

The Stop-Signal Paradigm

Dora Matzke¹, Frederick Verbruggen², and Gordon D. Logan³

¹UNIVERSITY OF AMSTERDAM, ²UNIVERSITY OF EXETER,

³VANDERBILT UNIVERSITY

Abstract

Response inhibition—the ability to stop responses that are no longer appropriate—is frequently studied with the stop-signal paradigm. In the stop-signal paradigm, participants perform a choice response time task that is occasionally interrupted by a stop signal. The stop signal prompts participants to withhold their response on that trial. Performance in the stop-signal paradigm is typically formalized as horse race between a go and a stop process. If the go process wins the race, the response is executed; if the stop process wins the race, the response is inhibited. The stop-signal paradigm owes its popularity to the underlying horse-race model that enables researchers to estimate the latency of the unobservable stop response. In this chapter, we present a theoretical review of the stop-signal paradigm and the corresponding horse-race model. First we focus on the standard independent horse-race model and discuss the most important measures of inhibitory control in the stop-signal paradigm. We then describe the latest developments in the model-based analysis of stop-signal data. We conclude the chapter with recommendations on how to run stop-signal experiments and how to report and interpret the results of stop-signal studies.

Key Terms: Horse-race model, Inhibition function, Response inhibition, Stop-signal paradigm, Stop-signal reaction time

DM is supported by a Veni grant (451-15-010) from the Netherlands Organization for Scientific Research (NWO).

FV is supported by a research grant from the European Research Council (ERC) under the European Union's Seventh Framework Programme (FP7/2007-2013)/ ERC Grant Agreement No. 312445.

GL is supported by a research grant from the National Eye Institute (R01 EY021833).

1. Introduction

Response inhibition is considered to be a key component of executive control (e.g., Aron, Robbins, & Poldrack, 2014; Logan, 1994; Miyake et al., 2000; Ridderinkhof, van den Wildenberg, Segalowitz, & Carter, 2004; Verbruggen, McLaren, & Chambers, 2014). The concept refers to the ability to suppress responses that are no longer required or inappropriate, which supports flexible and goal-directed behavior in ever-changing environments. In everyday life, there are many examples of the importance of response inhibition, such as stopping yourself from crossing a street when a car comes around the corner without noticing you, or withholding your reflex to grasp a hot pan falling from the stove. Furthermore, clinical research suggests that impairments in response inhibition may contribute to the development of a range of psychopathological and impulse-control disorders, such as attention deficit/hyperactivity disorder (ADHD), obsessive-compulsive disorder, substance abuse, pathological gambling, and eating disorders (e.g., Bechara, Noel, & Crone, 2006; Crews & Boettiger, 2009; de Wit, 2009; Fernie et al., 2013; Garavan & Stout, 2005; Nigg, 2001; Noël, Brevers, & Bechara, 2013). Response inhibition efficiency also correlates with the treatment outcome in people with such disorders (e.g., Nederkoorn, Jansen, Mulkens, & Jansen, 2007). Thus, response inhibition is crucial for flexible, adaptive, and goal-directed behavior.

A paradigm that is most suitable for the investigation of response inhibition in a laboratory setting is the stop-signal paradigm (Lappin & Eriksen, 1966; Logan & Cowan, 1984; Vince, 1948; for reviews, see Logan, 1994; Verbruggen & Logan, 2008b). In the standard stop-signal paradigm, participants usually perform a choice response time (RT) task (i.e., the *go task*; also referred to as the *primary task*), such as responding to the direction of an arrow (e.g., press a left key for a left-pointing arrow and a right key for a right-pointing arrow). Occasionally, the go stimulus is followed by a stop signal (e.g., an auditory tone or an additional visual stimulus) after a variable delay (*stop-signal-delay*; SSD), instructing subjects to withhold their response. Figure 1 depicts an example of the trial course of a stop-signal experiment. Typically, participants can inhibit their response when the stop signal is presented close to the moment of go stimulus presentation, but they cannot inhibit their response when the stop signal is presented close to the moment of response execution.

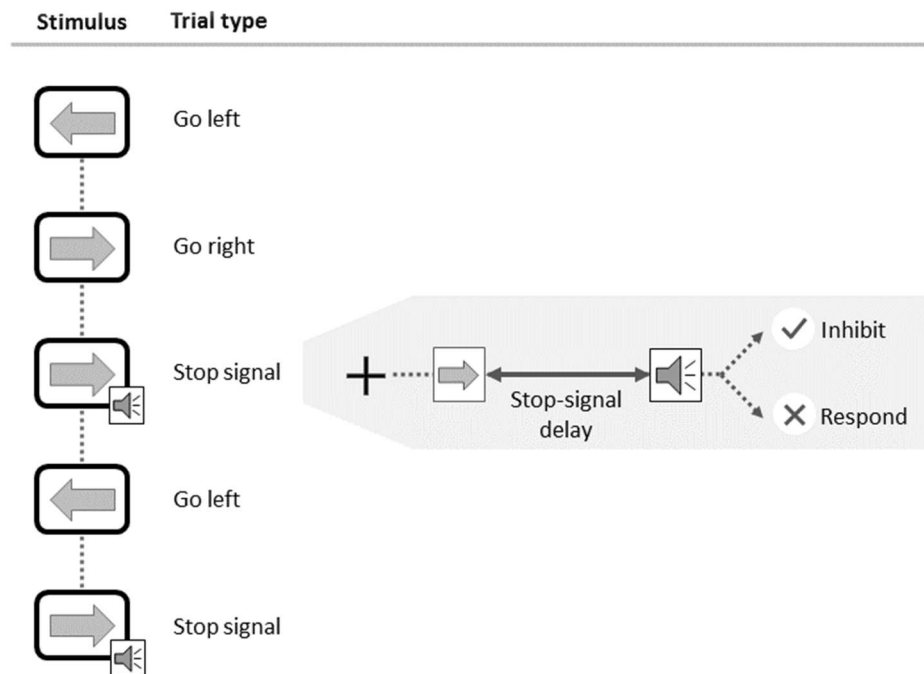


Figure 1. Depiction of a trial course in the stop-signal paradigm. In the go task, subjects respond to the direction of an arrow (the go stimulus; a “left arrow” requires a left response and a “right arrow” requires a right response). On a minority of the trials, the go stimulus is followed by an auditory stop signal after a variable stop-signal delay, instructing participants to withhold their response. Participants can successfully inhibit their response when the stop signal is presented close to the moment of go stimulus presentation, but they cannot inhibit their response when the stop signal is presented close to the moment of response execution.

The stop-signal paradigm is popular because it allows researchers to estimate the covert latency of the stop process: the *stop-signal reaction time* (SSRT). For example, SSRT has been used to explore the cognitive and neural mechanisms of response inhibition (e.g., Aron & Poldrack, 2006; Debey, De Schryver, Logan, Suchotzki, & Verschuere, 2015; Hanes, Patterson, & Schall, 1998; Logan & Cowan, 1984; van den Wildenberg, van der Molen, & Logan, 2002; Verbruggen, Stevens, & Chambers, 2014), the development and decline of inhibitory capacities across the life span (e.g., Chevalier, Chatham, & Munakata, 2014; Huizinga, Dolan, & van der Molen, 2006; Williams, Ponesse, Schachar, Logan, & Tannock, 1999), and correlations between individual differences in stopping and behaviors such as substance abuse, risk taking, and more generally, control of impulses and urges (e.g., Ersche et al., 2012; Schachar & Logan, 1990; Whelan et al., 2012). Furthermore, stop-signal studies have shown how response inhibition can be enhanced or impaired by a variety of factors, including motivational incentives, drugs, emotional stimuli, or neurological disorders (e.g., Aron, Fletcher, Bullmore, Sahaakian, & Robbins, 2003; Boehler, Schevernels,

Hopf, Stoppel, & Krebs, 2014; Fillmore, Rush, & Hays, 2002; Mulvihill, Skilling, & Vogel-Sprott, 1997; Tannock, Schachar, Carr, & Logan, 1989; Tannock, Schachar, & Logan, 1995; Verbruggen & De Houwer, 2007). These are just a few examples; for elaborate reviews, see Bari and Robbins (2013), Chambers, Garavan and Bellgrove (2009), Logan (1994), and Verbruggen and Logan (2008b).

SSRT can be estimated because performance in the stop-signal task can be formalized as an independent horse race between a go process, triggered by the presentation of the go stimulus, and a stop process, triggered by the presentation of the stop signal (Logan & Cowan, 1984; Logan, Van Zandt, Verbruggen, & Wagenmakers, 2014). When the stop process finishes before the go process, response inhibition is successful and no response is emitted; when the go process finishes before the stop process, response inhibition is unsuccessful and the response is incorrectly emitted.

The role of inhibitory processes in many executive control paradigms is debated (see e.g., MacLeod, Dodd, Sheard, Wilson, & Bibi, 2003), but most researchers have agreed that some kind of inhibition is involved in deliberately stopping a prepared motor response. The idea that responses are actively suppressed on stop-signal trials has received support from brain stimulation studies. These studies indicate that intracortical inhibitory circuits in primary motor cortex are recruited on stop-signal trials (e.g., Coxon, Stinear, & Byblow, 2006; van den Wildenberg et al., 2010). Furthermore, brain stimulation studies suggest that both task-relevant and irrelevant muscles are suppressed on stop-signal trials, indicating that stopping can have global effects on the motor system (Badry et al., 2009; Greenhouse, Oldenkamp, & Aron, 2011; Majid, Cai, George, Verbruggen, & Aron, 2012).

In this chapter, we present a theoretical review of the independent horse-race model and related models, and we discuss the most important measures of inhibitory control in the stop-signal paradigm. Up until Section 3.2, we focus on the standard independent horse-race model and related SSRT estimation techniques, and largely follow the structure and content of previous reviews by Logan (1994), Verbruggen and Logan (2008b), and Verbruggen and Logan (2009a). From Section 3.3 onwards, we describe the latest developments in the model-based analysis of stop-signal data, focusing on the estimation of SSRT distributions, process models of response inhibition, and variants of the stop-signal paradigm. We conclude the chapter with recommendations on how to run stop-signal experiments, and how to report and interpret findings from stop-signal studies.

2. Independent Horse-Race Model of

Response Inhibition

To formally account for performance in the stop-signal paradigm, Logan (1981) and Logan and Cowan (1984) formalized response inhibition as a horse race between two independent processes: a go process and a stop process. In this section, we briefly describe the precursors of the horse-race idea and then present the mathematical details of the independent horse-race model.

For simplicity, we first assume that SSRT is constant, but later we introduce the complete horse-race model that treats both go RTs and SSRTs as random variables. We assume throughout the chapter that the go process is entirely under the voluntary control of the participants, without a ballistic component that must run to completion once it has been launched, and therefore, cannot be inhibited. Although this is likely to be an unrealistic assumption, the contribution of ballistic processing to go RTs has been shown to be very brief and happen only very late in responding (e.g., de Jong, Coles, Logan, & Gratton, 1990; Logan & Cowan, 1984; McGarry & Franks, 1997; McGarry, Inglis, & Franks, 2000; Osman, Kornblum, & Meyer, 1986). Furthermore, we assume that the distribution of the stop signals is random and that stimuli in the go task are not consistently associated with stopping. (Note that this assumption is met in most stop-signal studies.) When the stimulus-stop mapping is consistent (e.g., when left arrows are always followed by a stop signal), participants can learn stimulus-stop associations (Verbruggen & Logan, 2008a; for a review, see Verbruggen, Best, Bowditch, Stevens, & McLaren, 2014). The retrieval of such associations will interfere with going and can influence SSRT estimates because responding may be suppressed before the stop signal is presented.

2.1 Early Horse-Race Models

The idea that response inhibition can be conceptualized as a race between two competing processes has been around well before Logan and Cowan's (1984) formal description of the horse-race model. The horse-race idea was qualitatively present in the work of Vince (1948) who observed that participants were unable to stop their responses to the go stimulus when the stop-signal delay was longer than 50 ms. Lappin and Eriksen (1966) used a visual stop-signal task and found that participant slowed their RT to the go stimulus in order to keep response rate constant across the stop-signal delays.

Although these findings suggest that participants' ability to stop is determined by the relative finishing times of their go and stop process, the formalization of response inhibition as a horse race had to await the work of

Ollman (1973), who applied the stop-signal procedure to a response timing task, in which participants were asked to produce a response of a given duration. Ollman proposed that participants perform the stop-signal task by setting a subjective deadline for the go response. If the stop signal is detected before the deadline, the go response is successfully inhibited; if the stop signal is detected after the deadline, the go response is incorrectly emitted. Ollman's model assumed that the finishing times of the go and the stop process follow a normal and exponential distribution, respectively. Although the model with its specific parametric assumptions was not supported by empirical data, Ollman's work paved the way for the quantitative description of response inhibition as a "horse race" between a go and a stop process, an idea that has dominated the literature even since.

2.2 Independent Horse-Race Model: The Basics

As mentioned above, the independent horse-race model (Logan, 1981; Logan & Cowan, 1984) assumes that response inhibition can be conceptualized as a race between two independent processes: a go process that is initiated by the go stimulus, and a stop process that is triggered by the stop signal. If the stop process finishes before the go process, the response is successfully inhibited; if the go process finishes before the stop process, the go response is erroneously emitted. Thus, the horse-race model posits that the outcome of response inhibition depends on the relative finishing times of the go and the stop process.

Logan and Cowan's (1984) conceptualization of response inhibition as a race between two competing processes is consistent with Ollman's (1973) model. Their horse-race model, however, is more general: It makes predictions about the interplay between RTs and response rate that do not require specifying the functional form of the go RT and SSRT distribution. The generality of the model and the precise mathematical description of the race allowed Logan and Cowan to develop distribution-free measures of the efficiency and the latency of the stop process (i.e., SSRT). This development has proved to be a milestone in the quantitative assessment of response inhibition in various scientific disciplines within as well as outside of psychology. For example, SSRT has been used in pharmacological, psychiatry, and neuroscience research (see the *Supplementary Information* of Verbruggen, Chambers, & Logan, 2013, for an overview of the different research areas).

The generality of the horse-race model, however, comes at a price. The model does not specify the underlying processes that produce behavior in the stop-signal paradigm. Thus, the horse-race model can describe but cannot explain differences in inhibition performance between individuals, populations or experimental conditions. Although the horse-race model cannot give direct

insights into the process of stopping (cf. Section 4), it can be used to test hypotheses about the nature of response inhibition if predictions are formulated in terms of the accuracy and the speed of the stop process and in terms of factors that affect these. In this respect, the horse-race model is similar to signal detection theory, a popular and very general model for analyzing decision making processes in the presence of uncertainty (Green & Swets, 1966; MacMillan & Creelman, 2008)

2.3 Independent Horse-Race Model with Constant SSRT

In its most simple form, the independent horse-race model assumes that go RT is a random variable and, conditional on stop-signal delay, SSRT is constant. Although the assumption of constant SSRT is implausible, ignoring variability in SSRT simplifies the derivation of the model.

Panel A in Figure 2 shows a graphical representation of the model. The go RT distribution represents the distribution of the finishing times of the go process. If T_{go} is a random variable representing the finishing times of the go process with continuous probability density function $f_{go}(t)$ for $t \geq 0$, then the mean and variance of the go RT distribution equal:

$$\bar{T}_{go} = \int_0^{\infty} t f_{go}(t) dt \quad (1)$$

and

$$\sigma_{go}^2 = \int_0^{\infty} (t - \bar{T}_{go})^2 f_{go}(t) dt, \quad (2)$$

respectively. The vertical dotted line in Figure 2 represents the unobservable response to the stop signal. On a given stop-signal delay, the white area to the right of the vertical line represents go RTs that are too slow to win the race; the white area under the curve therefore represents the probability of inhibiting the go response— $P_{Inhibit}(t_{SSD})$. The gray area to the left of the vertical line represents go RTs that are fast enough to win the race; the gray area under the curve therefore represents response rate, i.e., the probability of incorrectly responding to the go stimulus— $P_{Respond}(t_{SSD})$.

Panel B in Figure 2 illustrates how performance in the stop-signal paradigm is determined by the relative finishing times of the go and the stop process. The model assumes that the go response is successfully inhibited if $T_{go} > (t_{stop} +$

t_{SSD}), where t_{stop} and t_{SSD} are constants representing SSRT and stop-signal delay, respectively. Stop-signal trials resulting in successful inhibitions are called *signal-inhibit trials*. In contrast, the go response is incorrectly emitted if $T_{go} < (t_{stop} + t_{SSD})$. Stop-signal trials resulting in erroneous go responses are called *signal-respond trials*, and the corresponding RTs are called *signal-respond RTs*.

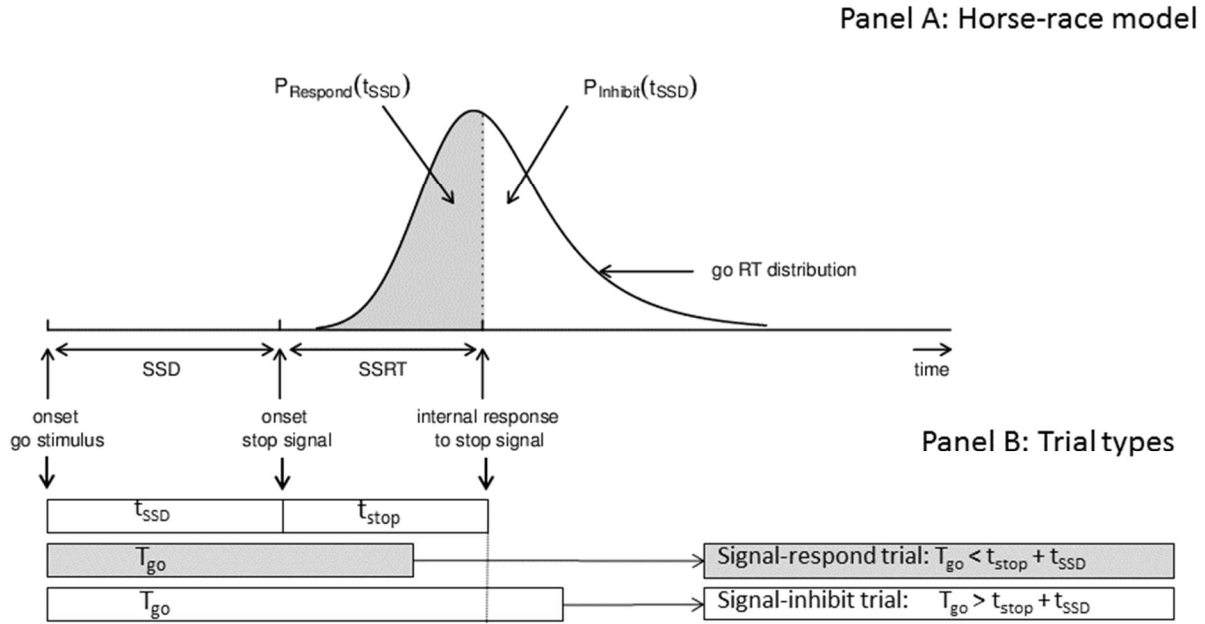


Figure 2. Graphical representation of the independent horse-race model with constant stop-signal reaction time. Panel A shows that response rate (i.e., $P_{Respond}(t_{SSD})$) and the probability of inhibition (i.e., $P_{Inhibit}(t_{SSD})$) are determined by the stop-signal delay (SSD), the stop-signal reaction time (SSRT), and the go RT distribution. Panel B shows that the go response is incorrectly emitted if $T_{go} < (t_{stop} + t_{SSD})$, resulting in a signal-respond trial. In contrast, the go response is successfully inhibited if $T_{go} > (t_{stop} + t_{SSD})$, resulting in a signal-inhibit trial.

The model predicts that the probability of responding on a given stop-signal delay is given by:

$$P_{Respond}(t_{SSD}) = \int_0^{t_{stop}+t_{SSD}} f_{go}(t) dt. \quad (3)$$

The mean of the signal-respond RTs is given by:

$$\bar{T}_{SR}(t_{SSD}) = \frac{1}{P_{Respond}(t_{SSD})} \int_0^{t_{stop}+t_{SSD}} t f_{go}(t) dt. \quad (4)$$

It follows from Equation 4 that mean signal-respond RT is necessarily faster than mean go RT. The model also predicts that mean signal-respond RT increases with increasing stop-signal delay and approaches mean go RT in the limit. The relationship between mean signal-respond RT and mean go RT is also evident from Panel A in Figure 2, where the gray area represents the signal-respond RT distribution. The mean of the signal-respond RTs is necessarily faster than the mean of the go RTs because mean signal-respond RT only represents the mean of those responses that were fast enough to finish before the stop signal (i.e., its calculation does not include the slow tail of the go RT distribution), whereas mean go RT represents the mean of all go responses. With increasing stop-signal delay, the stop response cuts off more of the go RT distribution (i.e., the vertical line shifts to the right), resulting in an increase in the gray area and therefore an increase in mean signal-respond RT (Logan & Cowan, 1984).

2.3.1 Inhibition Functions

According to the independent horse-race model, differences in inhibition performance can be entirely accounted for by the interplay between stop-signal delay, SSRT, and the location and variability of the go RT distribution. The interplay between these factors is often depicted using *inhibition functions*, functions that describe the relationship between stop-signal delay and response rate. These functions are important theoretically because they reflect the outcome of the race between the go process and the stop process (Logan and Cowan, 1984). They are important empirically because they reflect the ability to control responses; they can be used to compare inhibitory control in different groups, tasks, and conditions.

The effect of stop-signal delay on the inhibition function is shown in Panel A of Figure 3. The horse-race model posits that stop-signal delay biases the finishing time of the stop process relative to the go process. As stop-signal delay increases, the stop process is triggered later and later. The stop response therefore cuts off an increasingly larger portion of the go RT distribution, resulting in an increase in response rate. Theoretically, if the stop signal occurs sufficiently early, participants can always inhibit the go response, resulting in a response rate of 0 for short stop-signal delays. If the stop signal occurs sufficiently late, participants can never inhibit the go response, resulting in a response rate of 1 for very long stop-signal delays. As shown in the right panel, between these two extremes, response rate increases monotonically with increasing stop-signal delay.

The effect of increasing go RT on the inhibition function is shown in Panel B of Figure 3. The go RT distribution is shifted to longer RTs (i.e., it is shifted to the right) relative to the go RT distribution in Panel A. For the same stop-signal delay and SSRT, the stop response cuts off a smaller portion of the go RT distribution, resulting in a decrease in response rate. As shown in the right panel, the resulting inhibition function is shifted to the right relative to the inhibition function in Panel A (i.e., dashed line). This prediction of the race model resonates with the empirical finding that participants can slow their go RTs in order to keep response rate constant across the stop-signal delays (Lappin & Eriksen, 1966).

The effect go RT variability on the inhibition function is shown in Panel C of Figure 3. The variance of the go RT distribution is larger relative to the go RT distribution in Panel A. For the same stop-signal delay and SSRT, a smaller portion of the go RT distribution falls between any two consecutive stop-signal delays. As shown in the right panel, the resulting inhibition function is flatter than the inhibition function in Panel A.

The effect of SSRT on the inhibition function is shown in Panel D of Figure 3. SSRT is progressively increased relative to SSRT in Panel A. For the same stop-signal delay and go RT distribution, the stop response cuts off a larger portion of the go RT distribution, resulting in an increase in response rate. As shown in the right panel, the resulting inhibition function is shifted to the left relative to the inhibition function in Panel A.

2.3.2 Diagnosing Deficient Inhibition: Aligning Inhibition Functions

The goal of the quantitative analysis of stop-signal data is to detect differences in inhibition performance between populations, strategies, tasks, or experimental manipulations. Deficiencies in response inhibition may result from a slower or more variable stop process, or from a stop process that is not triggered reliably by the stop signal. All these possibilities impair participant's ability to stop and result in an increased response rate. However, an increase in response rate does not necessarily imply decreased inhibitory ability; for instance, two participants with similar inhibitory ability can differ in response rate as a result of differences in the speed of their go process.

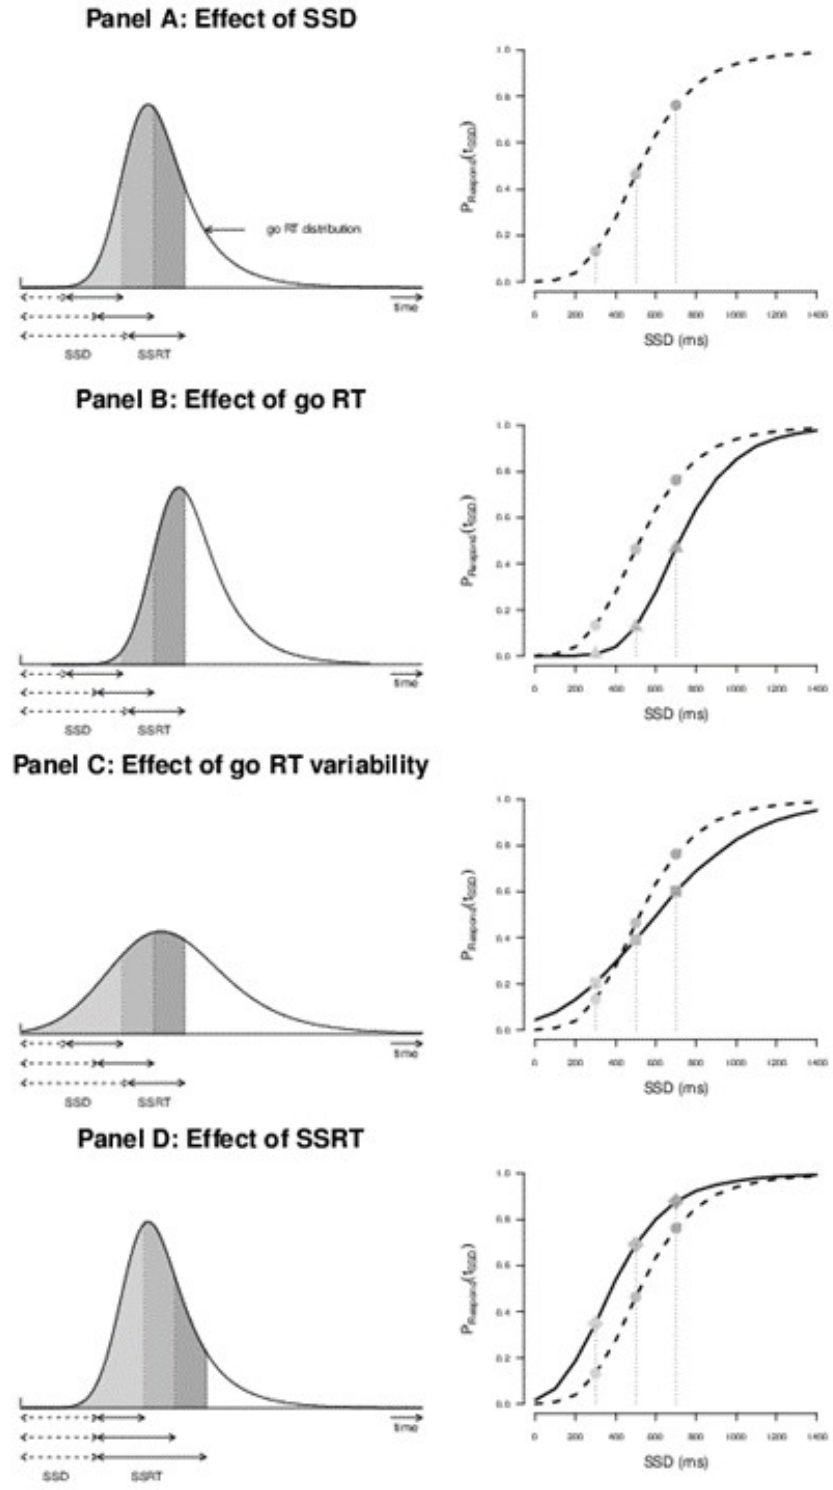


Figure 3. The effect of stop-signal delay (Panel A), go RT (Panel B), go RT variability (Panel C), and stop-signal reaction time (Panel D) on the inhibition function. SSD = stop-signal delay, SSRT= stop-signal reaction time.

When response rate is plotted against stop-signal delay, the horse-race model predicts that an increase in mean go RT shifts the inhibition function to the right (Figure 3, Panel B), an increase in go RT variability (Panel C) decreases the slope of the inhibition function, and an increase in SSRT shifts the inhibition function to the left (Panel D). Therefore inhibitory deficits can be diagnosed by testing whether inhibition functions in the different populations or conditions can be aligned by accounting for differences in mean go RT, differences in go RT variability, and differences in SSRT. Note that the tests are based on visual evaluation of the inhibition functions and not on quantitative assessment of the alignment. Successful alignment indicates that the same inhibitory process applies to all populations or conditions, albeit with differences in go RT and/or differences in SSRT (Logan, 1994; Logan & Cowan, 1984).

First, if inhibition functions can be aligned by plotting response rate against $\bar{T}_{go} - t_{SSD}$, then differences response rate between groups or conditions are only due to differences in mean go RT (e.g., Logan, Cowan, & Davis, 1984; Schachar & Logan, 1990). Note that the same reasoning does not apply to go RT variability; the horse-race model does not predict that accounting for go RT variability by plotting response rate against $(\bar{T}_{go} - t_{SSD})/\sigma_{go}$ should bring the inhibition functions into alignment (e.g., Logan et al., 1984). Second, if inhibition functions can be aligned by plotting response rate against $(\bar{T}_{go} - t_{SSD} - t_{stop})/\sigma_{go}$ (the so-called ZRFT transformation), then differences are due to differences in go performance as well as differences in SSRT (e.g., Logan & Cowan, 1984; Logan et al., 1984; Schachar & Logan, 1990; van der Schoot, Licht, Horsley, & Sergeant, 2000). Thus, differences in response rate only indicate differences in response inhibition ability if accounting for SSRT is necessary to bring the inhibition functions into alignment.

If inhibition functions cannot be aligned by these transformations, the independent horse-race model with constant SSRT cannot account for the data of one or more populations or conditions (Logan & Cowan, 1984). Misalignment is often manifested in differences in the slope of the transformed inhibition functions, and may indicate differences in the variability of the stop process or differences in the ability to trigger the inhibition mechanism (Badcock, Michie, Johnson, & Combrinck, 2002; Schachar & Logan, 1990; Tannock et al., 1995). Theoretically, estimates of SSRT variability from the complete horse-race model (see Section 3.2 and 3.3) and estimates of the probability of trigger failures (see Section 7.1.1) may be used to disentangle the effects of SSRT variability and triggering deficiencies on the slope of ZRFT transformed inhibition functions. Band, van der Molen, & Logan (2003) argued, however, that differences in ZRFT transformed inhibition functions cannot be uniquely attributed to differences in the variability of the stop process or differences in trigger failures because the ZRFT transformation fails to account sufficiently for go RT variability. Therefore,

differences in inhibition functions should be interpreted carefully because it is not always entirely clear what factors are causing the misalignment.

2.4 The Complete Independent Horse-Race Model

The complete independent horse-race model treats go RT, SSRT, and the time required for ballistic processes as independent random variables. For the formal derivation of the complete horse-race model the reader is referred to Logan and Cowan (1984). Here we reiterate their main results without accounting for the ballistic component, and set the stage for introducing approaches to SSRT estimation that do not rely on the over-simplified assumption of constant SSRT.




Figure 4. Graphical representation of the complete horse-race model. SSD = stop-signal delay; SSRT = stop-signal reaction time.

The complete horse-race model assumes that both go RT and SSRT are independent random variables. As shown in Figure 4, the underlying horse-race idea remains the same, but SSRT—just like go RT—can now take on a different value on every stop-signal trial. The model posits that the go response is successfully inhibited (resulting in a signal-inhibit trial) if $T_{go} > (T_{stop} + t_{SSD})$, where T_{go} and T_{stop} are independent random variables representing the finishing

time of the go and the stop process, respectively, and t_{SSD} is a constant representing stop-signal delay. In contrast, the go response is incorrectly emitted (resulting in a signal-respond trial) if $T_{go} < (T_{stop} + t_{SSD})$.

The model predicts that the probability of responding on a given stop-signal delay is given by:

$$P_{Respond}(t_{SSD}) = \int_0^{\infty} f_{go}(t) (1 - F_{stop}(t - t_{SSD})) dt, \quad (5)$$

where $F_{stop}(t - t_{SSD})$ is the cumulative distribution function of the finishing times of the stop process at t_{SSD} . It follows from Equation 5 that increasing stop-signal delay increases the probability of responding by decreasing $F_{stop}(t - t_{SSD})$. The distribution of signal-respond RTs on a given stop-signal delay is given by:

$$f_{SR}(t|t_{SSD}) = f_{go}(t) \frac{1 - F_{stop}(t - t_{SSD})}{P_{Respond}(t_{SSD})}. \quad (6)$$

The complete horse-race model predicts that the signal-respond RT distribution and the go RT distribution share a common lower bound. At higher quantiles, however, the cumulative distribution functions of the two distributions diverge; the shorter the stop-signal delay, the steeper the rise of the cumulative distribution function of the signal-respond RTs. The common lower bound also implies that mean signal-respond RT is shorter than mean go RT (Colonus, Ozyurt, & Arndt, 2001; Osman et al., 1986).

2.4.1 Inhibition Functions

According to the complete horse-race model, varying stop-signal delay in Equation 5 will produce the inhibition function. Similar to the horse-race model with constant SSRT, the complete model predicts that increasing mean go RT decreases the probability that the go process wins the race and results in a rightward shift in the inhibition function. In contrast, increasing mean SSRT decreases the probability that the stop process wins the race and results in a leftward shift in the inhibition function. Increasing go RT or SSRT variability influences the slope of the inhibition function (Logan & Cowan, 1984).

Logan and Cowan (1984) showed that treating the inhibition function as a cumulative distribution allows one to express its mean and variance in terms of the mean and variance of the go RTs and SSRTs. In particular, the mean of the inhibition function equals the difference between mean go RT and mean SSRT:

$$\bar{T}_{SSD} = \bar{T}_{go} - \bar{T}_{stop}. \quad (7)$$

The variance of the inhibition function equals the sum of the variances of the go RTs and SSRTs:

$$\sigma_{SSD}^2 = \sigma_{go}^2 + \sigma_{stop}^2. \quad (8)$$

As we show in Section 3.2, Equation 7 suggests an easy to use method to estimate mean SSRT that does not rely on the unlikely assumption of constant stopping latencies. The complete horse-race model is not limited to estimating the central tendency of the finishing time distribution of the stop process; the model enables the estimation of limits on the moments of the distribution of the stop process and the ballistic component. However, non-parametric estimation of moments of SSRT distributions higher than the first degree requires data quality that is often unavailable in typical stop-signal studies (Logan, 1994; Matzke, Dolan, Logan, Brown, & Wagenmakers, 2013)

2.5 Independence Assumptions

In order to simplify the derivation of the horse-race model, Logan and Cowan (1984) assumed that the go process and the stop process are independent of one another. The independence assumption allows one to treat the go RT distribution on go trials (i.e., trials without stop signal) as the underlying distribution of go RTs on stop-signal trials. The horse-race model relies on two types of independence: stochastic independence and context independence. According to the stochastic independence assumption, on a given trial, the finishing time of the go process is independent of the finishing time of the stop process: for all t_{go} and t_{stop} ,

$$P(T_{go} < t_{go} \cap T_{stop} < t_{stop}) = P(T_{go} < t_{go}) \times P(T_{stop} < t_{stop}). \quad (9)$$

According to the context independence (or signal independence) assumption, the distribution of the finishing times of the go process is the same on go trials and stop-signal trials: for all t_{go} and t_{SSD} ,

$$P(T_{go} < t_{go}) = P(T_{go} < t_{go} | t_{SSD}). \quad (10)$$

Importantly, the horse-race model does not assume functional independence between the go and the stop process. Functional independence means that factors that influence the finishing time distribution of the go process do not influence the finishing time distribution of the stop process, and vice versa. In fact, several neuropsychological and behavioral studies have shown that the go and the stop process are not functionally independent, for example, when the go task requires response selection (Logan et al., 1984; Szmalec, Demanet, Vandierendonck, & Verbruggen, 2009) or Stroop-like interference control (Chambers et al., 2007; Kramer, Humphrey, Larish, Logan, & Strayer, 1994; Ridderinkhof, Band, & Logan, 1999; Verbruggen, Liefoghe, & Vandierendonck, 2004, 2006). In contrast, other studies have provided evidence for the functional independence of the go and the stop process, for instance, in situations where the primary task involves interference due to task switching or dual-task interference due to multiple response alternatives (Logan et al., 2014; Verbruggen, Liefoghe, Szmalec, & Vandierendonck, 2005).

3. Stop-Signal Reaction Times

The stop-signal paradigm owes its popularity to the underlying horse-race model that enables researchers to estimate the latency of the stop process. SSRTs play a pivotal role in diagnosing deficient response inhibition in clinical populations and in assessing participants' stopping ability across different tasks and experimental conditions. Various methods are available to estimate SSRTs. The most popular methods focus exclusively on obtaining summary measures of the latency of stopping (Section 3.1), but the complete horse-race model also allows for the estimation of SSRT variability (Section 3.2). More recent methods provide researchers with the possibility to estimate the entire distribution of SSRTs (Section 3.3), to estimate the parameters of the underlying stop (and go) process (Section 4), and to quantify the relative contribution of trigger failures to stop-signal performance (Section 7.1.1).

3.1. Estimating Summary Measures of SSRT

Various methods are available to estimate summary measures, such as the mean, of the latency of the stop response. The methods differ in whether they treat SSRT as a constant or as a random variable. Which estimation method is most suitable also depends on how stop-signal delay is set. There are two procedures for setting stop-signal delay: (1) using some number of fixed stop-signal delays (i.e.,

the fixed-SSDs procedure) or (2) adjusting stop-signal delays dynamically (i.e., the tracking procedure). The most common tracking procedure involves adjusting stop-signal delay after every trial (i.e., the one-up/one down procedure; see Logan, Schachar, & Tannock, 1997; Verbruggen & Logan, 2009a; Verbruggen et al., 2013): At the beginning of the experiment, stop-signal delay is set to a specific value (e.g., 250 ms) and is then constantly adjusted after stop-signal trials, depending on the outcome of the race. When inhibition is successful, stop-signal delay increases (e.g., by 50 ms); when inhibition is unsuccessful, stop-signal delay decreases (e.g., by 50 ms). This one-up/one-down tracking procedure typically results in $P_{Respond} \approx 0.50$, which means that the race between the stop process and the go process is tied.

3.1.1 Fixed Stop-Signal Delays

The *integration method* is the most popular method when fixed stop-signal delays are used (Logan & Cowan, 1984). The integration method assumes that SSRT is constant and allows for the estimation of SSRT for each stop-signal delay separately. For any given stop-signal delay, the integration method involves finding the value of t_{stop} in the upper limit of the integral in Equation 3 for which the area of the go RT distribution equals $P_{Respond}(t_{SSD})$. In practice, go RTs are rank ordered and the n^{th} go RT is selected, where n is the number of go RTs multiplied by $P_{Respond}(t_{SSD})$. Stop-signal delay is then subtracted to arrive at an estimate of SSRT.

SSRTs estimated with the integration method decrease with increasing stop-signal delay (Logan & Burkell, 1986; Logan & Cowan, 1984). Estimates from different stop-signal delays are therefore averaged to arrive at a single SSRT estimate for each participant. Note that the decrease in estimated SSRT as a function of stop-signal delay is not necessarily at odds with the independence assumption, but can be explained by variability in SSRT. Suppose that SSRTs have a constant mean and nonzero variance. At short stop-signal delays, a large portion of the SSRT distribution will produce successful response inhibition; estimated SSRT therefore closely approximates the mean of the entire SSRT distribution. At long stop-signal delays, only a small portion of the SSRT distribution will produce successful inhibition; estimated SSRT is therefore lower than the mean of the entire SSRT distribution (de Jong et al., 1990; Logan & Burkell, 1986; Logan & Cowan, 1984).

Contrary to the integration method, the *mean method* assumes that SSRT is a random variable. As shown in Equation 7, mean SSRT can be computed by subtracting the mean of the inhibition function from mean go RT (Logan & Cowan, 1984). In the unlikely scenario that the observed inhibition function

ranges from 0 to 1, the mean of the inhibition function can be computed using the values of the $i, i = 2, \dots, n$, stop-signal delays and the corresponding response rates:

$$\bar{T}_{SSD} = \sum_{i=2}^n t_{SSD_i} (P_{Respond}(t_{SSD_i}) - P_{Respond}(t_{SSD_{i-1}})). \quad (11)$$

In case of truncated inhibition functions, the right side of Equation 11 must be divided by $(P_{Respond_{max}} - P_{Respond_{min}})$. However, truncated inhibition functions lose information about the tails of the distribution, which may affect estimates of the mean, particularly when the distribution is skewed.

If the inhibition function is symmetrical, the mean of the inhibition function in Equation 7 may be replaced by the median (Logan & Cowan, 1984). The use of the median is motivated by its ease of computation: The median of the inhibition function is the stop-signal delay where $P_{Respond} = 0.50$. In principle, two stop-signal delays are sufficient to estimate the median of the inhibition function, one with $P_{Respond} < 0.50$ and one with $P_{Respond} > 0.50$. The median can be then obtained by interpolation. If one is willing to assume a parametric form for the inhibition function, the median may be also obtained by fitting a Weibull or logistic function to the observed inhibition function (Chambers et al., 2006; Hanes & Schall, 1995). Note that the Weibull function is not always symmetric, in which case the median cannot replace the mean. A related method entails subtracting the median of the inhibition function from the median of the go RTs. This method is not justified by the mathematics of the race model unless the mean equals the median. As opposed to the integration method, the mean and the median methods do not provide SSRT estimates for each stop-signal delay separately.

The integration method and the mean method both produce reliable SSRT estimates in combination with fixed stop-signal delays, provided that the mean of the inhibition function (Equation 11) is estimated accurately. The use of fixed stop-signal delays, however, requires a relatively large number of observations. For instance, Band et al. (2003) advised researchers to present participants with at least 900 go trials and 60 stop-signal trials on five different stop-signal delays to obtain reliable estimates using the integration method.

3.1.2 Tracking Procedure

The mean method is the most popular method for estimating SSRTs when the tracking procedure is used to set stop-signal delays (Logan & Cowan, 1984; Logan et al., 1997). When tracking results in an overall $P_{Respond}$ of 0.50, the mean of the inhibition function is given by the mean of the stop-signal delays, provided

that the inhibition function is symmetrical. Once the mean of the inhibition function is computed, mean SSRT can be obtained using Equation 7. Due to its simplicity, the mean method has become the dominant method for estimating SSRTs (Verbruggen et al., 2013) and has been implemented in the popular STOP-IT software (Verbruggen, Logan, & Stevens, 2008).

The integration method in combination with tracking entails selecting the n^{th} go RT, where n equals the number of RTs in the go RT distribution multiplied by the overall P_{Respond} . SSRT is then obtained by subtracting mean stop-signal delay from the n^{th} go RT (e.g., Ridderinkhof et al., 1999; Verbruggen et al., 2004; Verbruggen, Stevens et al., 2014). The median method entails subtracting mean stop-signal delay from the median of the go RTs (e.g., Aron & Poldrack, 2006), however there is no justification for the median method in the race model. The race model makes predictions about mean RT and the mean of the inhibition function (Logan & Cowan, 1984). The relationship in Equation 7 does not hold for medians.

Methods relying on tracking require fewer observations for accurate and reliable SSRT estimation than methods that use fixed stop-signal delays (Band et al., 2003; Congdon et al., 2012; Williams et al., 1999). Researchers are recommended to present participants with approximately 120-150 go trials and 40-50 stop-signal trials in combination with the tracking procedure (Verbruggen & Logan, 2009a). Verbruggen et al. (2013) showed, however, that the mean method overestimates SSRTs when go RTs are right-skewed or when participants gradually slow their responses over the course of the experiment. The integration method is less sensitive to the skewness of the go RT distribution, but it underestimates SSRTs in the presence of response slowing. The bias as a result of response slowing disappears when the integration method is applied to smaller blocks of trials as opposed to the entire experiment. Verbruggen and colleagues therefore recommended that researchers use the block-wise integration method to estimate SSRTs in combination with the tracking procedure.

3.2 Estimating SSRT Variability

Two populations or experimental groups may not only differ in mean SSRT, but may also differ in the variability of the latency of the stop response. Logan and Cowan's (1984) treatment of the inhibition function as a cumulative distribution function provides a method for estimating SSRT variability. They observed that in symmetrical distributions the variance is proportional to the slope of the cumulative distribution function at the median. For instance, if we assume a normal distribution, the slope of the inhibition function at the median is given by:

$$B_{0.5} = \frac{1}{\sigma_{SSD}\sqrt{2\pi}}. \quad (12)$$

SSRT variability can be obtained by solving Equation 12 for σ_{SSD} and substituting the solution in Equation 8:

$$\sigma_{stop}^2 = \left(\frac{1}{B_{0.5}\sqrt{2\pi}} \right)^2 - \sigma_{go}^2. \quad (13)$$

Note that the computation of SSRT variability using Equations 12 and 13 assumes a particular parametric form for the inhibition function. This approach is therefore less general than Logan and Cowan's (1984) non-parametric method for deriving the limit on the second moment of the SSRT distribution (see Section 2.4.1). Moreover, reliability studies have shown that the parametric method overestimates the true variability in stopping latencies when inhibition functions are not symmetrical (Band et al., 2003).

3.3 Estimating SSRT Distributions

It is well-known in the response time modeling literature that relying on measures of central tendency, such as the mean, may miss important features of the data (e.g., Heathcote, Popiel, & Mewhort, 1991; Matzke & Wagenmakers, 2009). Likewise, using only summary measures of SSRT may mask crucial aspects of stop-signal data and may lead to erroneous conclusions about response inhibition. For instance, two clinical populations may have the same mean SSRT, but their SSRT distributions may follow markedly different shapes. The development of methods that enable researchers to estimate the entire distribution of SSRTs has been an important advance in the stop-signal literature.

3.3.1 Nonparametric Estimation

Colonus (1990) and de Jong et al. (1990) developed a general approach for estimating the entire distribution of SSRTs. They showed that the survival distribution of SSRTs on a given stop-signal delay is given by (see also Equation 6):

$$1 - F_{stop}(t - t_{SSD}) = P_{Respond}(t_{SSD}) \frac{f_{SR}(t|t_{SSD})}{f_{go}(t)}. \quad (14)$$

In line with the generality of the horse-race model, this method does not hinge on the specific parametric form assumed for the finishing times of the go and the stop

process; all that is required are non-parametric density estimates for $f_{go}(t)$ and $f_{SR}(t|t_{SSD})$. Once the survival distribution of SSRTs is obtained, the quantiles of the SSRT distribution can be easily derived.

The non-parametric formulation, however, comes at a price: The Colonus – de Jong method requires an unrealistically large number of observations to accurately capture the tail of the SSRT distribution (Band et al., 2003; Logan, 1994; Matzke et al., 2013a). As a result, the method has been never used with empirical data.

3.3.2 Parametric Estimation

Process models provide parametric ways of estimating SSRT distributions, which will be discussed below in Section 4. Matzke et al. (2013a) proposed a purely descriptive parametric method that allows researchers to estimate the entire distribution of SSRTs. By assuming a specific parametric form for the go RTs and SSRTs, this approach can provide accurate estimates of SSRT distributions even with relatively few observations. According to the parametric approach, the likelihood on the $r = 1, \dots, R$, signal-respond trials is given by:

$$L_{SR}(\theta_{go}, \theta_{stop} | t_r, t_{ssd}) = \prod_{r=1}^R f_{go}(t_r | \theta_{go}) (1 - F_{stop}(t_r - t_{ssd} | \theta_{stop})), \quad (15)$$

where $f_{go}(t_r | \theta_{go})$ is the probability density function of the finishing time distribution of the go process with parameters θ_{go} and $F_{stop}(t_r - t_{ssd} | \theta_{stop})$ is the cumulative distribution function of the finishing time distribution of the stop process at t_{ssd} with parameters θ_{stop} . The likelihood on the $i = 1, \dots, I$, signal-inhibit trials is given by:

$$L_I(\theta_{go}, \theta_{stop} | t_i, t_{ssd}) = \prod_{i=1}^I \int_0^{\infty} (1 - F_{go}(t_i | \theta_{go})) f_{stop}(t_i - t_{ssd} | \theta_{stop}) dt_i, \quad (16)$$

where $F_{go}(t_i | \theta_{go})$ is the cumulative distribution function of the finishing time distribution of the go process with parameters θ_{go} and $f_{stop}(t_i - t_{ssd} | \theta_{stop})$ is the probability density function of the finishing time distribution of the stop process at t_{ssd} with parameters θ_{stop} . Note that the likelihood on signal-inhibit trials requires integrating over t_i because RTs on signal inhibit-trials—the SSRTs—are by definition unobserved.

Matzke et al.'s (2013a) parametric approach relies on the ex-Gaussian distribution to quantify the shape of the go RT and SSRT distribution (e.g.,

Heathcote et al., 1991; Matzke & Wagenmakers, 2009). The ex-Gaussian distribution is a three-parameter convolution of a Gaussian and an exponential distribution: the μ and σ parameters quantify the mean and the standard deviation of the Gaussian component and reflect the leading edge and mode of the distribution; τ quantifies the mean of the exponential component and reflects the slow tail of the distribution. The model postulates six ex-Gaussian parameters for each participant: three parameters for the go RT distribution, $\theta_{go} = [\mu_{go}, \sigma_{go}, \tau_{go}]$, and three parameters for the SSRT distribution, $\theta_{stop} = [\mu_{stop}, \sigma_{stop}, \tau_{stop}]$. Mean go RT is given by $\mu_{go} + \tau_{go}$ and mean SSRT is given by $\mu_{stop} + \tau_{stop}$. Note that the ex-Gaussian distribution may be substituted with other (shifted) RT distributions, such as the Wald, the Weibull, or the lognormal distribution (e.g., Heathcote, 2004; Heathcote, Brown, & Cousineau, 2004). The model does not interpret the ex-Gaussian distribution as a two stage model, as the convolution might suggest, nor does the model interpret μ , σ , and τ as parameters of the underlying processes (Matzke & Wagenmakers, 2009). The model uses the ex-Gaussian distribution because it is easy to work with mathematically and computationally.

Parameter estimation may proceed by means of standard maximum likelihood estimation (e.g., Myung, 2003, Van Zandt, 2000). However, as the parametric approach was intended to handle individual as well as hierarchical data structures, Matzke et al. (2013a) relied on Bayesian parameter estimation instead (e.g., Lee & Wagenmakers, 2013). In the hierarchical approach, rather than estimating parameters separately for each participant, the participant-level parameters are modeled using truncated normal population-level distributions. The population-level distributions act as priors that adjust—shrink—poorly estimated extreme parameter values to more moderate ones. As a result, the hierarchical approach can provide more accurate and less variable estimates than individual estimation, especially if only scarce participant-level data are available (e.g., Farrell & Ludwig, 2008; Gelman & Hill, 2007; Rouder, Sun, Speckman, Lu, & Zhou, 2003). The posterior distribution of the model parameters can be approximated using Markov chain Monte Carlo sampling (e.g., Gilks, Richardson, & Spiegelhalter, 1996), which has been implemented in the BEESTS software (Matzke, Love, Wiecki, Brown, Logan, & Wagenmakers, 2013).

Regardless of the type of stop-signal delay setting, the Bayesian parametric approach requires relatively few observations per participant to produce reliable estimates of SSRT distributions. The individual approach provides accurate and precise parameter estimates with approximately 250 stop-signal trials. The hierarchical approach requires a sample size of approximately 25 participants, each performing as few as 100 stop-signal trials (Matzke et al., 2013a).

Chevalier et al. (2014) used the Bayesian parametric approach to examine the effects of practice on children's stop-signal performance. They found that

practice differentially affected the leading edge and the slow tail of the SSRT distribution: Practice decreased the μ_{stop} parameter, whereas it increased the τ_{stop} parameter. Colzato, Jongkees, Sellaro, van den Wildenberg and Hommel (2014) used the Bayesian parametric approach to show that the administration of tyrosine (i.e., a precursor of dopamine) selectively affects the μ_{stop} parameter of the SSRT distribution, resulting in a decrease in mean SSRT, but no change in the shape of the SSRT distribution.

4. Process Models of Response Inhibition

The independent horse-race model, including its parametric variants discussed so far, are purely descriptive; they enable researchers to quantify the latency of the unobservable stop response, but they do not specify the processes that give rise to the finishing time distribution of the go and the stop process. In order to explain how stopping occurs one has to rely on *process models* of response inhibition. Process models give direct insights into the mechanisms that implement going and stopping and explain the effects of experimental manipulations on stop-signal performance.

In this section we outline two classes of process models of response inhibition. The first class of models—the Hanes-Carpenter model and the race diffusion model—focuses on describing the properties of the go and the stop process in order to explain how the finishing time distributions arise. The second class of models—the interactive race model and blocked input models—does not only describe the nature of the processes that race against each other, but also attempts to explain *how* responses are stopped.

4.1 Describing the Properties of the Go and Stop Process

The Hanes-Carpenter model and the race diffusion model conceptualize response inhibition as a race between a set of evidence accumulators. The two models, however, differ in the mathematical formulation of the evidence accumulation process and the type of go task that they can be applied to. The Hanes-Carpenter model was developed to describe how saccades are inhibited and

applies exclusively to simple RT go tasks (i.e., go task with a single go response). It does not account for choice errors. The race diffusion model was developed to describe stopping of all kinds of responses in simple and choice RT tasks, accounting for accuracy as well as RT. Note that most stop-signal studies have used choice RT tasks (for reviews, see Logan, 1994; Verbruggen & Logan, 2008b). Both models can be considered as special cases of the Logan and Cowan (1984) independent horse-race model with specific parameterizations of the go and stop processes.

4.1.1 Hanes-Carpenter Model of Saccadic Inhibition

The Hanes-Carpenter model (Hanes & Carpenter, 1999; see also Hanes & Schall, 1995; Hanes et al., 1998) aims to explain the processes that are involved in a saccadic version of the stop-signal task. Participants fixate their gaze at a central fixation point, and when the fixation point disappears, they are required to make a saccade to a visual target that appears in one of two positions in the periphery. Occasionally, this go task is interrupted by a visual stop signal (e.g., reappearance of the fixation point) that instructs participants to withhold their eye movement on that trial. Performance is much like in stop-signal tasks with other responses, except that go RTs and SSRTs are shorter and participants never make choice errors.

The Hanes-Carpenter model is based on the Linear Approach to Threshold with Ergodic Rate (LATER; Carpenter, 1981; Carpenter & Williams, 1995) approach, a model that has been successfully used to describe the processes involved in the initiation of saccades in humans. LATER assumes that saccade initiation can be conceptualized as a signal that rises linearly towards a fixed threshold; when the signal reaches the threshold, the saccade is initiated. The rate of rise is assumed to vary from trial to trial according to a normal distribution. The Hanes-Carpenter model assumes that the inhibition of saccades can be similarly formalized as a rise-to-threshold mechanism, such as the one shown in Figure 5, where the go and the stop process rise linearly towards their respective thresholds. If the go process reaches the threshold first, the saccade is initiated; if the stop process reaches the threshold first, saccade initiation is inhibited. The Hanes-Carpenter model is similar to the Linear Ballistic Accumulator model, which allows multiple choices and variation in starting point to account for errors (Brown & Heathcote, 2008).

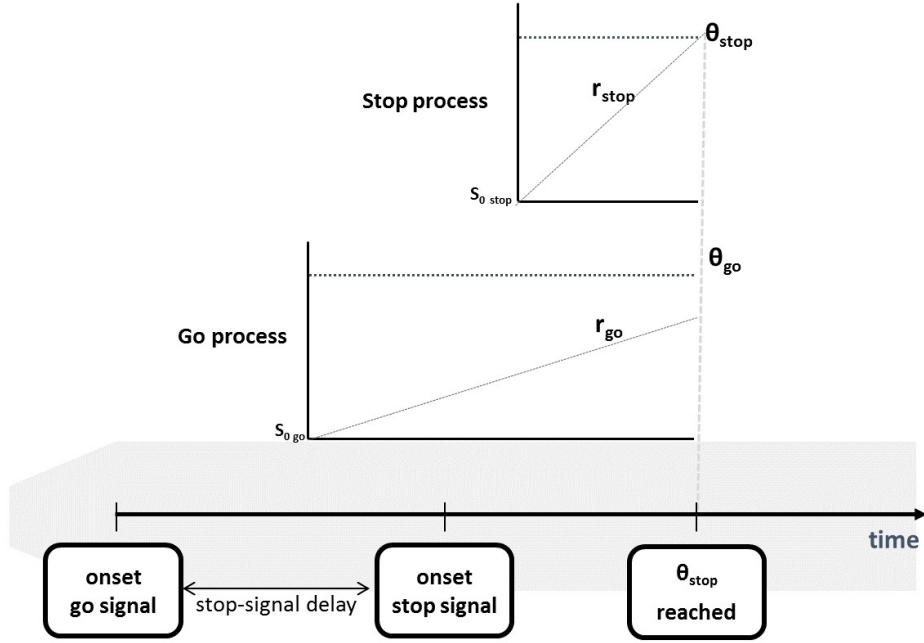


Figure 5. The Hanes-Carpenter model. The model assumes that the go process raises in a linear fashion with rate r_{go} from a predefined starting point $s_{0 go}$ to a fixed threshold θ_{go} . Similarly, the stop process raises in a linear fashion with rate r_{stop} from a starting point $s_{0 stop}$ to a fixed threshold θ_{stop} . In the illustration, the stop process reaches the threshold before the go process; saccade initiation is therefore inhibited.

Specifically, the Hanes-Carpenter model assumes that the go process t_{go} raises in a linear fashion with rate r_{go} from a predefined starting point $s_{0 go}$ to a fixed threshold θ_{go} :

$$s_{0 go} + r_{go}t_{go} = \theta_{go}. \quad (17)$$

If r_{go} is assumed to vary from trial to trial according to a normal distribution with mean μ_{go} and standard deviation σ_{go} , the probability density function of the finishing times of the go process is given by:

$$f_{go}(t) = \frac{\theta_{go} - s_{0 go}}{\sigma_{go}\sqrt{2\pi t^2}} \exp \left[-\frac{\left(\frac{\theta_{go} - s_{0 go}}{t} - \mu_{go} \right)^2}{2\sigma_{go}^2} \right]. \quad (18)$$

Similarly, the stop process t_{stop} is assumed to increase linearly with rate r_{stop} from a starting point $s_{0_{stop}}$ to a fixed threshold θ_{stop} , where the rate of rise is normally distributed with mean μ_{stop} and standard deviation σ_{stop} . The probability density function of the finishing times of the stop process is given by substituting the stop parameters in Equation 18. The probability density function of the signal-respond RTs and the survival distribution of SSRTs can be obtained by substituting into Equation 6 and 14, respectively. The model also features two fixed parameters that quantify the constant processing time of the go and the stop signals. The model parameters may be estimated with Monte Carlo simulations (Hanes & Carpenter, 1999) or with maximum likelihood estimation (e.g., Corneil & Elsley, 2005; Kornyló, Dill, Saenz, & Krauzlis, 2003) using analytic expressions for the density functions of the finishing time distributions and the signal-respond RTs (Colonius et al., 2001).

The Hanes-Carpenter model can be used to estimate SSRT. The mean finishing time for the stop process is simply $(\theta_{stop} - s_{0_{stop}})/r_{stop}$. The distribution of finishing times for the stop process can be obtained by substituting r_{stop} , $s_{0_{stop}}$, and θ_{stop} into Equation 18. These statistics describe the “parent” distribution of the stop runner in the race. To generate the distribution of finishing times when the stop process wins the race, the distribution from Equation 18 would have to be substituted into Equation 6.

Hanes and Carpenter (1999) successfully applied the model to the data of four participants and concluded that the process of saccade inhibition can be described with an independent race mechanisms with a linear rise to threshold. Colonius et al. (2001) used the Hanes-Carpenter model to show that saccade inhibition is more efficient in response to auditory stop signals than visual stop signals.

4.1.2 The Race Diffusion Model

The race diffusion model is a specific instantiation of the general independent race model developed by Logan et al. (2014). As the name suggests, the general independent race model is a generalization of the standard independent horse-race model that can account for go and stop performance in (multiple-) choice RT tasks. The model assumes a race between a set of stochastically independent evidence accumulators (Ratcliff & Smith, 2004), one accumulator that corresponds to the stop response and N accumulators that correspond to the N possible responses on the go task. The response and corresponding RT on a given trial is determined by the first accumulator that reaches its threshold. The standard independent horse-race mode is a special case of the general independent race

model with a single accumulator for the go process and another one for the stop process.

The model predicts that the probability of go response i , $i = 1, \dots, N$, is given by the probability that go process i finishes before all other runners in the race:

$$P_{Respond,i}(t_{SSD}) = \int_0^{\infty} f_{go,i}(t) \prod_{j \in N, j \neq i} (1 - F_{go,j}(t)) (1 - F_{stop}(t - t_{SSD})) dt, \quad (19)$$

where $f_{go,i}$ is the probability density function of the finishing times of the i^{th} go accumulator and $F_{stop}(t - t_{SSD})$ is the cumulative distribution function of the finishing times of the stop accumulator at t_{SSD} . On go trials, t_{SSD} is set to ∞ , and $F_{stop}(t - t_{SSD})$ equals 0. The probability of successful inhibition on a given stop-signal delay is given by:

$$P_{Inhibit}(t_{SSD}) = \int_0^{\infty} f_{stop}(t - t_{SSD}) \prod_{i \in N} (1 - F_{go,i}(t)) dt. \quad (20)$$

The joint probability density function of RTs given response i is then given by:

$$f_i(t|t_{SSD}) = \frac{f_{go,i}(t) \prod_{j \in N, j \neq i} (1 - F_{go,j}(t)) (1 - F_{stop}(t - t_{SSD}))}{1 - P_{Inhibit}(t_{SSD})}. \quad (21)$$

On go trials, $F_{stop}(t - t_{SSD})$ and $P_{Inhibit}(t_{SSD})$ both equal 0. On stop-signal trials, $t_{SSD} \ll \infty$, and Equation 21 gives the probability density function of signal-respond RTs for response i . The survival distribution of SSRTs at a given stop-signal delay can be obtained by substituting the probability density function of the go RTs

$$f_{go}(t) = \sum_{i \in N} f_{go,i}(t) \prod_{j \in N, j \neq i} (1 - F_{go,j}(t)) \quad (22)$$

and the probability density function of the signal-respond RTs

$$f_{SR}(t|t_{SSD}) = \frac{\sum_{i \in N} f_{go,i}(t) \prod_{j \in N, j \neq i} (1 - F_{go,j}(t)) (1 - F_{stop}(t - t_{SSD}))}{1 - P_{Inhibit}(t_{SSD})} \quad (23)$$

into Equation 14.

The general independent race model makes general predictions about the interplay between response rate and RT distributions without specifying the properties of the accumulators that give rise to the finishing time distributions. In order to specify the processes that generate the finishing time distributions, Logan et al. (2014) investigated three special independent race models: the diffusion model (Ratcliff, Van Zandt, & McKoon, 1999), the Linear Ballistic Accumulator (Brown & Heathcote, 2008), and the Poisson counter model (van Zandt, Colonius, & Proctor, 2000). All three models assume that each runner in the race is a stochastic accumulator, but make different assumptions about the nature of the evidence accumulation process. All the three models fit the data well, but the race diffusion model did slightly better. Here we follow Logan and colleagues and only consider the race diffusion model in more detail.

As shown in Figure 6, the race diffusion model assumes that the stop accumulator and each of the N go accumulators is a Wiener diffusion process with drift rate ξ , starting point 0, and a single threshold z . The model also assumes a nondecision time ω parameter that quantifies the time required for stimulus encoding and response execution and a drift coefficient that was set to 1. The finishing time distribution of each accumulator is a Wald (i.e., inverse Gaussian) distribution. The probability density function of the finishing time distribution of go accumulator $i, i = 1, \dots, N$, is thus given by:

$$f_i(t) = \frac{z_i}{\sqrt{2\pi t^3}} \exp\left(-\frac{1}{2t}(\xi_i t - z_i)^2\right) \text{ for } t > 0. \quad (24)$$

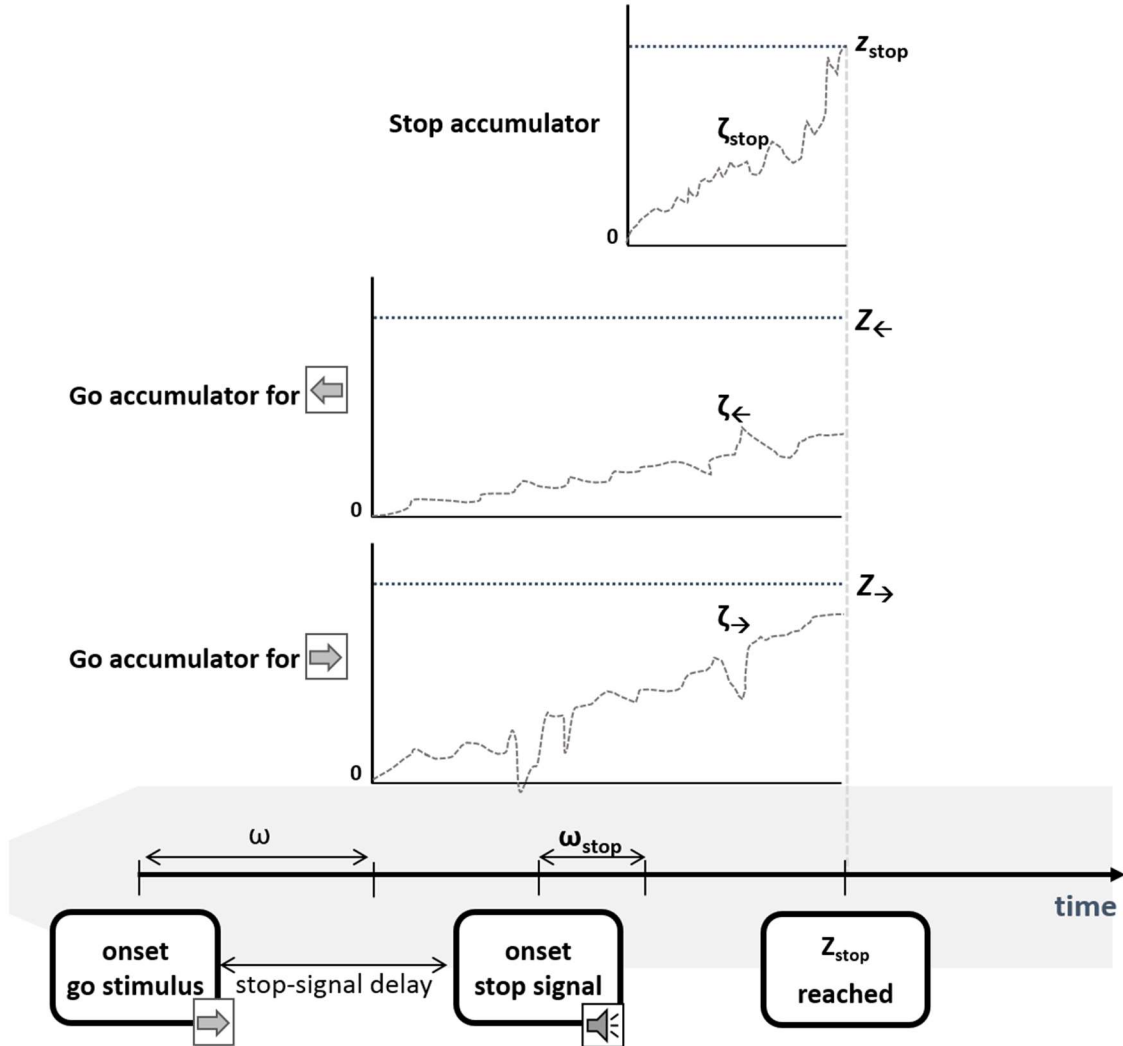


Figure 6. The race diffusion model. In the present illustration, the model assumes a race between a set of stochastically independent evidence accumulators, one accumulator for the stop response and $N=2$ accumulators that correspond to the 2 possible responses on the go task (i.e., left or right-pointing arrow). The model assumes that the accumulators can be described by a Wiener diffusion process with drift rate ξ , starting point 0, threshold z , and nondecision time ω . The stop accumulator reaches threshold before either of the go accumulators; the go response is therefore inhibited.

The probability density function of the finishing time distribution of the stop accumulator with support $t > t_{ssd}$ can be obtained by substituting $(t - t_{ssd})$ for t , and ξ_{stop} and z_{stop} for ξ_i and z_i in Equation 24. The finishing time distribution of the winner of the race is given by the distribution of the minima of the Wald distributions for all the runners.

In order to account for the RTs of fast error responses, Logan et al. (2014) extended the model and allowed the threshold parameter to vary across trials.

Threshold was assumed be a uniform random variable with support $[(z - a), (z + a)]$. In the extended model, the probability density function of the i^{th} go accumulator is given by:

$$g_i(t|\xi_i, z_i, a_i) = \frac{1}{2a_i} [\varphi(\alpha_i) - \varphi(\beta_i) - \xi_i(\Phi(\alpha_i) - \Phi(\beta_i))], \text{ for } \xi_i > 0 \text{ and } a_i > 0, \quad (25)$$

where $\varphi(x)$ and $\Phi(x)$ are the probability density and cumulative distribution function of the standard normal distribution, respectively, and $\alpha = \frac{-(z-a-t\xi)}{\sqrt{t}}$ and $\beta = \frac{(z+a-t\xi)}{\sqrt{t}}$.

Note that for $a = 0$, Equation 25 simplifies to Equation 24. For $\xi = 0$, Equation 25 simplifies to:

$$g_i(t|z_i, a_i) = \frac{1}{2a_i} [\varphi(\alpha_i) - \varphi(\beta_i)]. \quad (26)$$

After substituting Equation 25 and 26 into Equations 19-21, the model parameters can be estimated with maximum likelihood estimation (Van Zandt, 2000) using the correct and error go RT distributions, the signal-respond RT distributions, and the inhibition functions.

The race diffusion model can also be used to estimate mean SSRT and SSRT distributions. The parent SSRT distribution can be obtained by calculating Equation 24 with the best-fitting stop parameters. The distribution of winning SSRTs can be obtained by calculating Equation 25 or 26 with the best fitting stop parameters. Logan et al. (2014) found that SSRTs calculated from the model agreed well with SSRTs calculated from the data with the integration method.

Logan et al. (2014) applied the race diffusion model to investigate capacity limitations in the go and the stop process. To test the hypothesis that the go and stop processes share capacity, Logan and colleagues assumed that the threshold parameter is selectively influenced by strategic factors, whereas the drift rate parameter is selectively influenced by structural factors and can therefore be interpreted as a measure of processing capacity (Ratcliff & Smith, 2004; Ratcliff et al., 1999). Eight different versions of the race diffusion model were fit to the stop-signal data of six participants. Each participant performed three series of stop-signal trials, one with two choice alternatives, one with four choice alternatives, and one with six choice alternatives on the go task. The eight models imposed different combinations of constraints on the drift rate and threshold parameters of the go and the stop process as a function of the number of choice alternatives. In the best fitting model, the drift rate of the go process decreased as the number of choice alternatives increased but the drift rate of the stop process

did not vary with the number of alternatives. The modeling results led Logan and colleagues to conclude that (1) the go process has limited capacity and (2) that the stop process does not share capacity with the go process. These findings support the functional independence of the go and the stop process (see Section 2.5).

4.2 Describing How Responses are Inhibited

The Hanes-Carpenter model and the race diffusion model outlined in the previous section describe the nature of the go and the stop process, but do not specify how responses are stopped. The interactive race model and blocked-input models of saccadic inhibition address this limitation. The interactive race model is a neurally plausible instantiation of the standard independent horse-race model that assumes that responses are stopped by a mechanism that directly inhibits the growth of activation of the go process. In contrast, blocked-input models assume that stopping is not a result of inhibiting the growth of activation in the go process, but rather of blocking the input to the go process, possibly by inhibiting the process that generates drift rates or the process that communicates them to the response processes.

The interactive race model and blocked-input models were developed within the framework of neurally constrained cognitive modeling. Within this framework, evaluation of the descriptive accuracy of competing models is based on the models' ability to simultaneously account for behavioral and neurophysiological data. This approach is useful in situations when competing cognitive models are indistinguishable on grounds of the predictions they make for behavioral data (i.e., model mimicry; Logan, 2004; Myung, Pitt, & Kim, 2005; Townsend, & Ashby, 1983; Wagenmakers, Ratcliff, Gomez, & Iverson, 2004). In particular, the additional constraints provided by requiring models to account for essential features of the neurophysiological data can break model mimicry and can contribute to a more principled choice among models (Boucher, Palmeri, Logan, & Schall, 2007; Logan, Yamaguchi, Schall, & Palmeri, 2015; Hanes & Schall, 1996; Turner, Forstmann, Wagenmakers, Brown, Sederberg, & Steyvers, 2013).

4.2.1 Interactive Race Model of Saccadic Inhibition

The interactive race model (Boucher et al., 2007) is inspired by the apparent contradiction between the results of neurophysiological and behavioral studies of saccade inhibition. On the neurophysiological side, it is well established that saccades are produced by inhibitory interactions between gaze-shifting

neurons that are involved in saccade initiation and gaze-holding neurons that are involved in saccade inhibition (for a review, see Munoz & Schall, 2003). In contrast, on the behavioral and modeling side, there is substantial evidence for the independence of the go and the stop process. In particular, the standard independent horse-race model has been repeatedly shown to provide excellent description of behavior in the stop-signal paradigm in general (see Section 5) and the saccadic stop-signal task in particular (e.g., Hanes & Carpenter, 1999).

To resolve this paradox, Boucher et al. (2007) proposed a simple network shown in Panel A of Figure 7 that consists of a go (or move) and a stop (or fixation) unit that may interact via inhibitory links. The interactive race model conceptualizes the go unit as a stochastic accumulator that gathers evidence to a threshold θ . The saccade is initiated when activation in the go unit reaches threshold. The stop process is formalized as a stochastic evidence accumulator that stops saccade execution by inhibiting the growth of activation of the go unit and preventing it from reaching threshold. Inhibition is thus successful if the stop unit becomes active early enough and strongly enough to suppress the activation of the go unit before saccade initiation.

The model assumes constant rates of rise to threshold with noise terms that are drawn from zero-centered Gaussian distributions. The following differential equations (Usher & McClelland, 2001) describe the change in activation of the go and stop units within time step dt ($\frac{dt}{\tau}$ can be set to 1):

$$da_{move}(t) = \frac{dt}{\tau} (\mu_{move} - k_{move} * a_{move}(t) - \beta_{fix} * a_{fix}(t)) + \sqrt{\frac{dt}{\tau}} \varepsilon_{move} \quad (27)$$

$$da_{fix}(t) = \frac{dt}{\tau} (\mu_{fix} - k_{fix} * a_{fix}(t) - \beta_{move} * a_{move}(t)) + \sqrt{\frac{dt}{\tau}} \varepsilon_{fix}, \quad (28)$$

where μ_{move} and μ_{fix} represent the mean growth rates of the go and the stop units, respectively, and ε_{move} and ε_{fix} are Gaussian noise terms with standard deviation σ_{move} and σ_{fix} that reflect the amount of noise added in each step of the rise. The crucial ingredient of the model is the inhibitory link between the go and the stop unit: the β_{move} parameter reflects the inhibitory influence of the go unit on the stop unit; β_{fix} reflects the inhibitory influence of the stop unit on the go unit. The amount of inhibition is determined by the activation level a_{move} and a_{fix} at time point t . The leakage parameters k ensures that activation does not increase without bounds and can be set to 0. The model also features three parameters that quantify the time needed for stimulus encoding and for the ballistic stage of the go process, some of which were fixed to values derived from physiological measurements.

Model parameters can be estimated with optimizing the fit between observed and predicted data by minimizing a Pearson χ^2 statistic (Ratcliff & Tuerlinckx, 2002).

To assess the model's ability to describe the behavioral data (see also Section 5.2), Boucher et al. (2007) fit the model to behavioral data from two monkeys who performed the saccadic stop-signal task and found good fits to inhibition functions and go RT and signal-respond RT distributions. To assess the model's ability to predict the neurophysiological data, Boucher and colleagues proposed a set of linking propositions (Schall, 2004) that connects the model architecture to underlying physiology. In particular, they linked the go unit to movement-related neurons and the stop unit to fixation-related neurons in frontal eye fields and superior colliculus (Hanes & Schall, 1996; Pouget et al., 2011; Ratcliff, Cherian, & Segraves, 2003). They suggested that the inhibitory connections within the circuit of fixation and movement neurons were sufficient to explain the inhibition of responses.

The interactive race model that best satisfied the simultaneous constraints provided by the neural and behavioral data assumed that the inhibitory effect of the stop process on the go process is delayed and very brief. This result led Boucher et al. (2007) to conclude that response inhibition consists of two stages: during the first stage—the encoding stage—the go and stop process are independent; during the second stage—the interruption stage—the stop process potently inhibits the go process. As the interruption stage is very brief, SSRT estimates from the independent horse-race model are valid measures of the latency of stopping because it mostly reflects the encoding stage of response inhibition.

4.2.2 Blocked-Input Models of Saccadic Inhibition

Logan et al. (2015) proposed a family of alternatives to the interactive race model that provide different explanations of how saccades are stopped in the stop-signal task. Logan and colleagues focused on blocked-input models that postulate that saccades are not stopped by directly inhibiting the growth of activation of the go process, but rather by blocking the input to the go unit, operationalized as setting its drift rate to zero (Logan, 1983; Logan & Cowan, 1984).

The blocked input model conceptualizes the go (move) and the stop (fixation) units (see Panel B of Figure 7) as two stochastic accumulators that race towards their respective threshold θ_{move} and θ_{fix} . The change in activation of the go and stop unit can be described with Equations 27 and 28. According to the model, when the stop unit activation reaches threshold θ_{fix} , it blocks the input to the go unit by setting μ_{move} equal to 0. As a result, activation of the go unit will not reach threshold θ_{move} ; go activation will either hover (if $k_{move} = 0$) or will start to decay (if $k_{move} > 0$).

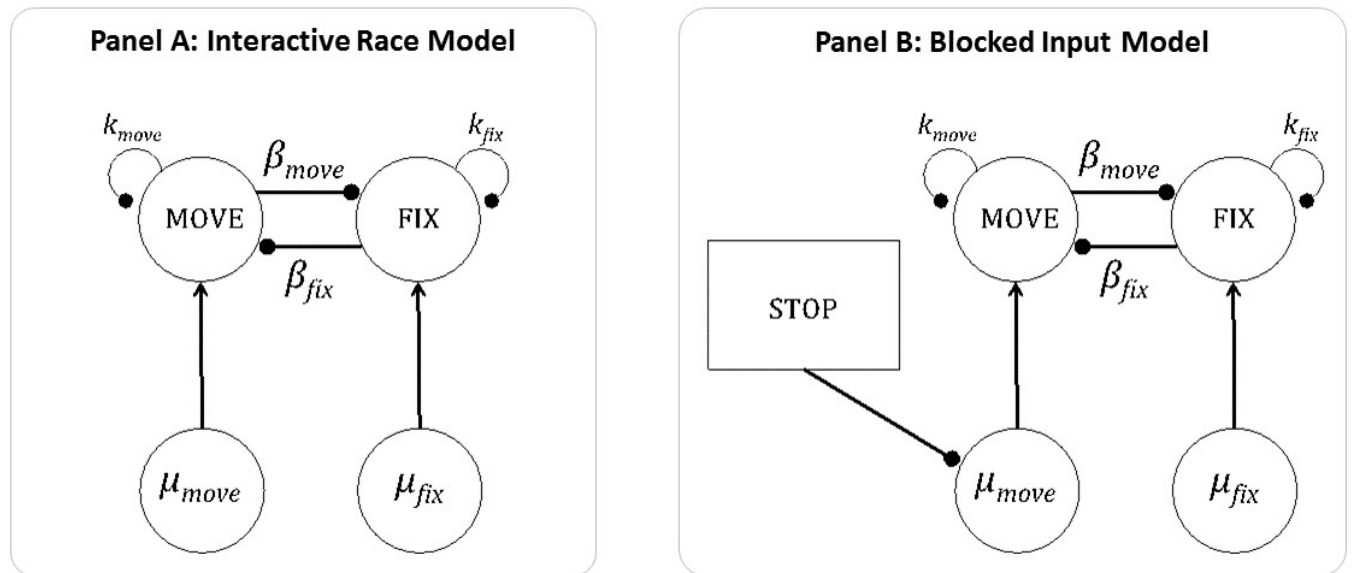


Figure 7. The architecture of the interactive race model and the blocked input model of saccadic inhibition. Panel A: Interactive race model of saccadic inhibition. The go process is identified with movement-related neurons and the stop process is identified with fixation-related neurons in frontal eye fields and superior colliculus. In model fits, β_{fix} is much larger β_{move} . Panel B: Blocked input model of saccadic inhibition. In the first version of the model $\beta_{move} = \beta_{fix} = 0$. In the second version of the model $\beta_{move} > 0$ and $\beta_{fix} > 0$ to account for fixation-related activity at the beginning of a trial. Stopping is accomplished by activating the stop process, which sets μ_{move} to zero. (Adapted from Figure 9 in Logan et al., 2015).

Logan et al. (2015) first considered a blocked input model in which $\beta_{move} = \beta_{fix} = 0$, and found that it fit the behavioral data as well as the interactive race model. The blocked input model provided a better description of the physiological data; the interactive race model predicted a reduction in go activation after the stop signal that was much steeper than observed in the neural

activity. Logan et al. then extended the models back in time to consider activity at the start of the trial when the eyes were fixated. Trials began with the fixation unit fully activated and inhibiting the go process, which had to overcome this inhibition when a go stimulus appeared. These models imposed strong constraints on the stop and go parameters. In particular, β_{fix} and μ_{fix} could not be so large that they inhibit all growth in go activation, or else saccades would never occur. These constraints led to equivalent predictions of physiological data but the blocked input model provided a better account of the behavioral data than the interactive race model. This led Logan et al. to re-evaluate the linking propositions that connected the stop process with fixation neurons in frontal eye fields and superior colliculus. They concluded that fixation neurons were not directly linked to the stop process and instead identified the stop process with a process outside the network that tips the balance in favor of stopping or going.

5. Testing the Goodness-of-Fit of the Horse-Race Model

Conclusions from the model-based analysis of response inhibition data are only warranted if the independent horse-race model indeed provides an adequate description of the data. Non-parametric methods for assessing the goodness-of-fit of the horse-race model focus on evaluating the context independence assumption by analyzing signal-respond RTs. Parametric methods for assessing goodness-of-fit also examine the descriptive accuracy of the chosen parametrization.

5.1 Non-Parametric Methods

Non-parametric methods for assessing the goodness-of-fit of the horse-race model rely on evaluating the context independence assumption. The analyses proceed by comparing the mean and the entire distribution of observed signal-respond RTs to predictions from the independent horse-race model.

First, as shown in Equation 4 and Figure 2, the independent horse-race model predicts that mean signal-respond RT should be faster than mean go RT. As explained in Section 2.4, this prediction should hold regardless whether SSRT is constant or it is a random variable (Colonius et al., 2001; Logan & Cowan, 1984). This prediction has been confirmed in many studies across a range of different populations and experimental manipulations (e.g., de Jong et al., 1990; Hanes &

Schall, 1995; Logan et al. 1984; Osman et al., 1986; van den Wildenberg, & van der Molen, 2004; Verbruggen et al., 2004; Verbruggen, Stevens, et al., 2014).

Second, as discussed in Section 2.3, the independent horse-race model predicts that mean signal-respond RT should increase with increasing stop-signal delay. This prediction can only be evaluated if a large number of stop-signal trials and hence a large number of signal-respond RTs are available on each stop-signal delay; with a small number of stop-signal trials, the estimation of the mean signal-respond RTs will be unstable. The increase in mean signal-respond RT as a function of stop-signal delay has been confirmed in many studies (de Jong et al., 1990; Hanes & Schall, 1995; Logan et al., 1984; Osman et al., 1986). Other studies have, however, reported violations of this prediction especially at short stop-signal delays that typically feature only a small number of signal-respond RTs (e.g., Logan, 1981, Logan et al, 1984).

Third, the independence assumption is often tested by comparing the observed mean signal-respond RTs to the mean signal-respond RTs predicted by the independent horse-race model. Predicted mean signal-respond RTs can be generated for each stop-signal delay by rank-ordering the go RTs and calculating the mean of the n fastest go RTs, where n is computed by multiplying the number of go RTs with $P_{Respond}(t_{SSD})$ (see Section 3.1.1). Several studies have reported only negligible differences (e.g., de Jong et al., 1990; Hanes and Schall, 1995; Logan and Cowan, 1984), whereas others have found substantial discrepancies between observed and predicted mean signal-respond RTs (e.g., Colonius et al., 2001, van den Wildenberg, van der Molen, & Logan, 2002; Verbruggen et al., 2004), especially at short stop-signal delays. However, testing differences between observed and predicted mean signal-respond RTs is not a conclusive test of the independence assumption of the horse-race model. The method of generating predicted signal-respond RTs is based on the unrealistic assumption of constant SSRT. As a result, signal-respond RTs that are longer than $(\bar{T}_{stop} + t_{SSD})$ are excluded from the computation of mean signal-respond RT, which results in a downward bias for the predictions. Moreover, Band et al. (2003) showed that the difference between observed and predicted mean signal-respond RT is not only sensitive to violations of context independence, but is also strongly influenced by SSRT variability; even if context independence holds, increasing SSRT variability increases the difference between observed and predicted mean signal-respond RT. Band and colleagues also showed that the difference between observed and predicted signal-respond RT is not sufficiently sensitive to violations of the stochastic independence assumption of the horse-race model.

Lastly, the independent horse-race model makes specific predictions for the entire distribution of signal-respond RTs. As discussed in Section 2.4, the model predicts that the signal-respond RT distribution and the go RT distribution share a common lower bound, and diverge at higher quantiles. Moreover, the shorter the

stop-signal delay, the steeper the rise of the cumulative distribution function of the signal-respond RTs. Although these predictions have been confirmed by several studies (e.g., Boucher et al., 2007; Camalier et al., 2007; Osman et al., 1986), others have reported violations of the distribution equality test (Colonius et al., 2001).

5.2 Parametric Methods

Parametric methods for assessing the goodness-of-fit of the horse-race model focus on the adequacy of the chosen architecture and the descriptive accuracy of the parametric form assumed for the finishing time distribution of the go and the stop process. Parametric methods proceed by comparing the observed data to data predicted by the model.

Matzke et al. (2013a) used Bayesian posterior predictive simulations (Gelman, Meng, & Stern, 1996) to examine the descriptive accuracy of their ex-Gaussian distributional approach by comparing predictions based on the joint posterior distribution of the model parameters to the observed data. Matzke and colleagues reported that the model provided an adequate description of the inhibition functions and the signal-respond RT distributions of most participants. Logan et al. (2014) confirmed the goodness-of-fit of the race diffusion model by comparing the observed inhibition functions, error rates, and go RT and signal-respond RT distributions to model predictions generated with the maximum likelihood estimates from the best fitting model. Logan et al. (2015, see also Boucher et al. 2007) assessed the descriptive accuracy of the interactive-race and the blocked-input model by comparing the observed inhibition functions and the go RT and signal-respond RT distributions to the ones predicted by the best fitting parameter values obtained by minimizing a Pearson χ^2 statistic. Similarly, Hanes and Carpenter (1999) relied on the comparison between observed and predicted inhibition functions and signal-respond RT distributions to verify the goodness-of-fit of their model. However, Colonius et al. (2001) reported that the horse-race model in general and the Hanes-Carpenter model in particular failed to account for the signal-respond RTs of one of their three participants, suggesting a violation of independence.

5.3 The Independence Assumption in Practice

Stop-signal data from simple stopping tasks are mostly consistent with the independence assumptions of the horse-race model, but more complex selective

stopping tasks have shown consistent violations of independence (e.g., Bissett & Logan, 2014; Verbruggen & Logan, 2015). Logan and Cowan (1984) introduced the independence assumptions to simplify the mathematical derivation of the horse-race model. Violations of the context and stochastic independence assumptions, however, should not be taken lightly as they invalidate calculations based on the race model. Band et al. (2003) showed that violations of stochastic independence may bias SSRT estimates and influence the slope of the ZRFT transformed inhibition function (see also de Jong et al., 1990).

Fortunately, traditional SSRT estimation methods that rely on central stop-signal delays where $P_{Respond}$ approximates 0.50 are relatively unaffected by minor violations of the independence assumptions (Band et al., 2003). Hence the mean method, whether used in combination with fixed stop-signal delays or delays resulting from tracking, is robust to violations of independence (although it may suffer from other problems, as discussed in Section 3.1 and Section 7.1.2). Similarly, the integration method results in reliable SSRT estimates as long as computations are based on the central part of the inhibition function. The midpoint of the inhibition function is automatically obtained with tracking, but can also be approximated with fixed stop-signal delays that fall in the central part of the inhibition function (see e.g., Logan et al., 1984; Logan et al., 2014). Note however that the integration method assumes that SSRT is constant, an assumption that is necessarily at odds with the possibility of a correlated go and stop process. Presently there are no reliability studies available for Matzke et al.'s (2013a) distributional approach and the complex process models of response inhibition discussed in Section 4.

6. Variants of the Stop-Signal Task

So far, we have focused on performance in the stop-signal task in which participants responded to a go stimulus, but withheld their response whenever a stop signal occurred. In this section, we will briefly discuss some variants of the standard stop-signal task.

6.1. Stopping in Stop-Change and Selective Stop Tasks

Two popular variants of the stop-signal task are the stop-change task and the selective stop task. In stop-change tasks, subjects are instructed to stop the originally planned go response and execute an alternative “change” response when a signal occurs (for reviews, see Boecker, Gauggel, & Druke, 2013; Logan & Burkell, 1986; Verbruggen & Logan, 2009a). Experimental, computational, and neuro-imaging work suggests that participants first inhibit the original go response and then execute the alternative change response (Boecker et al., 2013; Camalier et al., 2007; Jha et al., 2015; Verbruggen, Schneider, & Logan, 2008). In selective stop tasks, subjects are instructed to stop their response on some signal trials, but not on others (for a short review, see Bissett & Logan, 2014). There are two variants of the selective stop task: in stimulus selective stop tasks, different signals can be presented and subjects must stop if one of them occurs (valid signal), but not if the others occur (invalid signals); in motor selective stop tasks, subjects must stop some of their responses (critical responses) but not others (non-critical responses).

The independent horse-race model has been applied to the stop-change task and the selective stop task to estimate SSRT. Several studies indicate that going in the primary go task and stopping are independent in the stop-change paradigm. For example, Logan and Burkell (1986) directly compared performance in a stop-change task (with only valid signals) with performance in a dual-task paradigm. They found a standard dual-task effect in the dual-task task: When the delay between two go stimuli decreased, the latency of the second response increased (indicating dual-task interference). A similar dual-task effect was observed on signal-respond trials in the stop-change task: When the delay between the go stimulus and the change signal decreased, the latency of the change response increased (indicating dual-task interference). However, when inhibition of the first response was successful, stop-change performance was not affected much by the delay between the go stimulus and the change signal (see e.g., Hübner & Drucey, 2006, for a replication). In another study, Verbruggen, Schneider et al. (2008) manipulated the delay between the stop signal and a signal indicating which change response had to be executed. As this delay increased, the probability of stopping the primary task response changed very little, which indicates that the stop process was not influenced by the selection and execution of the change response. Combined, these studies indicate that stopping is largely independent from going in the primary task and going in the secondary task in the stop-change paradigm, which is consistent with the independent horse-race model.

Most researchers in the selective stop literature also assume that the decision to stop or not does not interact with ongoing go processes. (Note that they

have to make this assumption to estimate SSRT.) However, Bissett and Logan (2014) found that signal-respond RT and invalid-signal RT were sometimes longer than go RT in stimulus-selective stop tasks. A similar pattern of results was observed by de Jong, Coles, and Logan (1995) in a motor variant of the selective stop task: Signal-respond RTs for critical responses and signal RTs for non-critical responses were longer than go RT. These findings suggest that selecting the appropriate response to the signal may interact with ongoing go processes (violating the context independence assumption of the independence horse-race model; see above). Verbruggen and Logan (2015) tested the hypothesis that the go and stop process share capacity in selective stopping tasks by manipulating the consistency of mapping between signals and the requirement to stop or ignore in response to the signal. In consistent mapping conditions, each signal played the same role throughout the experiment; in varied mapping conditions, the role changed repeatedly over the course of the experiment. Following Shiffrin and Schneider (1977) and others, they assumed the varied mapping conditions would demand more capacity than the consistent mapping conditions, and so should produce larger violations of context independence. That is what they found.

These selective stopping results are interesting in contrast with simple stopping, where increasing the capacity demands of the go process has no effect on the stop process (see Section 4.1.2). We propose that this is due to the low selection demands in standard stop-signal tasks. This does not imply that capacity sharing can never occur in these tasks. The stop rate parameters depend on the discriminability, intensity, and modality of the stop signal (e.g., van der Schoot, Licht, Horsley, & Sergeant, 2005), which could be interpreted as a capacity limitation (Logan et al., 2014). Furthermore, competition between visual signals in the go and the stop tasks can influence stopping (Verbruggen, Stevens et al., 2014), which is consistent with the idea that stimuli have to compete for limited processing capacity (e.g., Bundesen, 1990; Desimone & Duncan, 1995). Finally, “functional dependence” (see Section 2.5) could also be interpreted as a capacity limitation. Thus, it seems that under certain circumstances, capacity sharing may occur in simple stop-signal and stop-change tasks.

6.2. Discrete versus Continuous Tasks

Most stop-signal tasks involve the execution and inhibition of discrete key presses. A few studies have also explored stopping in continuous stop-signal tasks (e.g., Morein-Zamir, Chua, Franks, Nagelkerke, & Kingstone, 2006; Morein-Zamir, Nagelkerke, Chua, Franks, & Kingstone, 2004; Morein-Zamir & Meiran, 2003). In such tasks, a target moves on the screen and participants are instructed to track it with a mouse or by pressing a force sensor. After a variable delay, a stop signal is presented, instructing the participant to stop the continuous response as

quickly as possible. SSRT can be defined as the moment at which substantial deceleration (Morein-Zamir & Meiran, 2003) or pressure offset (Morein-Zamir et al., 2004) occurs.

A main advantage of a continuous stop task is that the mean and the variability of SSRT can be measured directly. For example, Morein-Zamir, Hommersen, Johnston, and Kingstone (2008) examined performance of children with ADHD and matched control participants in a discrete (standard) stop-signal task and in a continuous (force-pressure) variant. In both tasks, SSRT was longer for children with ADHD than for the control children. This is consistent with other studies (for meta-analyses, see e.g., Oosterlaan, Logan, & Sergeant, 1997; Lipszyc & Schachar, 2010). Furthermore, the continuous variant revealed that stopping latency was also more variable in children with ADHD. Thus, stopping seems both slowed and more variable in children with ADHD.

The direct measurement of SSRT in continuous stop-signal tasks brings two additional advantages. First, fewer trials may be required to obtain a reliable SSRT measure. Second, SSRT can be measured even when the independence assumptions are violated. As discussed in Section 5.3, in discrete stop-signal tasks, SSRT estimates may be unreliable when the assumptions of the independent horse-race model are violated. Continuous stop tasks do not require the independence assumptions to estimate SSRT. Therefore, they can provide an index of inhibitory control (broadly defined) even when going and stopping interact or share processing capacity (for an alternative procedure, see Verbruggen & Logan, 2015).

In sum, continuous variants of the stop task seem to have certain advantages. However, only a few studies have used these tasks, and it remains unclear to what extent the same cognitive and neural mechanisms are involved in stopping discrete and continuous responses. Brunamonti, Ferraina, and Paré (2012) compared stop performance in tasks in which participants had to press a button with a finger, move a joystick with their wrists, or reach to a stimulus with their arms. SSRT was similar in all tasks, indicating that common inhibitory control mechanisms were involved (see also Chen, Scangos, & Stuphorn, 2010). Furthermore, Morein-Zamir et al. (2004) found that SSRTs in discrete and continuous tasks are highly correlated. These findings indicate an overlap in control mechanisms. But despite the large overlap, some studies indicate differences between controlling continuous and discontinuous movements (e.g., Spencer, Zelaznik, Diedrichsen, & Ivry, 2003). Furthermore, many processes are involved in stopping actions (see Section 7.3). Thus, further research is required to determine which control processes overlap and which processes differ.

7. Users' Guidelines

The soundness of conclusions from stop-signal studies depends on the quality of the data and the validity of the resulting SSRT estimates. In this section we present a number of recommendations on how to run, report, and interpret the results from stop-signal experiments.

7.1. How to Run Stop-Signal Experiments

7.1.1 How to Collect Stop-Signal Data

The stop signal paradigm is simple and elegant but conducting experiments is complicated by inherent tradeoffs between stopping and going: Participants succeed at the go task by going faster but they succeed at the stop task by going slower. Somehow, they must balance these demands. Many studies have shown how the balance they choose can be influenced by factors in the experimental design. The most important factor is the predictability of the stop signal: If the stop signal is predictable, participants will adjust their behavior to exploit the predictability.

Recommendation 1: Use a broad range of stop-signal delays. One important dimension of stop signal predictability is stop-signal delay. Participants adapt to the range of delays in the experiment (Lappin & Eriksen, 1966; Logan, 1981; Ollman, 1973), slowing go RT to increase the probability of stopping. Best performance is obtained with a broad range of delays that span the entire inhibition function (Logan, 1981). Under those conditions, the occurrence of the stop signal is maximally unpredictable, so participants have no predictability to exploit. This is easily accomplished by setting fixed delays and it is usually accomplished by the tracking procedure, which often produces bell-shaped distributions of stop-signal delays. We caution against more sophisticated tracking procedures that reduce the step size to converge on a single value, as that would reduce the range of stop-signal delays and increase the predictability of the stop signal. It may be better to combine them with two fixed delays, one so early that participants can nearly always stop and one so late that participants can rarely or never stop (e.g., Janssen, Heslenfeld, van Mourik, Logan, & Oosterlaan, 2015).

Recommendation 2: Present stop signals on a minority of trials. Another important dimension of stop signal predictability is the probability that a stop signal will occur on a given trial. Participants slow down as stop signal probability increases (Logan, 1981; Logan & Burkell, 1986), even in the tracking procedure, which keeps $P_{Respond}$ constant at 0.5 (Bissett & Logan, 2011; Verbruggen & Logan, 2009b). Stop signal probability typically varies between 0.1

and 0.3. Larger values produce greater slowing that may reflect strategic changes in the go task. Other things equal, we recommend choosing a stop signal probability between 0.1 and 0.3.

Recommendation 3: Take steps to avoid slowing in anticipation of stop signals. Participants almost always slow go RT when stop signals are presented. The slowing appears to result from a proactive strategy intended to increase probability of successful inhibition. It can be elicited by cues indicating that stop signals may occur on the next few trials: Slowing occurs on the trial immediately after the cue, before any stop signals have been presented (Verbruggen & Logan, 2009). Proactive slowing can be modeled successfully as an increase in the threshold for the go response, which is a strategically controlled parameter in the race diffusion model (Logan et al., 2014; Verbruggen & Logan, 2009b). Furthermore, recent findings indicate that participants also adjust attentional settings when they expect a stop signal (e.g. Elchlepp, Lavric, Chambers, & Verbruggen, 2016).

Proactive slowing is ubiquitous but it is often relatively stable over the experiment. When it is stable, the race model calculations can be applied using the RTs from no-stop-signal trials to estimate the go RT distribution. However, some participants slow progressively over the experiment, as if they are trying to beat the tracking algorithm. Progressive slowing presents challenges for analysis. It biases estimates of SSRT (Verbruggen et al., 2013). Sometimes the bias can be reduced by calculating SSRT in each block and collapsing across blocks (Verbruggen et al., 2013). However, some participants slow so dramatically that the tracking algorithm cannot keep up with them. Their response probabilities on stop trials do not converge on 0.5; response rates are usually much lower. Such data sets cannot be analyzed meaningfully with the race model and so should be discarded.

What can be done to control proactive slowing in anticipation of stop signals? The recommendations in Logan (1994) are still effective: Introduce the go task first without the stop task and allow participants some practice to get a feel for the task. Perhaps present them with some feedback about their RT and accuracy at the end of this practice. Then introduce the stop task and explicitly instruct participants not to slow their go RTs. Perhaps allow some practice before collecting the data for the main experiment. To control progressive slowing, we have been giving participants feedback about go RT and accuracy (number of incorrect trials and number of missed trials) and the probability of inhibition at the end of each block during the experiment (e.g., Bissett & Logan, 2011; Verbruggen, Stevens, et al., 2014). We have participants write down the numbers and give them to us to be sure they attend to them.

Recommendation 4: Look for trigger failures and correct for them. Participants sometimes ignore the stop signal entirely, responding whether or not a

stop signal occurs (Logan & Cowan, 1984). Such trigger failures can bias estimates of stopping latencies, let these be summary measures or SSRT distributions, and result in distorted inhibition functions (Band et. al., 2003; Matzke, Love, & Heathcote, 2015; Verbruggen, Stevens, et al., 2014). Failures to trigger the stop process on a constant proportion of the stop-signal trials compress the inhibition function. The lower asymptote of the compressed inhibition function equals the probability of a trigger failure. Formally, for a given trigger failure probability P_{TF} , the response rate on a given stop-signal delay is given by:

$$P_{Respond}(P_{TF}, t_{SSD}) = (1 - P_{TF})P_{Respond}(t_{SSD}) + P_{TF}. \quad (29)$$

Thus, a lower asymptote greater than zero is diagnostic of trigger failures. Few stop-signal studies include enough short stop-signal delays to estimate the lower asymptote accurately, however. Alternatively, one may fit a Weibull function to the inhibition function with the minimum and maximum point as free parameters; the estimated minimum point would reflect the probability of trigger failures (e.g., Hanes et al., 1998)

Trigger failures also result in signal-respond RT distributions that are mixtures of the “true” signal respond RT distribution and the go RT distribution:

$$f_{SR}(t|t_{SSD}, P_{TF}) = P_{TF}f_{go}(t) + (1 - P_{TF})f_{SR}(t|t_{SSD}) \quad (30)$$

Mixture distributions generally have larger variability than their parents, so the variance in signal-respond RTs is inflated by trigger failures. Greater variance in signal-respond RT than in go RT may be diagnostic of trigger failures. Plots of signal-respond and go RT distributions may also be informative. Without trigger failures, the upper tail of the signal-respond RT distribution (e.g., the 95th percentile) is shorter than the upper tail of the go RT distribution. With trigger failures, the upper tail of the observed signal-respond RT distribution is also the upper tail of the go RT distribution. Thus, failures of signal-respond and go RT distributions to diverge at the upper quantiles may be diagnostic of trigger failures.

In order to estimate the probability of trigger failures and correct the resulting bias in SSRT estimates, Matzke et al. (2015) proposed to parametrize the mixture in Equation 30 assuming ex-Gaussian distributions for the go RT and SSRT distributions (see also Matzke et al., 2013a). The Bayesian hierarchical implementation of the trigger-failure model provides accurate and precise parameter estimates with relatively scarce data. Matzke and colleagues reanalyzed two published stop-signal data sets (Badcock et al., 2002; Hughes, Fulham, Johnston, & Michie, 2012) and showed that the trigger-failure model provided a better description of the data than the standard ex-Gaussian Bayesian parametric

approach (Matzke et al., 2013a). On average, participants failed to trigger the stop process on 8-9% of the stop-signal trials. Although the probability of trigger failure was relatively modest, its presence was shown to severely distort SSRT estimates.

7.1.2 How to Analyze Stop-Signal Data

Recommendation 1: Fit process models to the data and interpret the data in terms of those processes. The process models described in Section 4 provide good accounts of observed behavior and the underlying physiology, describing performance as a stochastic decision (Boucher et al., 2007; Hanes & Carpenter, 1999; Logan et al., 2015; Logan et al., 2014). These models estimate the distribution of SSRTs as well as its mean, and the distributions may reveal interesting differences between conditions and groups. The models interpret performance in terms of drift rates, starting points, and thresholds. Concepts like strategic slowing, post stop-signal slowing, and inhibitory deficits might be better articulated in terms of these more fundamental properties of the decision process than simply in terms of mean SSRT. Stochastic decision models have provided tremendous insight into go processes and go RT (Ratcliff, Smith, Brown & McKoon, 2016). They should provide similar insights into the stop-signal task. Note that analyzing stop-signal data using process models requires more data points (and therefore longer experiments) than computing summary measures of SSRT using traditional estimation methods.

Recommendation 2: Estimate the distribution of SSRTs. If researchers are not interested in the details of the underlying process and are satisfied with measures of SSRT, we suggest they harness the Bayesian parametric approach (Matzke et al., 2013a) and its trigger-failure variant (Matzke et al., 2015) to estimate the distribution of SSRTs. User friendly software that implements powerful Bayesian methods is freely available (BEESTS; Matzke et al., 2013b), and analyses of distributions may reveal patterns of data that are obscured in the means. For example, mean $SSRT = \mu_{stop} + \tau_{stop}$ so μ and τ may differ between conditions that produce equivalent mean SSRTs. They have different effects on variability and can be separated by fitting BEESTS to the data.

Recommendation 3: Use the mean method with caution. The mean method produces valid, mathematically justified estimates of mean SSRT if the independence assumptions hold and the means of the go distribution and the inhibition function are estimated accurately (Logan & Cowan, 1984). Unfortunately, few studies estimate the mean of the inhibition function directly (e.g., using Equation 11). Instead, most estimate the median of the inhibition function with the tracking procedure. If the inhibition function is asymmetrical,

the median will underestimate the mean, and consequently, overestimate SSRT (Verbruggen et al., 2013). Inhibition functions are likely to be asymmetrical when the go RT distribution is skewed. This can be seen in Equation 3, which defines the inhibition function when SSRT is constant as the integral of the go RT distribution from 0 to $(t_{stop} + t_{SSD})$. The inhibition function is simply the go RT distribution shifted to the right by SSRT. Any skew in the go RT distribution will necessarily appear in the inhibition function. Since most RT distributions are skewed, we can expect most inhibition functions to be skewed, and thus, we can expect the median of most inhibition functions to underestimate the mean we need for calculating SSRT. What should researchers do about that?

The simplest possibility is to abandon the mean method and use the integration method instead, as Verbruggen et al. (2013) recommended, but the simplicity of the mean method is seductive. If researchers succumb to the seduction, we recommend that they check for skew in the go RT distributions. One method for checking skew is to fit the ex-Gaussian distribution to the go RTs and examine estimates of τ . Verbruggen et al. showed that estimation errors occurred primarily with large values of τ , so researchers might compare their values of τ with the values Verbruggen et al. used to determine whether the skew in the go RTs compromises estimation of SSRT. Similarly, researchers could compare τ between groups to see if group differences in SSRT might be artifacts of differences in skew. However, if researchers are willing to fit ex-Gaussian distributions to their go RTs, we recommend they fit BEESTS to the whole data set and get estimates of the entire distribution of SSRTs.

We also suggest trying to estimate the mean of the inhibition function directly, through Equation 11, and using the mean to calculate SSRT, as the race model dictates. This works best when the entire inhibition function can be estimated, as in experiments with a broad range of fixed delays. The effects of skew on the mean depend on the tails of the distribution, so estimating the mean of a truncated inhibition function may underestimate the actual mean. The tracking procedure typically produces a bell-shaped distribution of stop-signal delays with sparse tails that may not extend to the extremes of the inhibition function ($P_{Respond} = 0$ or 1). We have not explored this possibility through simulations, but it would be very informative to do so.

We recommend against using the median method, in which the median of the inhibition function (the mean stop-signal delay in the tracking procedure) is subtracted from the median go RT. This solves the problem of estimating the appropriate parameter of the inhibition function – the tracking procedure converges on the median – but the calculations are not justified in the race model. The race model calculations are in terms of means, not medians (Logan & Cowan, 1984). We have not explored the relation between the mean method and the

median method mathematically or with simulations, but researchers interested in using the median method instead of the mean method should do so.

Recommendation 4: Otherwise, use the integration method. The integration method calculated at stop-signal delays near the middle of the inhibition function yields accurate, unbiased estimates of SSRT (Band et al., 2003; Verbruggen et al., 2013). We recommend the integration method to researchers who are interested primarily in mean SSRT. With fixed delays, researchers should calculate SSRT at each stop-signal delay and average over stop-signal delays, as SSRT decreases with stop-signal delay. Note that SSRT cannot be calculated if $P_{Respond} = 1$ or 0. Stop-signal delays that produce $P_{Respond} = 1$ or 0 should be excluded from analysis. With delays set by the tracking procedure, researchers should calculate SSRT with the integration method, using the overall $P_{Respond}$ as the limit of integration and using the mean stop-signal delay as the stop-signal delay value. However, this method is compromised if there is progressive slowing over the experiment. If there is evidence of progressive slowing, integration SSRT should be calculated in each block and averaged over blocks. This can correct for progressive slowing if the slowing is not too extreme (Verbruggen et al., 2013).

7.2. How to Report Stop-Signal Experiments

We recommend that reports of stop-signal experiments contain sufficient information to allow an evaluation of the fit of the original horse-race model, on which most calculations will rely. We propose that every stop-signal study should report the following:

1. Report the procedure in enough detail that it can be evaluated. Report the number of trials overall, the number of stop-signal trials (i.e., the probability of a stop signal), the range and value of stop-signal delays used, the method used to calculate SSRT, and the number of observations used in that calculation.

2. Report mean signal-respond and go RT and confirm they are significantly different in each experimental condition. With tracking, this can be done collapsing over delays. With fixed delays, it should be done at each delay, noting that signal-respond RT becomes more similar to go RT as stop-signal delay increases, so differences need not be significant at the longest delay.

3. Confirm that signal-respond RT is shorter than go RT for every participant for whom SSRT is estimated. SSRT should not be estimated for participants with signal-respond RTs longer than go RTs, as these participants violate the independence assumptions of the race model. The number of participants excluded for this reason should be reported. The criterion for assessing the difference within participants is unclear. The simplest would be to

conduct a t test within each participant, but that may be too strict a criterion. The most minimal criterion would be to subtract signal response RT from go RT and conclude signal response RT was smaller if the difference was positive. Despite the ambiguity about the most appropriate criterion, we believe researchers should make this comparison and report it.

4. Report the response rate given a stop signal in each condition. With fixed delays, this means reporting the inhibition function in each condition (i.e., the probability of responding at each stop-signal delay). With tracking, the probability of responding should be calculated for each condition. Some researchers have reported inhibition functions from tracking procedures (e.g., Thakkar, Schall, Boucher, Logan, & Park, 2011) but they are often noisy at the tails where there are few observations so response rate estimates are unstable.

5. When using the tracking procedure, report the mean stop-signal delay for each condition so readers know the baseline from which SSRT was computed.

6. Use an appropriate method to estimate SSRT. We recommend process models, then parametric models, and then the integration method, depending on researchers' goals and interests. Researchers who use the mean method with the tracking procedure (where the mean stop-signal delay estimates the median of the inhibition function) should address concerns about skew compromising their SSRT estimates discussed above (Verbruggen et al., 2013).

7.3. How to Interpret Stop-Signal Data

A final note concerns the interpretation of the stop-signal data. In the stop-signal literature, individual or group differences are often attributed to variation in the effectiveness of a single inhibitory control function. But many processes contribute to stopping an action. As discussed in Section 4, response inhibition often requires an interplay between basic and computationally well-defined "reactive" processes, such as signal detection, action selection, and suppression of motor output or blocking go input. These processes can be regulated and influenced by sets of processes that take place on different timescales: outcome monitoring, advance preparation (i.e., proactive control), rule acquisition and maintenance, associative learning, and development (Verbruggen, McLaren et al., 2014).

Thus, it is important to realize that SSRT is a global concept that describes the chain of processes involved in an act of control that results in a response being withheld. More specifically, SSRT captures the duration of perceptual, decisional, and (inhibitory) motor-related processes. For example, previous behavioral studies and computational work have highlighted the role of perceptual processes (see

above). Successfully stopping a response also depends on decisional processes, such as response selection and memory retrieval (e.g. Logan et al., 2014; van de Laar, van den Wildenberg, van Boxtel, & van der Molen, 2010; Verbruggen & Logan, 2015). Finally, when the decision to stop is reached, motor output or other ongoing processing has to be suppressed (e.g., via a fronto-basal-ganglia network) or go input has to be blocked. Thus, in simple stop-signal tasks and their many variants, SSRT reflects more than the duration of a single neural inhibitory process, and researchers should consider at which processing stage(s) differences between groups or conditions arise (for a more elaborate discussion of this issue, see e.g., Verbruggen, McLaren, et al., 2014).

8. Concluding Comments

Response inhibition refers to the ability to suppress responses that are inappropriate or no longer required, such as rapidly halting when the traffic light turns red. Response inhibition is considered a key component of executive control and has received—and continues to receive—considerable attention in fields as diverse as psychology, pharmacology, psychiatry, neurology, and biology (Verbruggen et al., 2013). In laboratory settings, response inhibition is typically investigated with the stop-signal paradigm. The stop-signal paradigm owes its popularity to the underlying horse-race model (Logan & Cowan, 1984) that facilitates the estimation of the latency of the otherwise unobservable stop response.

We presented a theoretical review of the horse-race model and discussed the most important measures of response inhibition performance in the stop-signal paradigm. We first outlined the standard independent horse-race model and related SSRT estimation techniques, and showed that the independent race architecture typically offers an excellent description of stop-signal data across different populations, tasks, and experimental manipulations. We then described the latest developments in the model-based analysis of stop-signal data, focusing on the simultaneous estimation of SSRT distributions and trigger failures and variants of the standard horse-race model that give direct insights into the mechanisms of stopping. In particular, we discussed two classes of process models of response inhibition: models that describe the properties of the go and the stop process in order to explain how the finishing time distributions arise and models that attempt to explain how responses are stopped. Although these models lack the generality of the standard independent horse-race model, they provide fine-grained insights into the mechanisms of stopping. We believe that the application of process

models to more complex variants of the stop-signal task, such as the stop-change and selective stopping tasks, is a promising area for future research that may also benefit from recent developments in Bayesian hierarchical modeling and related model selection methods.

Definitions and Terms

Response inhibition

The cognitive concept of response inhibition refers to the ability to suppress responses that are inappropriate or no-longer required, which supports flexible and goal-directed behavior in ever-changing environments. Response inhibition is a key component of executive control.

Stop-signal paradigm

The stop-signal paradigm is a popular experimental paradigm to study response inhibition. The standard stop-signal paradigm consists of a two-choice response time task. The primary choice task is occasionally interrupted by a stop signal that instructs participants to withhold their response on that trial.

Horse-race model

The horse-race model posits that response inhibition in the stop-signal paradigm can be conceptualized as a horse race between two independent processes: a go and a stop process. If the go process wins the race, the response is executed; if the stop process wins the race, the response is inhibited. According to the horse-race model, response inhibition is thus determined by the relative finishing times of the go and the stop process.

Stop-signal reaction time

Stop-signal reaction time is the latency of the stop process. Although stop-signal reaction time cannot be observed directly, it can be estimated using the horse-race model. Stop-signal reaction times play a pivotal role in diagnosing deficient response inhibition in clinical populations and in assessing participants' stopping ability across different tasks and experimental conditions.

Inhibition function

Inhibition functions describe the relationship between response rate and the time interval between the onset of the primary task stimulus and the onset of the stop-signal (i.e., stop-signal delay). The horse-race model predicts that response rate increases with increasing stop-signal delay. Inhibition functions reflect the outcome of the race between the go and the stop process and can be used to compare inhibitory control across populations, tasks, or conditions.

type="list"

List of Abbreviations

- ADHD: Attention Deficit/Hyperactivity Disorder
- LATER: Linear Approach to Threshold with Ergodic Rate
- RT: Response Time
- SSRT: Stop-Signal Reaction Time
- SSD: Stop-Signal Delay

References

- Aron, A.R., Fletcher, P.C., Bullmore, E.T., Sahakian, B.J., & Robbins, T.W. (2003). Stop-signal inhibition disrupted by damage to right inferior frontal gyrus in humans. *Nature Neuroscience*, 6, 115–1166. <http://doi.org/10.1038/nm1003>
- Aron, A.R., & Poldrack, R.A. (2006). Cortical and subcortical contributions to stop signal response inhibition: Role of the subthalamic nucleus. *Journal of Neuroscience*, 26, 2424–2433. <http://doi.org/10.1523/JNEUROSCI.4682-05.2006>
- Aron, A.R., Robbins, T.W., & Poldrack, R.A. (2014). Inhibition and the right inferior frontal cortex: One decade on. *Trends in Cognitive Sciences*, 18, 177–185. <http://doi.org/10.1016/j.tics.2013.12.003>
- Badcock, J.C., Michie, P., Johnson, L., & Combrinck, J. (2002). Acts of control in schizophrenia: Dissociating the components of inhibition. *Psychological Medicine*, 32, 287–297.
- Badry, R., Mima, T., Aso, T., Nakatsuka, M., Abe, M., Fathi, D., ... , & Fukuyama, H. (2009). Suppression of human cortico-motoneuronal excitability during the stop-signal task. *Clinical Neurophysiology*, 120, 1717–1723. <http://doi.org/10.1016/j.clinph.2009.06.027>
- Band, G.P.H., van der Molen, M.W., & Logan, G.D. (2003). Horse-race model simulations of the stop-signal procedure. *Acta Psychologica*, 112, 105–142.
- Bari, A., & Robbins, T.W. (2013). Inhibition and impulsivity: Behavioral and neural basis of response control. *Progress in Neurobiology*, 108, 44–79. <http://doi.org/10.1016/j.pneurobio.2013.06.005>
- Bechara, A., Noel, X., & Crone, E.A. (2006). Loss of willpower: Abnormal neural mechanisms of impulse control and decision making in addiction. *Handbook of Implicit Cognition and Addiction*, 215–232.
- Bissett, P.G., & Logan, G.D. (2011). Balancing cognitive demands: Control adjustments in the stop-signal paradigm. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 37, 392–404.
- Bissett, P.G., & Logan, G.D. (2014). Selective stopping? Maybe not. *Journal of Experimental Psychology: General*, 143, 455–72. <http://doi.org/10.1037/a0032122>
- Boecker, M., Gauggel, S., & Druke, B. (2013). Stop or stop-change - Does it make any difference for the inhibition process? *International Journal of Psychophysiology*, 87, 234–243. <http://doi.org/10.1016/j.ijpsycho.2012.09.009>

- Boehler, C.N., Schevernels, H., Hopf, J.-M.M., Stoppel, C.M., & Krebs, R.M. (2014). Reward prospect rapidly speeds up response inhibition via reactive control. *Cognitive Affective & Behavioral Neuroscience, 14*, 593–609. <http://doi.org/10.3758/s13415-014-0251-5>
- Boucher, L., Palmeri, T.J., Logan, G.D., & Schall, J.D. (2007). Inhibitory control in mind and brain: An interactive race model of countermanding saccades. *Psychological Review, 114*, 376–397.
- Brown, S.D., & Heathcote, A.J. (2008). The simplest complete model of choice reaction time: Linear ballistic accumulation. *Cognitive Psychology, 57*, 153–178.
- Brunamonti, E., Ferraina, S., & Paré, M. (2012). Controlled movement processing: Evidence for a common inhibitory control of finger, wrist, and arm movements. *Neuroscience, 215*, 69–78.
- Bundesen, C. (1990). A theory of visual attention. *Psychological Review, 97*, 523–547. <http://doi.org/10.1037/0033-295X.97.4.523>
- Camalier, C.R., Gotler, R., Murthy, A., Thompson, K.G., Logan, G.D., Palmeri, T.J., & Schall, J.D. (2007). Dynamics of saccade target selection: Race model analysis of double step and search step saccade production in human and macaque. *Vision Research, 47*, 2187–2211.
- Carpenter, R.H.S. (1981). Oculomotor procrastination. In D.F. Fisher, R.A. Monty, & J.W. Senders (Eds.), *Eye movements: Cognition and visual perception* (pp. 237–246). Hillsdale, NJ: Erlbaum.
- Carpenter, R.H.S., & Williams, M.L.L. (1995). Neural computation of log likelihood in control of saccadic eye movements. *Nature, 377*, 59–62.
- Chambers, C.D., Bellgrove, M.A., Gould, I.C., English, T., Garavan, H., McNaught, E., Kamke, M., & Mattingley, J.B. (2007). Dissociable mechanisms of cognitive control in prefrontal and premotor cortex. *Journal of Neurophysiology, 98*, 3638–3647.
- Chambers, C.D., Bellgrove, M.A., Stokes, M.G., Henderson, T.R., Garavan, H., Robertson, I.H., Morris, A.P., & Mattingley, J.B. (2006). Executive “brake failure” following deactivation of human frontal lobe. *Journal of Cognitive Neuroscience, 18*, 444–455.
- Chambers, C.D., Garavan, H., & Bellgrove, M.A. (2009). Insights into the neural basis of response inhibition from cognitive and clinical neuroscience. *Neuroscience & Biobehavioral Reviews, 33*, 631–646. <http://doi.org/10.1016/j.neubiorev.2008.08.016>
- Chen, X., Scangos, K.W., & Stuphorn, V. (2010). Supplementary motor area exerts proactive and reactive control of arm movements. *Journal of Neuroscience, 30*, 14657–14675. <http://doi.org/10.1523/JNEUROSCI.2669-10.2010>

- Chevalier, N., Chatham, C.H., & Munakata, Y. (2014). The practice of going helps children to stop: The importance of context monitoring in inhibitory control. *Journal of Experimental Psychology: General*, *143*, 959-965. <http://doi.org/10.1037/a0035868>
- Colonius, H. (1990). A note on the stop-signal paradigm, or how to observe the unobservable. *Psychological Review*, *97*, 309–312.
- Colonius, H., Ozyurt, J., & Arndt, P.A. (2001). Countermanding saccades with auditory stop signals: Testing the race model. *Vision Research*, *41*, 1951–1968.
- Colzato, L.S., Jongkees, B.J., Sellaro, R., van den Wildenberg, W.P.M., & Hommel, B. (2014). Eating to stop: Tyrosine supplementation enhances inhibitory control but not response execution. *Neuropsychologia*, *62*, 398-402.
- Congdon, E., Mumford, J.A., Cohen, J.R., Galvan, A., Canli, T., & Poldrack, R.A. (2012). Measurement and reliability of response inhibition. *Frontiers in Psychology*, *3*:37. doi:10.3389/fpsyg.2012.00037
- Corneil, B.D., & Elsley, J.K. (2005). Countermanding eye-head gaze shifts in humans: Marching orders are delivered to the head first. *Journal of Neurophysiology*, *94*, 883–895.
- Coxon, J.P., Stinear, C.M., & Byblow, W.D. (2006). Intracortical inhibition during volitional inhibition of prepared action. *Journal of Neurophysiology*, *95*, 3371–3383. <http://doi.org/10.1152/jn.01334.2005>
- Crews, F.T., & Boettiger, C.A. (2009). Impulsivity, frontal lobes and risk for addiction. *Pharmacology, Biochemistry, and Behavior*, *93*, 237–247. <http://doi.org/10.1016/j.pbb.2009.04.018>
- de Jong, R., Coles, M.G., & Logan, G.D. (1995). Strategies and mechanisms in nonselective and selective inhibitory motor control. *Journal of Experimental Psychology: Human Perception and Performance*, *21*, 498–511.
- de Jong, R., Coles, M.G., Logan, G.D., & Gratton, G. (1990). In search of the point of no return: The control of response processes. *Journal of Experimental Psychology: Human Perception and Performance*, *16*, 164–182.
- de Wit, H. (2009). Impulsivity as a determinant and consequence of drug use: A review of underlying processes. *Addiction Biology*, *14*, 22–31. <http://doi.org/10.1111/j.1369-1600.2008.00129.x>
- Debey, E., De Schryver, M., Logan, G.D., Suchotzki, K., & Verschuere, B. (2015). From junior to senior Pinocchio: A cross-sectional lifespan investigation of deception. *Acta Psychologica*, *160*, 58-68.

- Desimone, R., & Duncan, J. (1995). Neural mechanisms of selective visual attention. *Annual Review of Neuroscience*, *18*, 193–222.
- Elchlepp, H., Lavric, A., Chambers, C.D., & Verbruggen, F. (2016). Proactive inhibitory control: A general biasing account. *Cognitive Psychology*, *86*, 27–61.
<http://doi.org/10.1016/j.cogpsych.2016.01.004>
- Ersche, K.D., Jones, P.S., Williams, G.B., Turton, A.J., Robbins, T.W., & Bullmore, E.T. (2012). Abnormal brain structure implicated in stimulant drug addiction. *Science*, *335*, 601–604.
<http://doi.org/10.1126/science.1214463>
- Farrell, S., & Ludwig, C.J.H. (2008). Bayesian and maximum likelihood estimation of hierarchical response time models. *Psychonomic Bulletin & Review*, *15*, 1209–1217.
- Fernie, G., Peeters, M., Gullo, M.J., Christiansen, P., Cole, J.C., Sumnall, H. & Field, M. (2013). Multiple behavioural impulsivity tasks predict prospective alcohol involvement in adolescents. *Addiction*, *108*, 1916–1923.
- Fillmore, M.T., Rush, C.R., & Hays, L. (2002). Acute effects of oral cocaine on inhibitory control of behavior in humans. *Drug and Alcohol Dependence*, *67*, 157–167.
- Garavan, H., & Stout, J.C. (2005). Neurocognitive insights into substance abuse. *Trends in Cognitive Sciences*, *9*, 195–201. <http://doi.org/10.1016/j.tics.2005.02.008>
- Gelman, A., & Hill, J. (2007). *Data analysis using regression and multilevel/hierarchical models*. Cambridge, UK: Cambridge University Press.
- Gelman, A., Meng, X., & Stern, H. (1996). Posterior predictive assessment of model fitness via realized discrepancies. *Statistica Sinica*, *6*, 733–807.
- Gilks, W.R., Richardson, S., & Spiegelhalter, D.J. (1996). *Markov chain Monte Carlo in practice*. Boca Raton, FL: Chapman & Hall/CRC.
- Green, D.M., & Swets, J.A. (1966). *Signal detection theory and psychophysics*. New York: John Wiley.
- Greenhouse, I., Oldenkamp, C.L., & Aron, A.R. (2012). Stopping a response has global or nonglobal effects on the motor system depending on preparation. *Journal of Neurophysiology*, *107*, 384–392. <http://doi.org/10.1152/jn.00704.2011>
- Hanes, D.P., & Carpenter, R.H.S. (1999). Countermanding saccades in humans. *Vision Research*, *39*, 2777–2791.
- Hanes, D.P., Patterson, W.F., & Schall, J.D. (1998). Role of frontal eye fields in countermanding saccades: Visual, movement, and fixation activity. *Journal of Neurophysiology*, *79*, 817–834.

- Hanes, D.P., & Schall, J.D. (1995). Countermanding saccades in macaque. *Visual Neuroscience*, *12*, 929–937.
- Hanes, D.P., & Schall, J.D. (1996). Neural control of voluntary movement initiation. *Science*, *274*, 427–430.
- Heathcote, A. (2004). Fitting Wald and ex-Wald distributions to response time data: An example using functions for the S-PLUS package. *Behavior Research Methods, Instruments & Computers*, *36*, 678-694.
- Heathcote, A., Brown, S., & Cousineau, D. (2004). QMPE: Estimating Lognormal, Wald, and Weibull RT distributions with a parameter-dependent lower bound. *Behavior Research Methods*, *36*, 277-290.
- Heathcote, A., Popiel, S.J., & Mewhort, D.J. (1991). Analysis of response time distributions: An example using the Stroop task. *Psychological Bulletin*, *109*, 340-347.
- Hübner, R., & Druwey, M.D. (2006). Response execution, selection, or activation: What is sufficient for response-related repetition effects under task shifting? *Psychological Research*, *70*, 245–261. <http://doi.org/10.1007/s00426-005-0219-8>
- Hughes, M.E., Fulham, W.R., Johnston, P.J., & Michie, P.T. (2012). Stop-signal response inhibition in schizophrenia: Behavioural, event-related potential and functional neuroimaging data. *Biological Psychology*, *89*, 220-231.
- Huizinga, M., Dolan, C.V., & van der Molen, M.W. (2006). Age-related change in executive function: Developmental trends and a latent variable analysis. *Neuropsychologia*, *44*, 2017–2036. <http://doi.org/10.1016/j.neuropsychologia.2006.01.010>
- Janssen, T.W.P., Heslenfeld, D.J., van Mourik, R., Logan, G.D., & Oosterlaan, J. (2015). Neural correlates of response inhibition in children with attention-deficit/hyperactivity disorder: A controlled version of the stop-signal task. *Psychiatry Research: Neuroimaging*, *233*, 278-284.
- Jha, A., Nachev, P., Barnes, G., Husain, M., Brown, P., & Litvak, V. (2015). The frontal control of stopping. *Cerebral Cortex*, *25*, 4392-406. <http://doi.org/10.1093/cercor/bhv027>
- Kornylo, K., Dill, N., Saenz, M., & Krauzlis, R.J. (2003). Canceling of pursuit and saccadic eye movements in humans and monkeys. *Journal of Neurophysiology*, *89*, 2984–2999.
- Kramer, A.F., Humphrey, D.G., Larish, J.F., Logan, G.D. & Strayer, D.L. (1994). Aging and inhibition: Beyond a unitary view of inhibitory processing in attention. *Psychology and Aging*, *9*, 491-512.

- Lappin, J.S., & Eriksen, C.W. (1966). Use of a delayed signal to stop a visual reaction-time response. *Journal of Experimental Psychology*, 72, 805–811.
- Lee, M.D., & Wagenmakers, E.-J. (2013). *Bayesian modeling for cognitive science: A practical course*. Cambridge, UK: Cambridge University Press.
- Lipszyc, J., & Schachar, R. (2010). Inhibitory control and psychopathology: A meta-analysis of studies using the stop signal task. *Journal of the International Neuropsychological Society*, 16, 1064-1076. <http://doi.org/10.1017/S1355617710000895>
- Logan, G.D. (1981). Attention, automaticity, and the ability to stop a speeded choice response. In J. Long & A.D. Baddeley (Eds.), *Attention and Performance IX*. Hillsdale, NJ: Erlbaum.
- Logan, G.D. (1983). On the ability to inhibit simple thoughts and actions: I. Stop signal studies of decision and memory. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 9, 585–606.
- Logan, G.D. (1994). On the ability to inhibit thought and action: A users' guide to the stop signal paradigm. In D. Dagenbach & T.H. Carr (Eds), *Inhibitory processes in attention, memory, and language*. (pp. 189-239). San Diego: Academic Press.
- Logan, G.D. (2004). Cumulative progress in formal theories of attention. *Annual Review of Psychology*, 55, 207–234.
- Logan, G.D., & Burkell, J. (1986). Dependence and independence in responding to double stimulation: A comparison of stop, change, and dual-task paradigms. *Journal of Experimental Psychology: Human Perception and Performance*, 12, 549–563.
- Logan, G.D., & Cowan, W.B. (1984). On the ability to inhibit thought and action: A theory of an act of control. *Psychological Review*, 91, 295–327.
- Logan, G.D., Cowan, W.B., & Davis, K.A. (1984). On the ability to inhibit simple and choice reaction-time responses: A model and a method. *Journal of Experimental Psychology: Human Perception and Performance*, 10, 276–291.
- Logan, G.D., Schachar, R.J., & Tannock, R. (1997). Impulsivity and inhibitory control. *Psychological Science*, 8, 60–64.
- Logan, G.D., Van Zandt, T., Verbruggen, F., & Wagenmakers, E.-J. (2014). On the ability to inhibit thought and action: General and special theories of an act of control. *Psychological Review*, 121, 66-95.
- Logan, G.D., Yamaguchi, M., Schall, J.D., & Palmeri, T.J. (2015). Inhibitory control in mind and brain 2.0: Blocked-input models of saccadic countermanding *Psychological Review*, 122, 115-147.

- MacMillan, N., & Creelman, C.D. (2004). *Detection Theory: A User's Guide* (2nd Edition). Hillsdale, NJ: Erlbaum.
- Majid, D.S.A., Cai, W., George, J.S., Verbruggen, F., & Aron, A.R. (2012). Transcranial magnetic stimulation reveals dissociable mechanisms for global versus selective corticomotor suppression underlying the stopping of action. *Cerebral Cortex*, *22*, 363–371. <http://doi.org/10.1093/cercor/bhr112>
- Matzke, D., Dolan, C.V., Logan, G.D., Brown, S.D., & Wagenmakers, E.-J. (2013a). Bayesian parametric estimation of stop-signal reaction time distributions. *Journal of Experimental Psychology: General*, *142*, 1047-1073.
- Matzke, D., Love, J., Wiecki, T., Brown, S.D., Logan, G.D., & Wagenmakers, E.-J. (2013b). Releasing the BEESTS: Bayesian estimation of stop-signal reaction time distributions. *Frontiers in Quantitative Psychology and Measurement*, *4*:918. doi: 10.3389/fpsyg.2013.00918
- Matzke, D., Love, J., & Heathcote, A. (in press). A Bayesian approach for quantifying the probability of trigger failures in the stop-signal paradigm. *Behavior Research Methods*.
- Matzke, D., & Wagenmakers, E.-J. (2009). Psychological interpretation of the ex-Gaussian and shifted Wald parameters: A diffusion model analysis. *Psychonomic Bulletin & Review*, *16*, 798-817.
- McGarry, T., & Franks, I.M. (1997). A horse race between independent processes: Evidence for a phantom point of no return in the preparation of a speeded motor response. *Journal of Experimental Psychology: Human Perception and Performance* *23*, 1533–1542.
- McGarry, T., Inglis, J.T., & Franks, I.M. (2000). Against a final ballistic process in the control of voluntary action: Evidence using the Hoffmann reflex. *Motor Control*, *4*, 469–485.
- Miyake, A., Friedman, N.P., Emerson, M.J., Witzki, A.H., Howerter, A., & Wager, T.D. (2000). The unity and diversity of executive functions and their contributions to complex “frontal lobe” tasks: A latent variable analysis. *Cognitive Psychology*, *41*, 49–100. <http://doi.org/10.1006/cogp.1999.0734>
- Morein-Zamir, S., Chua, R., Franks, I., Nagelkerke, P., & Kingstone, A. (2006). Measuring online volitional response control with a continuous tracking task. *Behavior Research Methods*, *38*, 638–647.
- Morein-Zamir, S., Hommersen, P., Johnston, C., & Kingstone, A. (2008). Novel measures of response performance and inhibition in children with ADHD. *Journal of Abnormal Child Psychology*, *36*, 1199–1210. <http://doi.org/10.1007/s10802-008-9243-7>

- Morein-Zamir, S., & Meiran, N. (2003). Individual stopping times and cognitive control: Converging evidence for the stop signal task from a continuous tracking paradigm. *Quarterly Journal of Experimental Psychology*, *56*, 469–489.
<http://doi.org/10.1080/02724980244000495>
- Morein-Zamir, S., Nagelkerke, P., Chua, R., Franks, I., & Kingstone, A. (2004). Inhibiting prepared and ongoing responses: Is there more than one kind of stopping? *Psychonomic Bulletin & Review*, *11*, 1034–1040.
- Mulvihill, L.E., Skilling, T.A., & Vogel-Sprott, M. (1997). Alcohol and the ability to inhibit behavior in men and women. *Journal of Studies on Alcohol*, *58*, 600–605.
- Munoz, D.P., & Schall, J.D. (2003). Concurrent, distributed control of saccade initiation in the frontal eye field and superior colliculus. In W.T. Hall & A. Moschovakis (Eds.), *The Superior Colliculus: New Approaches for Studying Sensorimotor Integration* (pp. 55–82). New York: CRC Press.
- Myung, I.J. (2003). Tutorial on maximum likelihood estimation. *Journal of Mathematical Psychology*, *47*, 90–100.
- Myung, I.J., Pitt, M.A., & Kim, K. (2005). Model evaluation, testing and selection. In K. Lambert & R. Goldstone (Eds.), *Handbook of Cognition* (pp. 422–436). Thousand Oaks, CA: Sage.
- Nederkoorn, C., Jansen, E., Mulkens, S., & Jansen, A. (2007). Impulsivity predicts treatment outcome in obese children. *Behaviour Research and Therapy*, *45*, 1071–1075.
<http://doi.org/10.1016/j.brat.2006.05.009>
- Nigg, J.T. (2001). Is ADHD a disinhibitory disorder? *Psychological Bulletin*, *127*, 571–598.
- Noël, X., Brevers, D., & Bechara, A. (2013). A neurocognitive approach to understanding the neurobiology of addiction. *Current Opinion in Neurobiology*, *23*, 632–638.
<http://doi.org/10.1016/j.conb.2013.01.018>
- Ollman, R.T. (1973). Simple reactions with random countermanding of the “go” signal. In S. Kornblum (Ed), *Attention and Performance IV*. New York: Academic Press.
- Oosterlaan, J., Logan, G.D., & Sergeant, J.A. (1998). Response inhibition in AD/HD, CD, comorbid AD/HD + CD, anxious, and control children: A meta-analysis of studies with the stop task. *Journal of Child Psychology and Psychiatry*, *39*, 411–25.
- Osman, A., Kornblum, S., & Meyer, D. E. (1986). The point-of-no-return in choice reaction-time—Controlled and ballistic stages of response preparation. *Journal of Experimental Psychology: Human Perception and Performance* *12*, 243–258.

- Pouget, P., Logan, G.D., Palmeri, T.J., Boucher, L., Paré, M., & Schall, J.D. (2011). Neural basis of adaptive response time adjustment during saccade countermanding. *The Journal of Neuroscience*, *31*, 12604–12612.
- Ratcliff, R., Cherian, A., & Segraves, M. (2003). A comparison of macaque behavior and superior colliculus neuronal activity to predictions from models of two-choice decisions. *Journal of Neurophysiology*, *90*, 1392–1407.
- Ratcliff, R., & Smith, P.L. (2004). A comparison of sequential sampling models for two-choice reaction time. *Psychological Review*, *111*, 333–367.
- Ratcliff, R., Smith, P.L., Brown, S.D., & McKoon, G. (2016). Diffusion decision model: Current Issues and history. *Trends in Cognitive Sciences*, *20*, 260-281.
- Ratcliff, R., & Tuerlinckx, F. (2002). Estimating parameters of the diffusion model: Approaches to dealing with contaminant reaction times and parameter variability. *Psychonomic Bulletin & Review*, *9*, 438–481.
- Ratcliff, R., Van Zandt, T., & McKoon, G. (1999). Connectionist and diffusion models of reaction times. *Psychological Review*, *106*, 261-300.
- Ridderinkhof, K.R., Band, G.P.H., & Logan, G.D. (1999). A study of adaptive behavior: Effects of age and irrelevant information on the ability to inhibit one's actions. *Acta Psychologica*, *101*, 315–337.
- Ridderinkhof, K.R., van den Wildenberg, W.P.M., Segalowitz, S.J., & Carter, C.S. (2004). Neurocognitive mechanisms of cognitive control: The role of prefrontal cortex in action selection, response inhibition, performance monitoring, and reward-based learning. *Brain and Cognition*, *56*, 129–140. <http://doi.org/10.1016/j.bandc.2004.09.016>
- Rouder, J.N., Sun, D., Speckman, P.L., Lu, J., & Zhou, D. (2003). A hierarchical Bayesian statistical framework for response time distributions. *Psychometrika*, *68*, 589–606.
- Schachar, R., & Logan, G.D. (1990). Impulsivity and inhibitory control in normal development and childhood psychopathology. *Developmental Psychology*, *26*, 710–720.
- Schall, J.D. (2004). On building a bridge between brain and behavior. *Annual Review of Psychology*, *55*, 23–50.
- Shiffrin, R.M., & Schneider, W. (1977). Controlled and automatic human information processing: II. Perceptual learning, automatic attending, and a general theory. *Psychological Review*, *84*, 127–190.

- Spencer, R.M.C., Zelaznik, H.N., Diedrichsen, J., & Ivry, R.B. (2003). Disrupted timing of discontinuous but not continuous movements by cerebellar lesions. *Science*, *300*, 1437–1439. <http://doi.org/10.1126/science.1083661>
- Szmalcew, A., Demanet, J., Vandierendonck, A., & Verbruggen, F. (2009). Investigating the role of conflict resolution in memory updating by means of the one-back choice RT task. *Psychological Research*, *73*, 390-406.
- Tannock, R., Schachar, R.J., Carr, R. P. & Logan, G.D. (1989). Dose response effects of methylphenidate on academic performance and overt behavior in hyperactive children. *Pediatrics*, *84*, 648-657.
- Tannock, R., Schachar, R., & Logan, G. (1995). Methylphenidate and cognitive flexibility: Dissociated dose effects in hyperactive-children. *Journal of Abnormal Child Psychology* *23*, 235–266.
- Thakkar, K.N., Schall, J.D., Boucher, L., Logan G.D., & Park, S. (2011). Response inhibition and response monitoring in a saccadic countermanding task in schizophrenia. *Biological Psychiatry*, *69*, 55-62.
- Townsend, J.T., & Ashby, F.G. (1983). *The stochastic modeling of elementary psychological processes*. Cambridge, England: Cambridge, University Press.
- Turner, B.M., Forstmann, B.U., Wagenmakers, E.-J., Brown, S.D., Sederberg, P.B., & Steyvers, M. (2013). A Bayesian framework for simultaneously modeling neural and behavioral data. *Neuroimage*, *72*, 193-206.
- Usher, M., & McClelland, J.L. (2001). The time course of perceptual choice: The leaky, competing accumulator model. *Psychological Review*, *108*, 550–592.
- van de Laar, M.C., van den Wildenberg, W.P.M., van Boxtel, G.J.M., & van der Molen, M.W. (2010). Processing of global and selective stop signals: Application of Donders' subtraction method to stop-signal task performance. *Experimental Psychology*, *57*, 149–159. <http://doi.org/10.1027/1618-3169/a000019>
- van den Wildenberg, W.P.M., Burle, B., Vidal, F., van der Molen, M.W., Ridderinkhof, K.R., & Hasbroucq, T. (2010). Mechanisms and dynamics of cortical motor inhibition in the stop-signal paradigm: A TMS study. *Journal of Cognitive Neuroscience*, *22*, 225–39. <http://doi.org/10.1162/jocn.2009.21248>
- van den Wildenberg, W.P.M. & van der Molen, M.W., (2004). Developmental trends in simple and selective inhibition of compatible and incompatible responses. *Journal of Experimental Child Psychology* *87*, 201–220.

- van den Wildenberg, W.P.M., van der Molen, M.W., Logan, G.D. (2002). Reduced response readiness delays stop signal inhibition. *Acta Psychologica 111*, 155–169.
- van der Schoot, M., Licht, R., Horsley, T.M., & Sergeant, J.A. (2000). Inhibitory deficits in reading disability depend on subtype: Guessers but not spellers. *Child Neuropsychology, 6*, 297–312.
- van der Schoot, M., Licht, R., Horsley, T.M., & Sergeant, J.A. (2005). Effects of stop signal modality, stop signal intensity and tracking method on inhibitory performance as determined by use of the stop signal paradigm. *Scandinavian Journal of Psychology, 46*, 331–341. <http://doi.org/10.1111/j.1467-9450.2005.00463.x>
- Van Zandt, T. (2000). How to fit a response time distribution. *Psychonomic Bulletin & Review, 7*, 424–465.
- Van Zandt, T., Colonius, H., & Proctor, R.W. (2000). A comparison of two response-time models applied to perceptual matching. *Psychonomic Bulletin & Review, 7*, 208–256.
- Verbruggen, F., Best, M., Bowditch, W.A., Stevens, T., & McLaren, I.P.L. (2014). The inhibitory control reflex. *Neuropsychologia, 65*, 263–278. <http://doi.org/http://dx.doi.org/10.1016/j.neuropsychologia.2014.08.014>
- Verbruggen, F., Chambers, C.D., & Logan, G.D. (2013). Fictitious inhibitory differences: How skewness and slowing distort the estimation of stopping latencies. *Psychological Science, 24*, 352–362.
- Verbruggen, F., & De Houwer, J. (2007). Do emotional stimuli interfere with response inhibition? Evidence from the stop signal paradigm. *Cognition & Emotion, 21*, 391–403. <http://doi.org/10.1080/02699930600625081>
- Verbruggen, F., Liefoghe, B., Szmalec, A., & Vandierendonck, A. (2005). Inhibiting responses when switching: Does it matter? *Experimental Psychology, 52*, 125–130.
- Verbruggen, F., Liefoghe, B., & Vandierendonck, A. (2004). The interaction between stop signal inhibition and distractor interference in the flanker and Stroop task. *Acta Psychologica 116*, 21–37.
- Verbruggen, F., Liefoghe, B., & Vandierendonck, A. (2006). The effect of interference in the early processing stages on response inhibition in the stop-signal task. *Quarterly Journal of Experimental Psychology 59*, 190–203.
- Verbruggen, F., & Logan, G.D. (2008a). Automatic and controlled response inhibition: Associative learning in the go/no-go and stop-signal paradigms. *Journal of Experimental Psychology: General, 137*, 649–672. <http://doi.org/10.1037/a0013170>

- Verbruggen, F., & Logan, G.D. (2008b). Response inhibition in the stop-signal paradigm. *Trends in Cognitive Sciences*, *12*, 418–424. <http://doi.org/10.1016/j.tics.2008.07.005>
- Verbruggen, F., & Logan, G.D. (2009a). Models of response inhibition in the stop-signal and stop-change paradigms. *Neuroscience & Biobehavioral Reviews*, *33*, 647–661. <http://doi.org/10.1016/j.neubiorev.2008.08.014>
- Verbruggen, F., & Logan, G.D. (2009b). Proactive adjustments of response strategies in the stop-signal paradigm. *Journal of Experimental Psychology: Human Perception and Performance*, *35*, 835–854. <http://doi.org/10.1037/a0012726>
- Verbruggen, F. & Logan, G.D. (2015). Evidence for capacity sharing when stopping. *Cognition*, *142*, 81-95.
- Verbruggen, F., Logan, G.D., & Stevens, M.A. (2008). STOP-IT: Windows executable software for the stop-signal paradigm. *Behavior Research Methods*, *40*, 479–483.
- Verbruggen, F., McLaren, I.P.L., & Chambers, C.D. (2014). Banishing the control homunculi in studies of action control and behavior change. *Perspectives on Psychological Science*, *9*, 497–524. <http://doi.org/10.1177/1745691614526414>
- Verbruggen, F., Schneider, D.W., & Logan, G.D. (2008). How to stop and change a response: The role of goal activation in multitasking. *Journal of Experimental Psychology: Human Perception and Performance*, *34*, 1212–1228. <http://doi.org/10.1037/0096-1523.34.5.1212>
- Verbruggen, F., Stevens, T., & Chambers, C.D. (2014). Proactive and reactive stopping when distracted: An attentional account. *Journal of Experimental Psychology: Human Perception and Performance*, *40*, 1295–1300. <http://doi.org/10.1037/a0036542>
- Vince, M.E. (1948). The intermittency of control movements and the psychological refractory period. *British Journal of Psychology*, *38*, 149-157.
- Wagenmakers, E.-J., Ratcliff, R., Gomez, P., & Iverson, G.J. (2004). Assessing model mimicry using the parametric bootstrap. *Journal of Mathematical Psychology*, *48*, 28-50.
- Whelan, R., Conrod, P.J., Poline, J.-B.B., Lourdusamy, A., Banaschewski, T., Barker, G. J., ..., & the IMAGEN Consortium (2012). Adolescent impulsivity phenotypes characterized by distinct brain networks. *Nature Neuroscience*, *15*, 920–925. <http://doi.org/10.1038/nn.3092>
- Williams, B.R., Ponsesse, J.S., Schachar, R.J., Logan, G.D., Tannock, R. (1999). Development of inhibitory control across the life span. *Developmental Psychology*, *35*, 205–213.