,” a manipulation whereby the fixation point is extinguished before a saccadic target is shown (Saslow, 1967; Munoz et al., 2000). In the countermanding task in particular, extinguishing the fixation point 200 ms before the onset of the saccadic target produces a decrease in the mean RT of no-stop trials of ~ 50 ms compared with the standard version of the task (Stevenson et al., 2009), in which the fixation point disappears at the same time as the target appears (i.e., no gap; Fig. 1). Furthermore, the 200 ms gap also produces a decrease of ~ 40 ms in the SSRT (Stevenson et al., 2009). These results would seem to imply that the gap effect in the countermanding task speeds up two dynamic processes, the oculomotor preparation and its active cancellation.

However, these two changes in behavior can be easily reproduced by the model (Fig. 7) simply by assuming that the early extinction of the fixation point produces both a moderate increase in the mean buildup rates and, more importantly, a decrease in the afferent delays. That is, when the fixation spot disappears early, the screen remains blank for 200 ms and the appearance of the saccadic target becomes the go signal; then, once the target is displayed, the oculomotor plan starts to develop slightly sooner (shorter latency) and slightly faster (higher buildup rate) than in the standard condition without a gap; finally, when the stop signal is shown, it is also detected slightly sooner (shorter latency). This combination of effects leads to shorter RTs in no-stop trials (Fig. 7b,c) and fewer cancellations (Fig. 7a) and shorter SSRTs (Fig. 7d) in stop trials, as reported experimentally. In this way, Figure 7a–d closely replicates the results of Stevenson et al. (2009; their Fig. 3) without assuming any changes in inhibitory action.

In this case, the mean afferent delay must change both for the detection of the target (through parameter μ_T) and of the stop signal (through parameter μ_S). The former is necessary to account for the drop in no-stop RTs, and the latter is necessary to account for the drop in SSRT. The reason is this: the target latency cannot affect the SSRT because the SSRT depends on when the motor plan turns around relative to its onset—it does not matter when the onset occurs—and the stop-signal latency cannot affect the RT in no-stop trials because in those the stop signal is just not given. When both latencies decrease equally, the preparatory activity starts sooner and turns around sooner, so all the RT distributions, for stop and no-stop trials (Fig. 7b), shift. The shift of the stop trials, in particular, causes a corresponding shift in the rPT distributions (Fig. 7e). This, in turn, shifts the midpoint of the tachometric curve, which correlates tightly with SSRT, as discussed earlier (Fig. 6c).

So, the model predicts that in the gap condition the latencies to both visual stimuli, the target and the stop, must be shortened approximately equally relative to the standard no-gap condition.

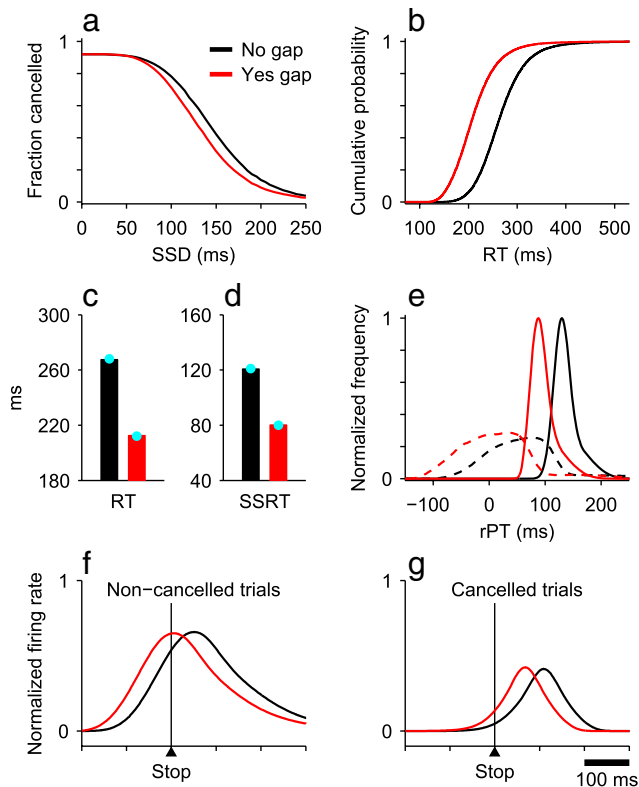


Figure 7. The cancellable rise-to-threshold model reproduces the gap effect in the countermanding task. Results are shown for two runs with two sets of parameters corresponding to no-gap (black) and yes-gap (red) conditions (Table 1). **a**, Psychometric curves. **b**, Cumulative distributions of RTs in no-stop trials. **c**, Mean RTs in no-stop trials. **d**, SSRTs obtained with the integration method. **e**, Distributions of rPT values in cancelled (solid lines; based on CTs) and noncancelled (dashed lines; based on RTs) trials. **f**, Expected oculomotor responses for noncancelled trials aligned on the stop signal. **g**, Same as in **f**, but for cancelled trials. **c**, **d**, Blue dots indicate mean experimental values reported by Stevenson et al. (2009).

This refers to the latencies of the motor-related responses, which is what the model represents (not to the latencies of purely visual neurons). The neurophysiological prediction is that all motor plans aligned on the go and on the stop signal should shift to the left in the gap condition relative to the no-gap (Fig. 7*f,g*). In general, the gap indeed decreases the latencies of saccade-related responses (Dorris et al., 1997; Everling and Munoz, 2000; Sparks et al., 2000), but the untested prediction for the countermanding task is specifically that the average traces should shift for all outcomes, no-stop, cancelled, and noncancelled trials.

The inverse of the gap effect

Nouraei et al. (2003) conceived an experiment that produced exactly the opposite effects obtained in the gap paradigm. Human subjects performed the countermanding task after receiving varying doses of an anesthetic (sevoflurane) in the sedative range, and both the RTs in no-stop trials and SSRTs in stop trials increased monotonically with increasing levels of anesthetic. In five of six subjects, the magnitude of the changes in RT and SSRT was indistinguishable at all doses. Therefore, as with the gap effect, parallel changes in RT and SSRT were produced, except in the reverse direction: all responses slowed down. This result can be easily reproduced by the model very much as in Figure 7 simply by assuming that the anesthetic leads to a systematic increase in the afferent delays and perhaps a moderate decrease in the mean buildup rate of the premotor activity, predictions that, like those

mentioned above, are verifiable through recordings in trained monkeys.

The changes in SSRT in the gap version of the countermanding task and in the experiment with sedated subjects were originally interpreted as affecting a complex series of neuronal processes associated with stopping a planned action. Although its biophysical basis may indeed be complex, our results demonstrate that the stopping process itself, understood as inhibitory activation, does not need to change at all; excellent quantitative agreement with the experiments can be obtained without invoking any changes in the way the presaccadic activity is decelerated by the stop signal. The data can be explained simply as a consequence of changes in the perceptual and motor-planning processes that are certain to be engaged during task performance.

Effects of redundant signals

A smaller but qualitatively similar combination of effects was reported using redundant go or stop signals (Cavina-Pratesi et al., 2001). In a way, this study performed a double dissociation of the two latency effects just discussed.

In this case, the motor response was to press a spacebar on a keyboard, and there were two complementary experiments (Cavina-Pratesi et al., 2001). First, the behavior of subjects was compared using one versus two visual go signals, with the stop always given by an auditory cue. In this condition, both visual stimuli indicated a movement to the same target. Slightly faster RTs (~10 ms decrease) in no-stop trials and slightly fewer cancelled trials were seen when two visual stimuli gave the go instruction instead of one. As in the previous section, this can be explained in our framework by a small decrease in the afferent delay of the target stimulus, perhaps with an additional small increase in the mean buildup rate of the saccadic plans. In contrast, in the second experiment, the go signal was always a single auditory stimulus but the stop signal could be cued either by one or by two visual stimuli. Now the result was a slightly shorter SSRT (~10 ms decrease) and a few more cancellations with two simultaneous stop signals than with just one; no change in no-stop RTs was seen, as expected. We propose that this decrease in SSRT is likely the result of a corresponding decrease in the afferent delay of the stop signal.

Classic experiments in visual detection show that RTs typically decrease with the intensity of the stimulus and that they are slightly shorter when two visual stimuli are displayed instead of one (Welford, 1980; Luce, 1986; Mordkoff et al., 1990; Marino and Munoz, 2009). The two experimental results of Cavina-Pratesi et al. (2001) are consistent with these observations and can be understood as a consequence of the exact same change in stimulation conditions; the only difference is that in one case the change applies to the detection of the go signal, whereas in the other it applies to the detection of the stop.

More generally, the results in this and the two preceding sections depend critically on large, systematic variations in afferent delays, which in the most literal interpretation would correspond to axonal delays, or synaptic/dendritic integration delays. Thus, it is worth pondering for a moment how modulations on the order of several tens of milliseconds could be accounted for in biological terms.

There are two important aspects to this question. First, the experimental evidence. It is well established that the response latencies of visual neurons vary widely and depend on stimulus intensity (Maunsell and Gibson, 1992), but the dependence is much stronger centrally than in the periphery; robust variations of hundreds of milliseconds have been documented in inferior

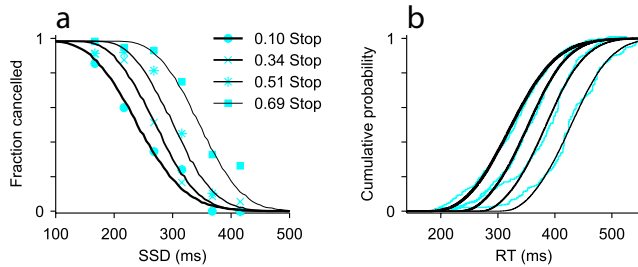


Figure 8. The dependence of performance on the overall fraction of stop trials is reproduced by variations in the buildup rate of the preparatory saccadic activity. Four experiments with different overall fractions of stop trials were simulated. Only two parameters changed; the mean and the SD of the initial buildup rate distribution (Table 1). **a**, Psychometric curves. **b**, Cumulative RT distributions for no-stop trials. Blue dots and lines represent experimental data redrawn from Emeric et al. (2007). Black lines are from model simulations.

temporal cortex (van Rossum et al., 2008; Oram, 2010). Latencies also decrease with multimodal compared with unimodal stimuli (Rowland et al., 2007; Fuhrmann et al., 2008). Psychophysically, differences in stimulus intensity alone may give rise to changes of many tens of milliseconds in the perceived temporal order of two sensory events, consistent with corresponding changes in sensory transmission delays (e.g., Gibbon and Rutschmann, 1969). Further, the response latencies of oculomotor neurons may depend on the outcomes of preceding trials (Fecteau and Munoz, 2003; Pouget et al., 2011; see below), so such effects are not exclusively stimulus-driven.

Second, it is important to recognize that neuronal latencies reflect the speed at which a signal propagates through multiple network layers, and theoretical work has shown that propagation speed is determined not only by axonal conduction times but also by the level of background activity, synaptic strength, balance between excitation and inhibition, and other biophysical factors (Idiart and Abbott, 1993; Van Vreeswijk and Sompolinsky, 1996; van

Rossum et al., 2008; Leblois et al., 2009; York and van Rossum, 2009). Any one of such factors could conceivably be altered internally or in a stimulus-dependent way to produce strong variations in latency. So, overall, latency modulations are not implausible, and may represent a prevalent mechanism controlling the timing of neural events.

Accounting for large, simultaneous shifts in psychometric and chronometric curves

In the countermanding task, performance depends strongly on the proportion of trials with a stop signal within a session of the task. This proportion has an enormous impact both on the probability of cancelling a trial and on the RTs of nonstop trials (Ramautar et al., 2004; Emeric et al., 2007). In one experiment, when the percentage of stop trials increased from 10% to 69% of all trials (stop plus no-stop), both the psychometric and chronometric curves of a subject shifted by ~ 100 ms (Emeric et al., 2007) (Fig. 8). So, when stop trials are more common, the subject systematically slows down all responses and increases substantially the probability of a successful cancellation. Notably, the SSRT does not seem to change at all under such conditions (Ramautar et al., 2004). Why not?

Seen as a perceptual decision-making problem, the answer is simple: because the perceptual demands of the task remain constant, and that is what the SSRT tracks. Indeed, the experimental data of Emeric et al. (2007, their Fig. 6) can be replicated accurately with the cancellable rise-to-threshold model by varying just

the distribution of initial buildup rates and keeping all other parameters constant (Fig. 8). If a higher proportion of stop trials leads, on average, to lower buildup rates, the result follows in a straightforward way: when stop trials are more frequent, RTs increase and a concomitant increase in the probability of a successful cancellation is produced. In this case, the afferent delay of the go/target stimulus may or may not vary too as a function of the probability of a stop; this cannot be determined with certainty based on the behavioral data alone. Either way, the tachometric curve is highly insensitive to both the latency of the go signal and the distribution of buildup rates, so in this example its shape does not change and its onset does so by a very small amount in comparison (~ 8 ms for the results in Fig. 8; see below). The model predicts that the buildup rates of the oculomotor plans, and perhaps their onset too, should vary as functions of the stop-signal probability, and the dependence should be the same for all trial outcomes.

These results are strongly reminiscent of motor-bias experiments, in which the subject prefers one response more than another even though the perceptual demands of the task remain fixed (Takikawa et al., 2002; Rorie et al., 2010; Stanford et al., 2010). Here, as in those studies, the behavioral effects on RT and success rate can be understood simply as a consequence of adjustments in the preparatory motor activity, which determines the distribution of RTs that is effectively available; the subsequent processes (stimulus detection or discrimination) do not need to change at all.

History-dependent effects

Another aspect of the countermanding task that has been noted in the literature is that subjects, both humans and monkeys, often show a statistically significant sequential adaptation: if the RTs of no-stop trials are sorted according to the preceding trial type, small but consistent differences are seen between the resulting groups. The major effect is that in no-stop trials preceded by another no-stop trial, RTs are typically lower than the overall average, whereas in no-stop trials preceded by a stop trial, RTs are typically higher than average (Cabel et al., 2000; Kornyló et al., 2003; Emeric et al., 2007). In addition, there may be other, more subtle, regularities. For instance, the effect may be more pronounced for no-stop trials preceded by two or more stop trials rather than by just one (Emeric et al., 2007), or the magnitude of the increase may be different depending on whether the preceding stop-trial was correct (cancelled) or incorrect (noncancelled) (Cabel et al., 2000; Nelson et al., 2010).

The results in the previous section (Fig. 8) suggest that the distribution of initial buildup rates can change quite dramatically over blocks of multiple trials. In contrast, these history-dependent variations in performance may reflect much smaller but still reliable adjustments in the same initial motor plans taking place on a much faster time scale of a single trial. Such sequential adaptation can be readily reproduced by the model by varying either the mean buildup rate or the afferent delay of the go signal as a function of trial history. For instance, suppose that after a no-stop trial the delay decreases toward a low value and that after a stop trial it increases toward a higher value. These low and high values may be adjusted to replicate a given experiment or subject.

Indeed, a neural correlate of precisely this sequential effect was recently reported in recordings from the FEF and superior colliculus (Pouget et al., 2011). Differences of 17 and 5 ms in the RTs of two monkeys were fully accounted for by differences of 18 and 5 ms in neuronal latency, respectively. That is, the onset of the buildup in oculomotor activity before a saccade depended sys-

tematically on the outcome of the previous trial, and notably, other response features, saccade threshold, mean buildup rate, and baseline activity, did not change significantly.

So, sequential effects can be explained by variations in motor preparation; no changes in inhibition or in the detection of the stop signal need to be invoked. This suggests that the oculomotor plans themselves must be quite flexible and that the outcome of each trial may have a slight but reproducible effect on the preparation for the next one, in agreement with observations in other tasks (Barracough et al., 2004; Sugrue et al., 2004; Busse et al., 2011).

Demystifying a paradoxical result about stop-signal modality

Various studies have shown that the modality of the stop signal is another important factor determining performance in the countermanding task (e.g., Akerfelt et al., 2006; Stevenson et al., 2009). Manipulations of the stop signal itself are interesting, and potentially quite informative, because they probe events that happen after the motor plan starts developing. Therefore, in contrast to the experimental data in the two previous sections, any observed effects must be independent of the afferent delay of the go/target stimulus and of the distribution of initial buildup rates.

In an experiment by Cabel et al. (2000), two types of stop signals were interleaved and presented with equal probability: (1) a red visual stimulus displayed at the location of the previous fixation point, the standard condition in most studies; and (2) a broadband noise burst emitted from a speaker 2 m above the subject's head. In both cases, the stop signal was meant to convey the same standard instruction, to inhibit a saccade to a previously presented target. When trials were sorted according to the two conditions, two major differences were observed. First, many fewer saccades were cancelled with the auditory stop signal than with the visual one (Fig. 9*a*, blue points). Most notably, the fraction of noncancelled trials was close to 0.5 for auditory stop-signal delays between 0 and 100 ms. And second, SSRTs were ~90 ms longer with auditory than with visual stimuli (Fig. 9*c*, blue points). Both results are puzzling because, in separate runs, the subjects' RTs to the same stimuli were measured in a simple button-press task, and the mean auditory RT was found to be 51 ms shorter than the mean visual RT. Therefore, although the subjects should have detected the auditory stop signal sooner than the visual, the stopping process seemed to occur later, according to the SSRT, and the probability of a successful cancellation was paradoxically low.

We tried to replicate these data with the cancellable rise-to-threshold model and concluded that the most natural parameter manipulations, including changing the magnitude and time constant of the deceleration, failed miserably. We found, instead, that two mechanistically plausible adjustments were sufficient to fully account for the experimental results. First, when the stop signal is auditory, the probability that it initiates the cancellation process at all must be <1. That is, in contrast to the visual signal, which always engages the stopping process, whether it is ultimately suc-

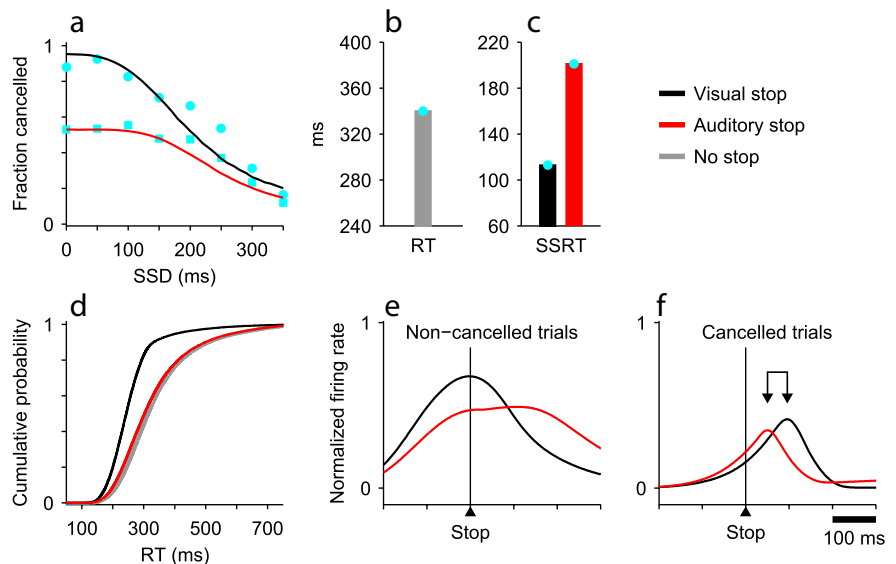


Figure 9. The effect of stop-signal modality is reproduced by the cancellable rise-to-threshold model. The model was run with two sets of parameters replicating the behavior observed with either visual (black lines) or auditory (red lines) stop signals (Table 1). *a*, Psychometric curves. *b*, Mean RT in no-stop trials. *c*, SSRTs for visual and auditory stop trials. *d*, Cumulative distributions of RTs for no stop trials (gray) and stop trials with visual (black) or auditory (red) stop signals. For this plot, stop trials with SSD ≤ 200 ms were considered. *e*, Predicted oculomotor activity in noncancelled trials. Traces are average model responses aligned on the stop signal. *f*, Same as in *e*, but for cancelled trials. Arrows indicate a difference of 46 ms between the peaks of the curves. *a*–*c*, Blue points indicate experimental values averaged across subjects, as reported by Cabel et al. (2000).

cessful or not, in a fraction p_e of the auditory trials the stop signal must be simply ignored. This type of lapse may be the result of attentional load, perhaps because in this case the subject needs to attend simultaneously to two stimuli of different modalities, the target on the screen and the sound from the speaker, but the actual origin makes no difference to the model. The second requirement of the model is that the afferent delay of the auditory stop signal must be lower than that of the visual signal by ~50 ms.

With these two conditions, the model replicates the experimental results of Cabel et al. (2000); extremely accurately (Fig. 9*a*–*c*). It accounts for both the higher rate of noncancelled trials (Fig. 9*a*) and the much longer SSRT (Fig. 9*c*) obtained in auditory trials, and at the same time it is fully consistent with the shorter latency of auditory versus visual responses found in the button-press task. This contrasts sharply with the conclusions drawn from the original study, in which the effects were attributed to the existence of multiple modality-specific, and possibly location-specific, stop processes. Our interpretation of the task as a motor plan that is modulated by a perceptual judgment is more parsimonious: it resolves the apparent paradox in the data without requiring multiple cancellation circuits or different deceleration dynamics across modalities. Rather, it is the perceptual detection step that is different in the visual and auditory conditions; it varies in effectiveness and in onset.

These conclusions can be tested because the model makes various specific predictions. First, behaviorally, the RTs for non-cancelled auditory trials at relatively short SSDs (≤ 200 ms) should be very similar to those in no-stop trials (Fig. 9*d*). This is because, in contrast to the visual case, in which noncancelled saccades are significantly faster than average (Fig. 5*b*), noncancelled saccades in auditory trials in that SSD range are not accelerated or decelerated at all; according to the model, the stop signal is simply missed, so it should have no effect.

Second, neurophysiologically, the oculomotor activity associated with cancelled and noncancelled saccades should reflect the

underlying assumptions of the model. When aligned on the stop signal, the mean activity produced during noncancelled saccades should be relatively localized in time (i.e., peaked) in the visual condition, as reported previously (Hanes et al., 1998; Boucher et al., 2007), because most noncancelled trials in this case have high buildup rates (Fig. 9e, black trace). In contrast, in the auditory condition, the average activity profile should be more spread out over time because it includes many trials, all those with auditory signals that were not processed, sampled from the full RT distribution (Fig. 9e, red trace). The third and crucial test, however, lies in the responses generated during cancelled trials: the activity should peak earlier in the auditory than in the visual condition, to reflect the shorter latency of the auditory stop signal (Fig. 9f). For these data, this is the key prediction of the model.

The empirical tachometric curve

We suggest that it is better to compute the tachometric curve than the SSRT; the curve is more informative and is based on a principled approach that can be easily related to the essential neurophysiological features of preparatory motor activity. Unfortunately, in the countermanding task, processing times (just like SSRTs) cannot be measured for individual cancelled trials. In practice, there are two ways around this limitation. One is to fit the cancellable rise-to-threshold model to the data and then generate the ideal tachometric curve based on simulated trials, as was done in Figure 5d. This is a viable approach if there are enough data to adequately constrain the model parameters. There is, however, an alternative that is better, in the sense that it requires no model fitting and can be implemented using just the psychophysical data. It produces a tachometric curve that is somewhat broader than the ideal one but otherwise has essentially the same properties and even some advantages. The procedure goes as follows.

The RTs in no-stop trials (Fig. 4d, gray areas) are, to an excellent approximation, a subset of the RTs in noncancelled trials (Fig. 4d, black traces) because the saccadic plans that develop quickly are precisely those that cannot be cancelled when the stop signal is presented (Fig. 3a). Conversely, the RTs that are left over in the no-stop distribution (Fig. 4d, red traces) correspond to the saccadic plans that started developing slowly enough, such that the system was able to cancel them when the stop signal was shown. The key insight here is that these can be considered as the unobserved RTs, or uRTs, of the successfully cancelled trials. As schematized in Figure 4a, each cancelled trial can be associated with a uRT, which is the RT that would have been observed had the stop signal not been presented. Although for any given cancelled trial it is impossible to say what that uRT was, because there is no overt action to measure, the distribution of uRTs for any SSD can be estimated by “subtracting” the noncancelled RTs from the full RT distribution obtained in no-stop trials (see Materials and Methods). The uRT can be thought of as determining the maximum processing time potentially available in a trial.

The empirical tachometric curve obtained in this way (Fig. 4f, red trace; see Materials and Methods) starts later and is somewhat broader than the ideal curve (Fig. 4f, purple trace), so it underestimates the speed of the detection/deceleration process. However, it has two advantages: (1) it has no tail, so once it reaches its maximum value it tends to stay flat; and (2) it is independent of the set of SSD values used to generate it (the ideal curve shows a slight dependence at the tail). Indeed, the empirical curve and the RT distribution in no-stop trials together provide a full characterization of a subject’s behavioral performance in the task (next section). In any case, the two tachometric curves behave very

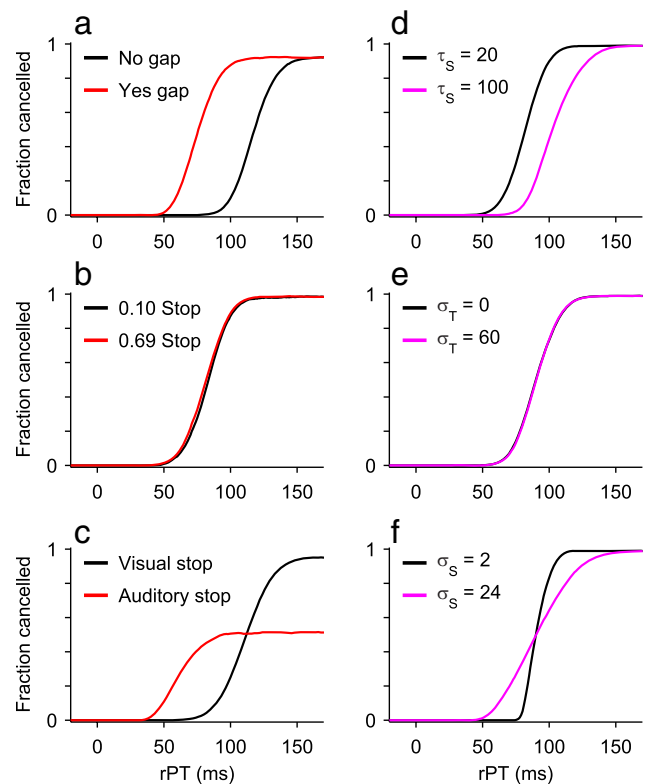


Figure 10. Empirical tachometric curves corresponding to various task manipulations. *a*, No-gap versus yes-gap conditions. Parameters as in Figure 7. *b*, Low versus high fraction of stop trials. Parameters as in Figure 8. *c*, Visual versus auditory stop signals. Parameters as in Figure 9. *d*, Slow versus fast deceleration. *e*, Low versus high variability in the afferent delay of the go signal. *f*, Low versus high variability in the afferent delay of the stop signal. *d–f*, Parameters are as in Figure 5, except as indicated in each panel.

similarly, and the empirical curve is just as sensitive to task manipulations (Fig. 10).

When there is a difference across conditions in the mean afferent delay with which the stop signal is detected, as in the case of the gap effect studied by Stevenson et al. (2009) (Fig. 7), the curve shifts horizontally by an amount that reflects the latency change quite closely (Fig. 10a). In contrast, when the only changes in the circuit across experimental conditions occur in the distribution of initial buildup rates, as was the case in Figure 8, which replicated data reported by Emeric et al. (2007), the tachometric curve stays essentially constant (Fig. 10b). Variations in the initial motor preparation have little effect on it. Also, the asymptotic level of the curve (i.e., its maximum value) is indicative of the proportion of errors, captured by parameter p_e , that cannot be attributed to limited viewing times (Fig. 10c), and perhaps involve lapses in attention, noise-limited integration, or other cognitive mechanisms (Shankar et al., 2011). This corresponds to the model recreation of the experiment by Cabel et al. (2000) (Fig. 9), in which stop signals of different modalities were used. The curve in this case also shifts to the left, again reflecting the decrease in afferent delay specific to the stop signal.

To explore other possible outcomes of behavioral manipulations, empirical tachometric curves were computed for simulations in which other model parameters were varied individually. The effect, if there was one, was almost always a shift of the curve, with little or no change in slope. This was true even for the time constant of deceleration τ (Fig. 10d) and is in agreement with the simulation results discussed earlier in which thousands of param-

Table 2. Error in the midpoint of the empirical tachometric curve for various numbers of trials

Midpoint error	Total number of trials				
	250	500	1000	2000	4000
Total SD (ms)	10	7	5	4	3
Sampling SD (ms)	4	3	2	1	1

The sampling error is due to numbers of trials only; the total error includes, in addition, the contribution from computations used to infer the rRTs of cancelled trials.

eter combinations were generated randomly (Fig. 6a). The only parameter that had an appreciable effect on the slope of the curve was σ_s , which determines the variability of the afferent delay of the stop signal (Fig. 10f). In contrast, neither the mean nor the variability of the afferent delay of the go signal had any impact on the curve (Fig. 10e), as mentioned earlier. All these effects were qualitatively the same with the ideal and empirical tachometric curves.

Finally, calculating the empirical curve from experimental data requires statistical manipulations (see Materials and Methods) that may introduce additional variability. We evaluated this variability by running simulations with different numbers of trials and found that such computations inject a relatively small amount of uncertainty into the tachometric curve. Across multiple runs with a fixed set of parameters (those used in Fig. 5), we compared the rRTs at which the curves reached 50% correct and found that the error in this midpoint value was quite reasonable even for a few hundred trials per run (Table 2). For example, with 1000 total trials, the midpoint of the tachometric curve can be determined within ± 5 ms (SD) of the actual value; of these 5 ms of total error margin, 2 ms is simply the result of the limited numbers of trials available, whereas 3 ms can be attributed to the computations required to process the experimental data. These results suggest that, in practice, shifts of the empirical tachometric curve can be resolved with relatively high precision.

Two curves characterize psychophysical performance

In this section, we demonstrate that two curves, the distribution of RTs in no-stop trials and the empirical tachometric curve, suffice to fully characterize a subject's performance in the countermanding task. This means that knowledge of these two fundamental curves is enough to derive any other psychophysical measures.

To see why this is indeed the case, suppose that the two curves are known and note that a subject's performance can be recreated (in a statistical sense) based those two curves alone. This is as follows (Fig. 11). Consider multiple stop trials of the task. In each trial, an RT is drawn randomly using the subject's cumulative RT distribution from no-stop trials and a uniform random number generator (Fig. 11a). Next, the corresponding rRT for that trial is obtained by subtracting the SSD (Fig. 11, blue arrow). This is regardless of whether the RT ends up being observed or not because that is how the rRT is defined for the empirical tachometric curve. Next, the probability of a correct cancellation is read off from the tachometric curve at the given rRT (Fig. 11b). Then, as a final step, a coin is flipped, biased according to that probability, to determine whether the saccade is actually cancelled or not. If not, then the RT is indeed observed and goes into the distribution for noncancelled trials for the given SSD; if yes, then the RT goes unobserved. In this way, RTs and outcomes are produced exactly as in a real experiment (or very nearly so; the match to the observed RT distribution for noncancelled trials is not perfect because of partial deceleration in some of those trials, but the

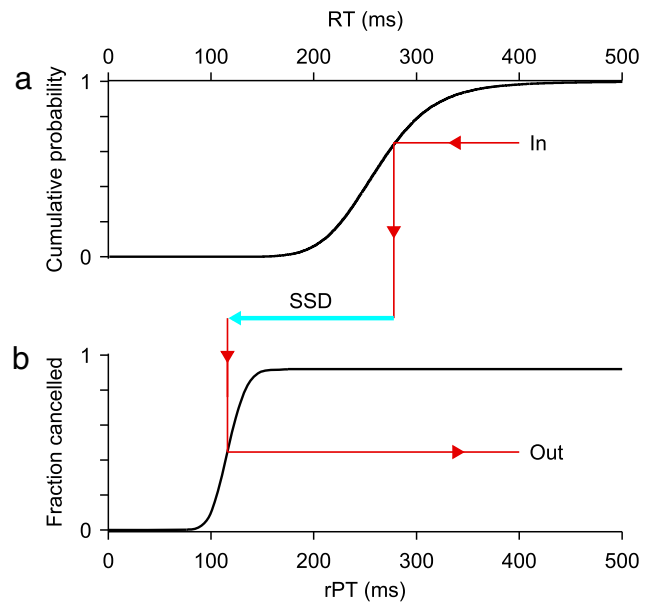


Figure 11. Performance in the countermanding task is completely specified by two psychophysical curves. *a*, Cumulative RT distribution from no-stop trials. *b*, Empirical tachometric curve. Red and blue arrows indicate how a randomly drawn RT (278 ms) determines, after subtracting the SSD (162 ms), a specific rRT (116), which in turn determines the probability of a correct cancellation (0.45).

discrepancy is minuscule). Thus, by iterating this process, one can calculate the fraction of cancelled trials at any SSD, or the RT distributions for noncancelled trials, or any other statistic based on the subject's responses.

This analysis is interesting for two reasons: (1) it demonstrates that performance can be divided almost perfectly into a purely motor process that determines the distribution of RTs, and a second process that reflects perceptual capacity (for stimulus detection) and other cognitive elements (interpreting the stimulus as a stop command; inhibitory action); and (2) this functional characterization is independent of the cancellable rise-to-threshold model, but the two are consistent, in that four parameters of the model (μ_B , σ_B , μ_T , σ_T) determine the resulting distribution of RTs and the other four (μ_S , σ_S , τ , p_e) determine the empirical tachometric curve. Thus, regardless of the mechanistic interpretations of their parameters, simpler or more complex models would be those that use either fewer or more parameters to describe the two fundamental curves.

Distinguishing alternative models

By means of the cancellable rise-to-threshold model, we have argued that perceptual mechanisms are key for interpreting numerous experimental results obtained with the countermanding task and that the tachometric curve provides a more robust assessment of performance than does the SSRT. These observations, however, do not directly speak to the validity of the original go/stop race model (Logan and Cowan, 1984; Boucher et al., 2007; Verbruggen and Logan, 2009). In the race framework, stop-related activity rises to threshold just like saccade-related activity; the eye movement is executed if the motor plan reaches threshold first, or it is cancelled if the stopping plan reaches threshold first. Inhibition acts fully or not at all, and the two competing processes do not interact until the winner is determined. This is indeed quite different from the gradual deceleration proposed here. So, beyond the existence of a threshold-sensitive trigger for inhibition, can the two models be distinguished based on their respec-

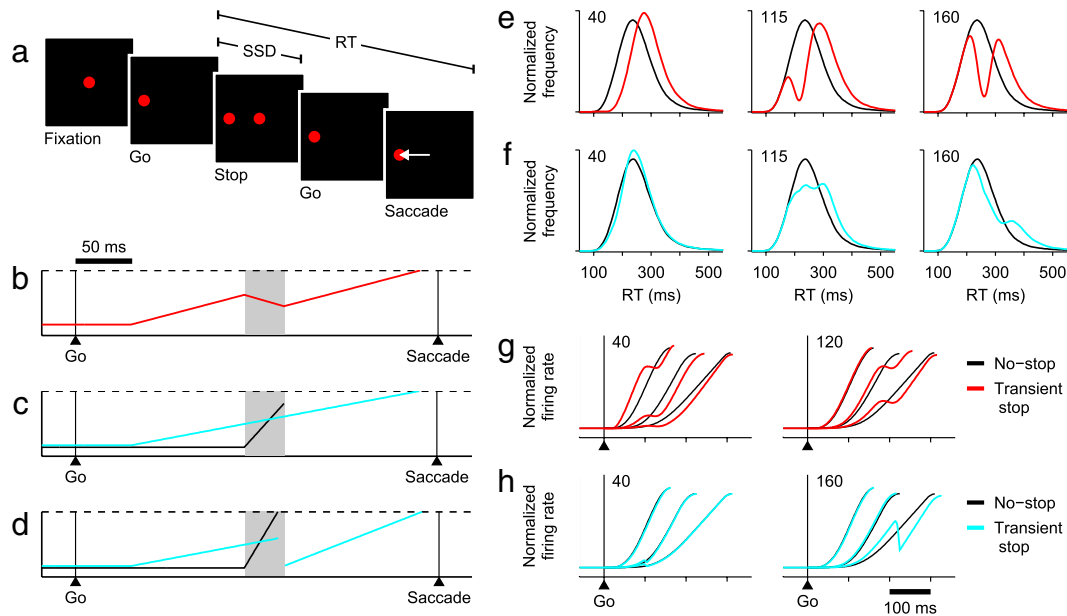


Figure 12. Contrasting predictions of two alternative models. *a*, A transient-stop trial. The stop signal is shown briefly, and the go instruction resumes after it. The correct response is a saccade. *b*, A transient-stop trial of the cancellable rise-to-threshold model. The gray shadow represents the period (after the SSD and the afferent delay) during which the stop signal is known to the model circuit and the motor plan may decelerate. *c*, A transient-stop trial of the go/stop race model in which the stopping plan (black trace) does not reach threshold. The motor plan (blue trace) follows its original trajectory unperturbed. *d*, A trial of the race model in which the stopping plan reaches threshold. The initial motor plan is cancelled, and a new one is issued after the stop command ends. *e*, Distributions of RTs in transient-stop (red traces) and no-stop (black traces) trials predicted by the cancellable rise-to-threshold model. Corresponding SSDs, in milliseconds, are indicated (upper left). *f*, Same as in *e*, but for the race model. *g*, Mean neural responses predicted by the cancellable rise-to-threshold model. Data are aligned on the initial go signal. In each plot, the three curves of each color are from three groups of trials sorted by RT, short ($RT \leq 200$ ms), intermediate ($RT 200 - 300$ ms), and long ($RT > 300$ ms), going from left to right. For clarity, only the rising segments of the responses are shown. Corresponding SSDs are indicated. *h*, Same as in *g*, but for the race model.

tive saccade-related responses? In this section, we discuss a potential experiment that leads to contrasting behavioral and neurophysiological predictions, and thus serves to highlight the main mechanistic differences between the two models.

The proposed experiment is a simple variant of the countermanding task in which, on some trials, the stop signal appears just briefly and is then itself countermanded (Fig. 12*a*). In these transient-stop trials, a saccade is always produced, and so the resulting RT distributions can be compared with those of no-stop trials. Crucially, the transient stop signal has very different consequences for the two models.

In the cancellable rise-to-threshold model, the developing motor plan is decelerated for the duration of the stop command, but the motor activity continues to rise once the influence of the stop signal ceases (Fig. 12*b*). How exactly and how much the firing rate drops during the transient stop (gray shade) is relatively unimportant; the key is that the trajectory toward threshold resumes with the same buildup rate as before, so the effect of the transient stop signal is essentially to interrupt momentarily the ongoing motor plan. As a consequence of this interruption, all the saccades that are not initiated before the stop command reaches the model circuit are delayed, and this manifests as a characteristic, SSD-dependent shift in the distribution of RTs (Fig. 12*e*). When the SSD is short (Fig. 12*e*, 40 ms), nearly all RTs are lengthened and the full distribution shifts to the right; and when the SSD is long (Fig. 12*e*, 160 ms), a large fraction of the saccades are initiated before the transient stop has had an effect; thus, only the longer RTs are shifted, which results in a conspicuous dip in the RT distribution.

By contrast, in the classic race model (see Materials and Methods) (Logan and Cowan, 1984), the motor and stopping plans remain independent up until the moment when the stop-related

activity crosses threshold; hence, the transient stop command has an all-or-none effect. In trials in which the stopping plan builds up relatively slowly, it never reaches threshold, so the motor plan proceeds undisturbed (Fig. 12*c*). In contrast, when the stopping plan builds up rapidly, it reaches threshold and cancels the motor plan; then, after the stop command ends, the motor plan has to be restarted (Fig. 12*d*). Here it does not matter much whether the renewed motor plan proceeds with the same or with a new, independent buildup rate, or how fast the saccade-related activity drops once the cancellation is issued; what is key is that, for just a fraction of the trials, the progress made before the stopping plan crosses threshold is completely erased and therefore has no impact on what happens subsequently. The effect of these dynamics on the resulting RT distributions (Fig. 12*f*) is subtle because for brief stop signals the fraction of inhibited trials is always small. Clearly, however, in this case, there is no shift at short SSDs and no dip at long SSDs.

These differential effects on the RT distributions have corresponding neural correlates in the expected oculomotor activity leading to a saccade. To see this, simulated motor plans were separated into short-, intermediate-, and long-RT groups and were then averaged within each group (Fig. 12*g,h*). According to the cancellable rise-to-threshold model, a kink in the average motor plans should be evident in all RT groups when comparing transient-stop versus no-stop trials, and the timing of the kink should depend systematically on the SSD (Fig. 12*g*). In contrast, the race model predicts either no effect or a very weak effect in all cases, except when both the RT and the SSD are relatively long (Fig. 12*h*). These differences between models are quite robust to parameter variations.

There are two important considerations about these comparisons. First, the results apply to the standard race model, in which

the motor and stopping plans develop independently of each other and the SSRT is clearly defined (it is the time interval between the onset of the stop signal and the moment when the stopping plan crosses threshold). The model can be modified to include a continuous interaction between the two processes, for instance, through mutual inhibition (Boucher et al., 2007). When this is done, however, the model changes quite drastically: the motor plan is suppressed without the stopping plan ever having to reach a threshold, so the interaction no longer corresponds to that of a race; furthermore, the event that defines the SSRT is gone. Indeed, the result is a gradual deceleration of the motor plan, much as in our case, except with a different parameterization that corresponds to a specific form of inhibitory neural activity. Thus, the threshold-sensitive inhibitory trigger is the key dynamic element to test.

Second, the psychophysical predictions based on the brief deceleration or interruption of an ongoing motor plan (Fig. 12*e*) are quite distinct and very likely to be correct. It is known that a saccade to a visual target can be delayed by a nearly simultaneous distracter or transient change in the visual scene. This phenomenon, known as “saccadic inhibition” (or the “remote distracter effect”), produces shifts and dips in the saccadic RT distributions precisely like those in Figure 12*e* and occurs over a wide range of experimental conditions (Reingold and Stampe, 2002; Buonocore and McIntosh, 2008, 2012; Bompas and Sumner, 2011). This supports the idea that perceptual information informs oculomotor plans as they develop, as proposed by the cancellable rise-to-threshold model; and if the model predictions are correct, it would suggest that, with respect to imminent saccades, distracters may act as transient stop signals.

Discussion

In the cancellable rise-to-threshold model, the outcome of each trial of the countermanding task depends fundamentally on three factors: the initial buildup rate of the saccadic plan, the SSD, and the perceptual speed of the subject. Their interplay determines when and how fast a developing oculomotor plan can be stopped. Rather than making mechanistic assumptions about inhibitory activity, this model only takes into account its ultimate consequence, the deceleration of the ongoing motor plan. We do not preclude the possibility that task manipulations may cause direct changes to inhibitory processes; we simply argue that the results of most previous manipulations can be explained parsimoniously by changes in distinct premotor and perceptual mechanisms that must act first.

The tachometric curve gives the average amount of stimulus viewing time necessary to achieve a given success rate. Strictly speaking, the curve is not a pure measure of perceptual capacity because success implies not only detecting a stimulus, but also interpreting it as a command to stop and activating the appropriate inhibitory circuitry. Thus, true changes in inhibition would indeed have an impact on the curve. Still, this curve is crucial because, in contrast to the RT, it is insensitive to motor execution; and in contrast to the SSRT, it provides a more complete, model-free characterization of performance.

Generality of the modeling framework

The model developed here is essentially a simplified version of a race-to-threshold model that matches psychophysical and neurophysiological data in the compelled-saccade task, a choice task that simulates an urgent decision (Salinas et al., 2010; Stanford et al., 2010; Shankar et al., 2011). Thus, the phenomenological description presented here is consistent with a wider set of experi-

ments. The current analyses can also be adapted to other tasks designed to elucidate various issues in visuomotor control, in which (1) the instruction to execute an action precedes the sensory cue that determines whether that action is the correct one and (2) the relevant variable is processing time (e.g., Becker and Jürgens, 1979; Camalier et al., 2007; Farooqui et al., 2011; Noorani et al., 2011).

We inferred that performance in the countermanding task is based on a stimulus-detection process that starts ~50–100 ms after the stop signal is given and develops over the course of ~20 ms. This reinforces earlier findings indicating that perceptual processes effectively develop very quickly and may affect motor plans that are already ongoing (Cisek, 2006; Cisek and Kalaska, 2010; Stanford et al., 2010; Shankar et al., 2011).

Interplay between neural mechanisms

The majority of the experimental results reproduced by our model were based on two distinct neural mechanisms: (1) variations in the average buildup rate of the motor plan, and (2) variations in response latency associated with detection of the target and stop signals. The former may be adjusted across blocks of trials (Fig. 8) or from one trial to the next (Fig. 7), and provide a straightforward handle for regulating the timing of motor signals. Perhaps more surprisingly, the latter also play a key role; they depend both on modality (Fig. 9) and on task conditions (Fig. 7), and as explained earlier, may represent another common strategy for controlling the timing of actions.

Thus, performance in the countermanding task depends on the level of commitment of the saccadic plans and on the speed of the perceptual detection and subsequent deceleration. This also explains the results of Leotti and Wager (2010), who manipulated the monetary rewards and penalties associated with correct and incorrect responses to precisely control each subject's speed-accuracy trade-off. They found that a higher priority for stops versus fast saccades produced more cancellations in stop trials, higher RTs in no-stop trials, and a decrease in SSRT. According to our model, these findings require both types of neural mechanism: (1) a decrease in the average buildup rates, to account for the higher RTs; and (2) either a decrease in the latency of the stop signal, or an increase in the maximum value of the tachometric curve, or both, to account for the decrease in SSRT. This combination of effects may seem contradictory because it simultaneously involves slower responses (higher RTs) and faster detection/deceleration (lower SSRTs), but indeed it makes sense because all three changes (i.e., lower buildup rates, faster detection, and fewer lapses) increase the probability of a successful cancellation.

Interestingly, manipulations of motivation in the compelled-saccade task led to very similar results: compared with a condition in which a small reward was at stake, when monkey subjects had the opportunity to earn a large reward—and presumably were highly motivated—they slowed down their responses (higher RTs), initiated the color discrimination sooner (leftward shift of the tachometric curve), and made fewer errors at long viewing times (increased asymptotic value of the tachometric curve), and each of these changes tended to improve performance (Shankar et al., 2011). This suggests that the neural mechanisms engaged during performance of these two tasks are rather similar, and perhaps are generally the same for other tasks involving fast responses.

Limitations of the SSRT

The countermanding task is often used to characterize or diagnose anomalies in the mechanisms underlying the capacity to

inhibit a response (Armstrong and Munoz, 2003; Nouraei et al., 2003; DeHaan et al., 2007; Lipszyc and Schachar, 2010; Eagle et al., 2011; Farooqui et al., 2011; Thakkar et al., 2011; Ashare and Hawk, 2012). In this context, the SSRT is the key quantity being monitored, and it is interpreted as a direct indication of the amount of time required by the corresponding inhibitory process to act. Instead, we found that the SSRT correlates closely with the timing of the perceptual detection process, as revealed by the tachometric curve (Fig. 6c). The correspondence is not perfect, though. It is possible, for instance, to observe a large increase in SSRT, suggesting a delayed stopping process, when the true effect is exactly the opposite and the saccade cancellation is, indeed, initiated substantially earlier (Fig. 9f).

Although this particular problem can be ameliorated by taking into account the possibility of cognitive lapses in the calculation of the SSRT (Tannock et al., 1989), in general the empirical tachometric curve is a much better benchmark of performance than the SSRT, for three reasons: (1) the tachometric curve is explicitly based on processing time, the variable that primarily determines the probability of success in each stop trial; (2) the curve reveals the onset and the speed of the perceptual process, as well as the rate of cognitive/attentional errors; it simply provides more information than a single number; and (3) the numerical value of the SSRT depends on the method chosen to compute it (e.g., Hanes et al., 1998; Kornyló et al., 2003; Stevenson et al., 2009; Verbruggen and Logan, 2009). Different methods produce different results, and there simply is no principled way to resolve the ambiguity.

Evidence of inhibitory action

One may wonder whether the SSRT truly has a specific behavioral or neural manifestation (Walton and Gandhi, 2006; Goonetilleke et al., 2011). This, however, is not essential to the current results. As far as we know, all previous models of the countermanding task propose specific neural interactions for cancelling the developing motor plan (Logan and Cowan, 1984; Boucher et al., 2007; Lo et al., 2009; Verbruggen and Logan, 2009; Noorani et al., 2011), the most common element being a threshold-sensitive trigger for inhibition. But experimental evidence of such a trigger is lacking, and in many cases detailed mechanistic assumptions may be unnecessary. Our results show that, to obtain an accurate phenomenological description of countermanding performance, it is enough to specify that the turnaround in activity occurs with a particular magnitude and time constant. The deceleration of the preparatory activity may be the result of complex underlying neural dynamics, but the biophysical details are not needed to describe it parametrically. The fact is that the suppression of the impending motor plan requires at least two prior steps: the perceptual detection of the stop stimulus and the interpretation of such stimulus as a stop command—and then inhibition may kick in. Further, numerous experimental manipulations can indeed be explained by the interplay between perceptual and motor-planning processes that must be engaged during task performance. Therefore, any inference about inhibitory control in the countermanding task is potentially confounded by effects on perception.

A recent imaging study explicitly attempted to avoid such a confound (Sharp et al., 2010). Using a manual version of the countermanding task, the idea was to compare activity during cancelled trials not to activity in no-stop trials, which have different sensory and attentional demands, but to activity in trials in which a sensory cue similar to the stop is detected but does not instruct the subject to stop (e.g., stop if the reappearing fixation

spot is red but not if it is green). Under these conditions, only one area, the presupplementary motor area, displayed significantly higher activation that could be interpreted as related to inhibitory control, consistent with other results (Stuphorn and Schall, 2006; Aron, 2011; Boehler et al., 2011). Such comparisons, as well as different task designs (Ross et al., 1994; Narayanan and Laubach, 2006; Narayanan et al., 2006; Hayton et al., 2010; Aron, 2011), may be necessary to effectively disambiguate inhibitory and perceptual signals in future studies.

References

- Akerfelt A, Colonius H, Diederich A (2006) Visual-tactile saccadic inhibition. *Exp Brain Res* 169:554–563. [CrossRef Medline](#)
- Armstrong IT, Munoz DP (2003) Inhibitory control of eye movements during oculomotor countermanding in adults with attention-deficit hyperactivity disorder. *Exp Brain Res* 152:444–452. [CrossRef Medline](#)
- Aron AR (2011) From reactive to proactive and selective control: developing a richer model for stopping inappropriate responses. *Biol Psychiatry* 69:e55–e68. [CrossRef Medline](#)
- Ashare RL, Hawk LW Jr (2012) Effects of smoking abstinence on impulsive behavior among smokers high and low in ADHD-like symptoms. *Psychopharmacology (Berl)* 219:537–547. [CrossRef Medline](#)
- Barracough DJ, Conroy ML, Lee D (2004) Prefrontal cortex and decision making in a mixed-strategy game. *Nat Neurosci* 7:404–410. [CrossRef Medline](#)
- Becker W, Jürgens R (1979) An analysis of the saccadic system by means of double step stimuli. *Vision Res* 19:967–983. [CrossRef Medline](#)
- Boehler CN, Appelbaum LG, Krebs RM, Chen LC, Woldorff MG (2011) The role of stimulus salience and attentional capture across the neural hierarchy in a stop-signal task. *PLoS One* 6:e26386. [CrossRef Medline](#)
- Bompas A, Sumner P (2011) Saccadic inhibition reveals the timing of automatic and voluntary signals in the human brain. *J Neurosci* 31:12501–12512. [CrossRef Medline](#)
- Boucher L, Palmeri TJ, Logan GD, Schall JD (2007) Inhibitory control in mind and brain: an interactive race model of countermanding saccades. *Psychol Rev* 114:376–397. [CrossRef Medline](#)
- Brown JW, Hanes DP, Schall JD, Stuphorn V (2008) Relation of frontal eye field activity to saccade initiation during a countermanding task. *Exp Brain Res* 190:135–151. [CrossRef Medline](#)
- Brown SD, Heathcote A (2008) The simplest complete model of choice response time: linear ballistic accumulation. *Cogn Psychol* 57:153–178. [Medline](#)
- Bruce CJ, Goldberg ME (1985) Primate frontal eye fields: I. Single neurons discharging before saccades. *J Neurophysiol* 53:603–635. [Medline](#)
- Buonocore A, McIntosh RD (2008) Saccadic inhibition underlies the remote distractor effect. *Exp Brain Res* 191:117–122. [CrossRef Medline](#)
- Buonocore A, McIntosh RD (2012) Modulation of saccadic inhibition by distractor size and location. *Vision Res* 69:32–41. [CrossRef Medline](#)
- Busse L, Ayaz A, Dhruv NT, Katzner S, Saleem AB, Schölvinck ML, Zaharia AD, Carandini M (2011) The detection of visual contrast in the behaving mouse. *J Neurosci* 31:11351–11361. [CrossRef Medline](#)
- Cabel DW, Armstrong IT, Reingold E, Munoz DP (2000) Control of saccade initiation in a countermanding task using visual and auditory stop signals. *Exp Brain Res* 133:431–441. [CrossRef Medline](#)
- Camalier CR, Gotler A, Murthy A, Thompson KG, Logan GD, Palmeri TJ, Schall JD (2007) Dynamics of saccade target selection: race model analysis of double step and search step saccade production in human and macaque. *Vision Res* 47:2187–2211. [CrossRef Medline](#)
- Carpenter RH, Williams ML (1995) Neural computation of log likelihood in control of saccadic eye movements. *Nature* 377:59–62. [CrossRef Medline](#)
- Cavina-Pratesi C, Bricolo E, Prior M, Marzi CA (2001) Redundancy gain in the stop-signal paradigm: implications for the locus of coactivation in simple reaction time. *J Exp Psychol Hum Percept Perform* 27:932–941. [CrossRef Medline](#)
- Cisek P (2006) Integrated neural processes for defining potential actions and deciding between them: a computational model. *J Neurosci* 26:9761–9770. [CrossRef Medline](#)
- Cisek P, Kalaska JF (2010) Neural mechanisms for interacting with a world full of action choices. *Annu Rev Neurosci* 33:269–298. [CrossRef Medline](#)
- Colonius H, Ozzyurt J, Arndt PA (2001) Countermanding saccades with au-

- ditory stop signals: testing the race model. *Vision Res* 41:1951–1968. [CrossRef Medline](#)
- DeHaan A, Halterman C, Langan J, Drew AS, Osternig LR, Chou LS, van Donkelaar P (2007) Cancelling planned actions following mild traumatic brain injury. *Neuropsychologia* 45:406–411. [CrossRef Medline](#)
- Dorris MC, Paré M, Munoz DP (1997) Neuronal activity in monkey superior colliculus related to the initiation of saccadic eye movements. *J Neurosci* 17:8566–8579. [Medline](#)
- Eagle DM, Wong JC, Allan ME, Mar AC, Theobald DE, Robbins TW (2011) Contrasting roles for dopamine D1 and D2 receptor subtypes in the dorsomedial striatum but not the nucleus accumbens core during behavioral inhibition in the stop-signal task in rats. *J Neurosci* 31:7349–7356. [CrossRef Medline](#)
- Emeric EE, Brown JW, Boucher L, Carpenter RH, Hanes DP, Harris R, Logan GD, Mashru RN, Paré M, Pouget P, Stuphorn V, Taylor TL, Schall JD (2007) Influence of history on saccade countermanding performance in humans and macaque monkeys. *Vision Res* 47:35–49. [CrossRef Medline](#)
- Everling S, Munoz DP (2000) Neuronal correlates for preparatory set associated with prosaccades and anti-saccades in the primate frontal eye field. *J Neurosci* 20:387–400. [Medline](#)
- Farooqui AA, Bhutani N, Kulashkhar S, Behari M, Goel V, Murthy A (2011) Impaired conflict monitoring in Parkinson's disease patients during an oculomotor redirect task. *Exp Brain Res* 208:1–10. [CrossRef Medline](#)
- Fecteau JH, Munoz DP (2003) Exploring the consequences of the previous trial. *Nat Rev Neurosci* 4:435–443. [CrossRef Medline](#)
- Fuhrmann Alpert G, Hein G, Tsai N, Naumer MJ, Knight RT (2008) Temporal characteristics of audiovisual information processing. *J Neurosci* 28:5344–5349. [CrossRef Medline](#)
- Gibbon J, Rutschmann R (1969) Temporal order judgment and reaction time. *Science* 165:413–415. [CrossRef Medline](#)
- Goonetilleke SC, Doherty TJ, Corneil BD (2010) A within trial measure of the stop signal reaction time in a head-unrestrained oculomotor countermanding task. *J Neurophysiol* 104:3677–3690. [CrossRef Medline](#)
- Hanes DP, Schall JD (1996) Neural control of voluntary movement initiation. *Science* 274:427–430. [CrossRef Medline](#)
- Hanes DP, Patterson WF 2nd, Schall JD (1998) Role of frontal eye fields in countermanding saccades: visual, movement, and fixation activity. *J Neurophysiol* 79:817–834. [Medline](#)
- Hayton SJ, Lovett-Barron M, Dumont EC, Olmstead MC (2010) Target-specific encoding of response inhibition: increased contribution of AMPA to NMDA receptors at excitatory synapses in the prefrontal cortex. *J Neurosci* 30:11493–11500. [CrossRef Medline](#)
- Heitz RP, Schall JD (2012) Neural mechanisms of speed-accuracy tradeoff. *Neuron* 76:616–628. [CrossRef Medline](#)
- Idiart MAP, Abbott LF (1993) Propagation of excitation in neural network models. *Network* 4:285–294. [CrossRef](#)
- Ito S, Stuphorn V, Brown JW, Schall JD (2003) Performance monitoring by the anterior cingulate cortex during saccade countermanding. *Science* 302:120–122. [CrossRef Medline](#)
- Jacobson L, Javitt DC, Lavidor M (2011) Activation of inhibition: diminishing impulsive behavior by direct current stimulation over the inferior frontal gyrus. *J Cogn Neurosci* 23:3380–3387. [CrossRef Medline](#)
- Kornylo K, Dill N, Saenz M, Krauzlis RJ (2003) Cancelling of pursuit and saccadic eye movements in humans and monkeys. *J Neurophysiol* 89:2984–2999. [CrossRef Medline](#)
- Leblois A, Bodor AL, Person AL, Perkel DJ (2009) Millisecond timescale disinhibition mediates fast information transmission through an avian basal ganglia loop. *J Neurosci* 29:15420–15433. [CrossRef Medline](#)
- Leotti LA, Wager TD (2010) Motivational influences on response inhibition measures. *J Exp Psychol Hum Percept Perform* 36:430–447. [CrossRef Medline](#)
- Lipszyc J, Schachar R (2010) Inhibitory control and psychopathology: a meta-analysis of studies using the stop signal task. *J Int Neuropsychol Soc* 16:1064–1076. [CrossRef Medline](#)
- Lo CC, Wang XJ (2006) Cortico-basal ganglia circuit mechanism for a decision threshold in reaction time tasks. *Nat Neurosci* 9:956–963. [CrossRef Medline](#)
- Lo CC, Boucher L, Paré M, Schall JD, Wang XJ (2009) Proactive inhibitory control and attractor dynamics in countermanding action: a spiking neural circuit model. *J Neurosci* 29:9059–9071. [CrossRef Medline](#)
- Logan GD, Cowan WB (1984) On the ability to inhibit thought and action: a theory of an act of control. *Psychol Rev* 91:295–327. [CrossRef](#)
- Luce RD (1986) Response times: their role in inferring elementary mental organization. Oxford: Oxford UP.
- Marino RA, Munoz DP (2009) The effects of bottom-up target luminance and top-down spatial target predictability on saccadic reaction times. *Exp Brain Res* 197:321–335. [CrossRef Medline](#)
- Maunsell JH, Gibson JR (1992) Visual response latencies in striate cortex of the macaque monkey. *J Neurophysiol* 68:1332–1344. [Medline](#)
- Mordkoff JT, Yantis S, Egeth HE (1990) Detecting conjunctions of color and form in parallel. *Percept Psychol* 48:157–168. [CrossRef Medline](#)
- Munoz DP, Dorris MC, Paré M, Everling S (2000) On your mark, get set: brainstem circuitry underlying saccadic initiation. *Can J Physiol Pharmacol* 78:934–944. [CrossRef Medline](#)
- Narayanan NS, Laubach M (2006) Top-down control of motor cortex ensembles by dorsomedial prefrontal cortex. *Neuron* 52:921–931. [CrossRef Medline](#)
- Narayanan NS, Horst NK, Laubach M (2006) Reversible inactivations of rat medial prefrontal cortex impair the ability to wait for a stimulus. *Neuroscience* 139:865–876. [CrossRef Medline](#)
- Nelson MJ, Boucher L, Logan GD, Palmeri TJ, Schall JD (2010) Nonindependent and nonstationary response times in stopping and stepping saccade tasks. *Atten Percept Psychophys* 72:1913–1929. [CrossRef Medline](#)
- Noorani I, Gao MJ, Pearson BC, Carpenter RH (2011) Predicting the timing of wrong decisions with LATER. *Exp Brain Res* 209:587–598. [CrossRef Medline](#)
- Nouraei SA, De Pennington N, Jones JG, Carpenter RH (2003) Dose-related effect of sevoflurane sedation on higher control of eye movements and decision making. *Br J Anaesth* 91:175–183. [CrossRef Medline](#)
- Oram MW (2010) Contrast induced changes in response latency depend on stimulus specificity. *J Physiol Paris* 104:167–175. [CrossRef Medline](#)
- Paré M, Hanes DP (2003) Controlled movement processing: superior colliculus activity associated with countermanded saccades. *J Neurosci* 23:6480–6489. [Medline](#)
- Port NL, Wurtz RH (2009) Target selection and saccade generation in monkey superior colliculus. *Exp Brain Res* 192:465–477. [CrossRef Medline](#)
- Pouget P, Logan GD, Palmeri TJ, Boucher L, Paré M, Schall JD (2011) Neural basis of adaptive response time adjustment during saccade countermanding. *J Neurosci* 31:12604–12612. [CrossRef Medline](#)
- Ramautar JR, Kok A, Ridderinkhof KR (2004) Effects of stop-signal probability in the stop-signal paradigm: the N2/P3 complex further validated. *Brain Cogn* 56:234–252. [CrossRef Medline](#)
- Ratcliff R, Hasegawa YT, Hasegawa RP, Smith PL, Segraves MA (2007) Dual diffusion model for single-cell recording data from the superior colliculus in a brightness-discrimination task. *J Neurophysiol* 97:1756–1774. [CrossRef Medline](#)
- Ray S, Pouget P, Schall JD (2009) Functional distinction between visuo-movement and movement neurons in macaque frontal eye field during saccade countermanding. *J Neurophysiol* 102:3091–3100. [CrossRef Medline](#)
- Reddi BA, Carpenter RH (2000) The influence of urgency on decision time. *Nat Neurosci* 3:827–830. [CrossRef Medline](#)
- Reingold EM, Stampe DM (2002) Saccadic inhibition in voluntary and reflexive saccades. *J Cogn Neurosci* 14:371–388. [CrossRef Medline](#)
- Roitman JD, Shadlen MN (2002) Response of neurons in the lateral intraparietal area during a combined visual discrimination reaction time task. *J Neurosci* 22:9475–9489. [Medline](#)
- Rorie AE, Gao J, McClelland JL, Newsome WT (2010) Integration of sensory and reward information during perceptual decision-making in lateral intraparietal cortex (LIP) of the macaque monkey. *PLoS One* 5:e9308. [CrossRef Medline](#)
- Ross RG, Hommer D, Breiger D, Varley C, Radant A (1994) Eye movement task related to frontal lobe functioning in children with attention deficit disorder. *J Am Acad Child Adolesc Psychiatry* 33:869–874. [CrossRef Medline](#)
- Rowland BA, Quessy S, Stanford TR, Stein BE (2007) Multisensory integration shortens physiological response latencies. *J Neurosci* 27:5879–5884. [CrossRef Medline](#)
- Salinas E, Shankar S, Costello MG, Zhu D, Stanford TR (2010) Waiting is the hardest part: comparison of two computational strategies for performing a compelled-response task. *Front Comput Neurosci* 4:153. [CrossRef Medline](#)
- Saslow MG (1967) Effects of components of displacement-step stimuli

- upon latency for saccadic eye movement. *J Opt Soc Am* 57:1024–1029. [CrossRef Medline](#)
- Shankar S, Massoglia DP, Zhu D, Costello MG, Stanford TR, Salinas E (2011) Tracking the temporal evolution of a perceptual judgment using a compelled-response task. *J Neurosci* 31:8406–8421. [CrossRef Medline](#)
- Sharp DJ, Bonnelle V, De Boissezon X, Beckmann CF, James SG, Patel MC, Mehta MA (2010) Distinct frontal systems for response inhibition, attentional capture, and error processing. *Proc Natl Acad Sci U S A* 107:6106–6111. [CrossRef Medline](#)
- Sparks D, Rohrer WH, Zhang Y (2000) The role of the superior colliculus in saccade initiation: a study of express saccades and the gap effect. *Vision Res* 40:2763–2777. [CrossRef Medline](#)
- Stanford TR, Shankar S, Massoglia DP, Costello MG, Salinas E (2010) Perceptual decision making in less than 30 milliseconds. *Nat Neurosci* 13:379–385. [CrossRef Medline](#)
- Stevenson SA, Elsley JK, Corneil BD (2009) A “gap effect” on stop signal reaction times in a human saccadic countermanding task. *J Neurophysiol* 101:580–590. [CrossRef Medline](#)
- Stuphorn V, Schall JD (2006) Executive control of countermanding saccades by the supplementary eye field. *Nat Neurosci* 9:925–931. [CrossRef Medline](#)
- Sugrue LP, Corrado GS, Newsome WT (2004) Matching behavior and the representation of value in the parietal cortex. *Science* 304:1782–1787. [CrossRef Medline](#)
- Takikawa Y, Kawagoe R, Itoh H, Nakahara H, Hikosaka O (2002) Modulation of saccadic eye movements by predicted reward outcome. *Exp Brain Res* 142:284–291. [CrossRef Medline](#)
- Tannock R, Schachar RJ, Carr RP, Chajczyk D, Logan GD (1989) Effects of methylphenidate on inhibitory control in hyperactive children. *J Abnorm Child Psychol* 17:473–491. [CrossRef Medline](#)
- Thakkar KN, Schall JD, Boucher L, Logan GD, Park S (2011) Response inhibition and response monitoring in a saccadic countermanding task in schizophrenia. *Biol Psychiatry* 69:55–62. [CrossRef Medline](#)
- Thompson KG, Hanes DP, Bichot NP, Schall JD (1996) Perceptual and motor processing stages identified in the activity of macaque frontal eye field neurons during visual search. *J Neurophysiol* 76:4040–4055. [Medline](#)
- van Rossum MC, van der Meer MA, Xiao D, Oram MW (2008) Adaptive integration in the visual cortex by depressing recurrent cortical circuits. *Neural Comput* 20:1847–1872. [CrossRef Medline](#)
- van Vreeswijk C, Sompolinsky H (1996) Chaos in neuronal networks with balanced excitatory and inhibitory activity. *Science* 274:1724–1726. [CrossRef Medline](#)
- Verbruggen F, Logan GD (2009) Models of response inhibition in the stop-signal and stop-change paradigms. *Neurosci Biobehav Rev* 33:647–661. [CrossRef Medline](#)
- Walton MM, Gandhi NJ (2006) Behavioral evaluation of movement cancellation. *J Neurophysiol* 96:2011–2024. [CrossRef Medline](#)
- Welford AT (1980) Reaction times. New York: Academic.
- York LC, van Rossum MC (2009) Recurrent networks with short term synaptic depression. *J Comput Neurosci* 27:607–620. [CrossRef Medline](#)