On the Ability to Inhibit Thought and Action: General and Special Theories of an Act of Control

Gordon D. Logan
Vanderbilt University

Trisha Van Zandt
The Ohio State University

Frederick Verbruggen
University of Exeter

Eric-Jan Wagenmakers
University of Amsterdam

Response inhibition is an important act of control in many domains of psychology and neuroscience. It is often studied in a stop-signal task that requires subjects to inhibit an ongoing action in response to a stop signal. Performance in the stop-signal task is understood as a race between a *go process* that underlies the action and a *stop process* that inhibits the action. Responses are inhibited if the stop process finishes before the go process. The finishing time of the stop process is not directly observable; a mathematical model is required to estimate its duration. Logan and Cowan (1984) developed an independent race model that is widely used for this purpose. We present a *general race model* that extends the independent race model to account for the role of choice in go and stop processes, and a *special race model* that assumes each runner is a stochastic accumulator governed by a diffusion process. We apply the models to 2 data sets to test assumptions about selective influence of capacity limitations on drift rates and strategies on thresholds, which are largely confirmed. The model provides estimates of distributions of stop-signal response times, which previous models could not estimate. We discuss implications of viewing cognitive control as the result of a repertoire of acts of control tailored to different tasks and situations.

Keywords: cognitive control, stop-signal, countermanding, race model, diffusion model

Supplemental materials: http://dx.doi.org/10.1037/a0035230.supp

The cognitive system can deploy many acts of control to direct thought and action toward its goals. These acts include shifting attention (Posner, 1980), changing task sets (Logan & Gordon, 2001), resolving and adapting to conflict (Botvinick, Braver, Barch, Carter, & Cohen, 2001; Cohen, Dunbar, & McClelland, 1990), trading speed for accuracy (Forstmann et al., 2008, 2010; Ratcliff, 2006), detecting and preventing errors (Holroyd & Coles, 2002), and inhibiting inappropriate responses (Logan & Cowan,

1984). Theorists often address each act of control separately, focusing on one empirical phenomenon and proposing the control processes and subordinate processes that produce it. We present a theory that allows us to address different acts of control within the same mathematical framework and suggest constraints on the set of acts of control in the cognitive system's repertoire.

We address the acts of control underlying the ability to inhibit inappropriate responses. We focus on the stop-signal paradigm, in which subjects must respond to a *go task* as quickly as they can but inhibit their response to the go task when they hear an occasional stop signal. The stop-signal paradigm is widely used in studies of response inhibition, elucidating the underlying neural structures (e.g., Aron & Poldrack, 2006; Hanes, Patterson, & Schall, 1998), the development and decline of inhibitory ability over the life span (e.g., Huizinga, Dolan, & van der Molen, 2006; Williams, Ponesse, Schachar, Logan, & Tannock, 1999), individual differences in inhibitory ability (e.g., Friedman et al., 2008; Miyake et al., 2000), and the deleterious effects of psychopathology (e.g., Chambers, Garavan, & Bellgrove, 2009; Schachar & Logan, 1990) and neurological disorders (e.g., Aron, Fletcher, Bullmore, Sahakian, & Robbins, 2003; Dimitrov et al., 2003).

The purpose of this article is to propose a theory of response inhibition in the stop-signal paradigm that accounts for choice. Choice is pervasive in stop-signal experiments (for reviews, see Logan, 1994; Logan & Cowan, 1984; Verbruggen & Logan, 2008c), but no current theory of response inhibition accounts for it

(for a review, see Verbruggen & Logan, 2009b). We account for choice by integrating race models with stochastic accumulator models (Usher & McClelland, 2001) to create general and special theories of response inhibition in the stop-signal paradigm. We test the theories by fitting them to data from a new experiment that varied the number of choice alternatives and a previous experiment that manipulated strategies (Verbruggen & Logan, 2009c).

The general theory is a race model. It describes stop-signal performance as a race between a stop process and a go process (Logan & Cowan, 1984), and it describes choice in the go process as a race between alternative responses (Brown & Heathcote, 2005, 2008; Van Zandt, 2000b). It makes minimal assumptions about the underlying processes, predicting relations among response time (RT) distributions and response probabilities that hold for all distributions. It accommodates but does not explain changes in RT with strategies and conditions.

The special theories are stochastic accumulator models embedded in a general race model. The special theories make specific assumptions about the underlying processes: Each runner in the race is a stochastic accumulator that accumulates information to a threshold (Ratcliff & Smith, 2004; Teodorescu & Usher, 2013). The special theories predict the shapes of the RT distributions and explain changes in RT with strategies and conditions as the result of changes in the parameters of stochastic accumulation: the rate of approach to threshold (*drift rate*), the *threshold,* and the time for perceptual and motor processing (*nondecision time*).

The special theories allow us to formulate alternative models that test strong hypotheses about essential properties of subordinate and executive processing. We develop models within each theory that hold parameters constant or vary them between conditions, and we compare their fit to the data to determine which parameters produce which effects. The parameters map directly onto psychological processes, and that allows us to test hypotheses about which effects are due to subordinate processing and which are due to executive processing. This is an advance over previous models of the stop-signal paradigm, which have focused primarily on a single condition or have described but not explained differences between conditions (Boucher, Palmeri, Logan, & Schall, 2007; Logan & Cowan, 1984).

We use the theory to explain two acts of control in the stop-signal task. One act is the stop process that is the main focus of much stop-signal research. It begins with the stop signal and ends with an attempt to inhibit the current response that succeeds or fails. The other act of control modulates the balance between stopping and going, which is becoming a popular topic of research. This act of control occurs before trials and between trials. It appears as proactive slowing of go RT when stop signals become relevant (Verbruggen & Logan, 2009c) or occur more frequently (Bissett & Logan, 2011; Logan, 1981; Ramautar, Kok, & Ridderinkhof, 2004) and as reactive slowing of go RT after a stop signal occurs (Bissett & Logan, 2011; Emeric et al., 2007; Rieger & Gauggel, 1999; Verbruggen, Logan, Liefooghe, & Vandierendonck, 2008).

Our theory explains the stop process as another runner in the race (Logan & Cowan, 1984). We model its duration (stop-signal response time or SSRT), and we model its effects on performance: If it wins the race, the go response is inhibited; if it loses, the go response is executed. Our theory explains modulatory acts of control as adjustments of the parameters of stochastic accumula-

tion (Logan & Gordon, 2001). We model the effects of these acts of control as changes in thresholds, drift rates, or nondecision times in the subordinate processes, but we do not model the time it takes to implement these effects. We hypothesize that experimental manipulations will selectively influence parameters of the stochastic accumulators (cf. Sternberg, 1969). Manipulations of structure, like the number of choice alternatives, the difficulty of perceptual processing, and the load on capacity, should affect drift rate. Manipulations of strategies, such as those that produce proactive slowing, should affect threshold. We test the selective-influence hypothesis by fitting special race models to data from an experiment that manipulates the number of choice alternatives and an experiment that manipulates strategies (Verbruggen & Logan, 2009c).

The special theories allow us to estimate the entire distribution of SSRT. This has not been feasible in previous theories of response inhibition (Colonius, 1990; De Jong, Coles, Logan, & Gratton, 1990; Logan & Cowan, 1984; but see Matzke, Dolan, Logan, Brown, & Wagenmakers, 2013). The distribution of SSRTs may be useful in studies of neuroscience, life-span development, individual differences, psychopathology, and neurological disorders, as RT distributions have provided useful information in these domains (Balota & Yap, 2011; Ratcliff & Smith, 2004; Van Zandt, 2000a).

The Stop-Signal Paradigm

The first stop-signal experiment was published in 1948 by Margaret Vince. A few experiments were reported in the 1960s and 1970s (Lappin & Eriksen, 1966; Ollman, 1973; Slater-Hammel, 1960), but stop-signal research did not begin in earnest until the 1980s, when it was organized around independent race models of the stop and go processes (Logan, 1981; Logan & Cowan, 1984; Osman, Kornblum, & Meyer, 1986). The 1990s saw the first applications to clinical psychology (Schachar & Logan, 1990; Tannock, Schachar, Carr, Chajczyk, & Logan, 1989), developmental psychology (Kramer, Humphrey, Larish, Logan, & Strayer, 1994; Schachar & Logan, 1990), neuroscience (De Jong et al., 1990; Hanes et al., 1998), and individual differences (Logan, Schachar, & Tannock, 1997). Since the turn of the century, stop-signal research has gained momentum, and the stop-signal paradigm is now a popular procedure for the study of response inhibition and cognitive control in cognitive science, clinical science, and neuroscience (see Verbruggen, Chambers, & Logan, 2013; for reviews, see Logan, 1994; Logan & Cowan, 1984; Verbruggen & Logan, 2008b). Stop-signal performance is considered to be an endophenotype for attention-deficit disorder (Schachar et al., 2005), drug addiction (Ersche et al., 2012), and obsessive-compulsive disorder (Chamberlain & Sahakian, 2007).

The stop-signal paradigm requires the deliberate inhibition of a voluntary response. Subjects are engaged in a go task that requires a speeded response. The go task usually involves choice between alternative responses, but some studies have addressed simple RT tasks (Logan, Cowan, & Davis, 1984). Occasionally, a *stop signal* is presented that instructs subjects to withhold their response on that trial. The stop signal is usually a tone, but some studies have used visual (Lappin & Eriksen, 1966; Verbruggen, Aron, Stevens, & Chambers, 2010) or tactile stop signals (Åkerfelt, Colonius, & Diederich, 2006). The most important independent variable is the

delay between the onset of the stimulus for the go task and the onset of the stop signal (*stop-signal delay*).

When given a stop signal, subjects either inhibit their response to the go task, producing a *signal-inhibit* trial, or fail to inhibit their response, producing a *signal-respond* trial. The probability of inhibiting the response—*P*(inhibit)— depends on stop-signal delay. It decreases as stop-signal delay increases. Many researchers plot the complementary probability of responding given a stop signal—*P*(respond|signal)—which increases as stop-signal delay increases. The plot of either probability against stop-signal delay is called the *inhibition function*. Typical inhibition functions are plotted in Figure 1 (for a discussion of inhibition functions, see Logan & Cowan, 1984; Verbruggen & Logan, 2009b).

The inhibition function is determined by stop-signal delay, but it also depends strongly on RT in the go task; the probability of responding given a stop signal is lower the longer the go RT (Logan, 1981; Logan & Cowan, 1984). Differences in go RT shift the inhibition function to the left or right along the stop-signal
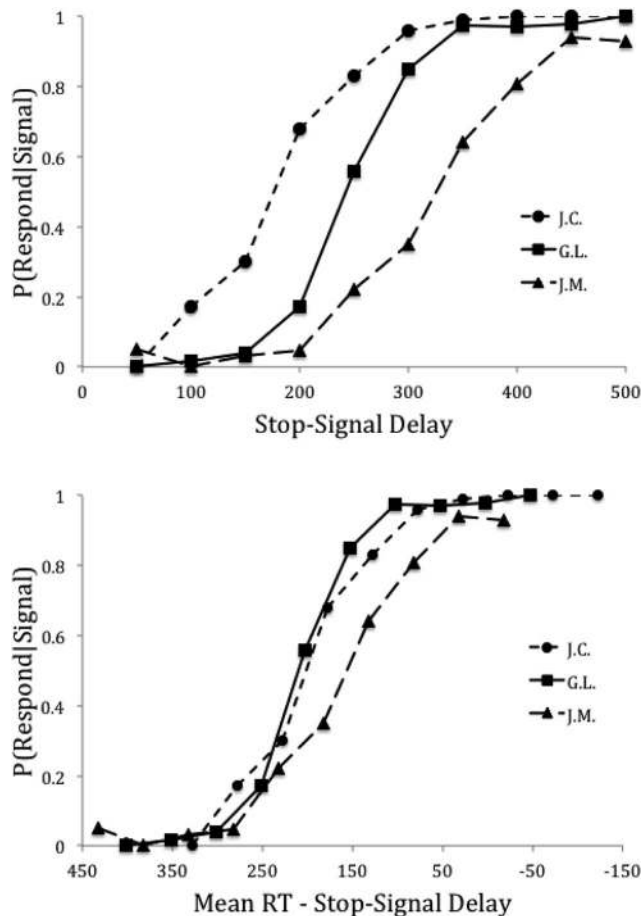

delay axis. Often, shifts due to go RT differences between conditions, strategies, tasks, and subjects can be compensated for precisely by replotting the inhibition function as a function of the difference between go RT and stop-signal delay (Logan, 1981; Logan & Cowan, 1984; cf. Salinas & Stanford, 2013). Researchers have also adjusted stop-signal delay with a tracking procedure to produce equivalent probabilities of inhibition in different conditions, strategies, tasks, and subjects, essentially aligning the inhibition functions (e.g., Logan et al., 1997; Osman et al., 1986).

The second panel of Figure 1 shows typical inhibition functions plotted against RT minus stop-signal delay. Note that the functions for J.C. and G.L. align better than the function for J.M. This happened because J.M. had greater variability in go RT than J.C. or G.L. Transformations that take go variability into account produce better alignment (see Logan & Cowan, 1984, Figure 3).

The alignment of inhibition functions suggests that it is profitable to think of response inhibition as an act of control with a specific latency. The difference between go RT and stop-signal delay reflects the time that is available to execute the act of control, that is, to detect the stop signal and heed it before inappropriately executing the go response. The alignment also motivates race models of stop-signal performance: The response can be inhibited only if the act of control finishes before the go response (Logan, 1981).

Another important dependent variable is go RT on signal-respond trials. It is usually faster than go RT on trials with no stop signal and faster for shorter stop-signal delays than for longer ones. When plotted as a cumulative distribution, the minimum go RTs are about the same for signal-respond RTs from different stop-signal delays and for go RTs from no-stop-signal trials. The distributions fan out at the higher quantiles, rising more slowly the longer the stop-signal delay (Osman et al., 1986). Example distributions of go RTs from signal-respond and no-stop-signal trials are plotted in Figure 2.

The relation between signal-respond RT and no-stop-signal RT also suggests that the act of control underlying response inhibition has a specific latency. Signal-respond RTs are the go responses that are faster than the act of control that underlies response inhibition. The shorter the stop-signal delay, the sooner the act of control finishes, so the faster the go RT has to be in order to finish before it. This relationship and the relationship between inhibition functions and go RT motivated the development of race models of stop-signal performance.

The Independent Race Model

Logan and Cowan (1984) proposed an independent race model to account for stop-signal performance. The model assumes that a stop process, initiated by the stop signal, races in parallel against a go process, initiated by the go stimulus, and performance is determined by the process that wins the race. If the stop process wins, the go response is inhibited; if the go process wins, the go response escapes inhibition. The finishing times of the stop process and the go process are assumed to be independent random variables, whose probability density functions are $f_{stop}(t)$ and $f_{go}(t)$, respectively. We assume $f_{stop}(t)$ and $f_{go}(t)$ are continuous and 0 for all $t < 0$. The model predicts the probability of responding, P_r, given a stop signal at delay t_d as

Figure 1. Top panel: Inhibition functions for three subjects plotted as a function of stop-signal delay. Bottom panel: Inhibition functions for the same three subjects replotted as a function of mean go response time minus stop-signal delay. Data are taken from "On the Ability to Inhibit Thought and Action: A Theory of an Act of Control," by G. D. Logan and W. B. Cowan, 1984, *Psychological Review, 91,* p. 298. Copyright 1984 by the American Psychological Association.

Perhaps the most important contribution of the independent race model was to provide methods for estimating the time it takes to inhibit a response (i.e., SSRT). Estimates of SSRT have served as measures of cognitive control in studies of cognition, life-span development, individual differences, psychopathology, and neuro-pathology (for reviews, see Logan, 1994; Verbruggen & Logan, 2008b). This is an important contribution because SSRT is not directly observable.

The *mean method* involves using the tracking procedure to tie the race between stopping and going, so each wins 50% of the time (Logan et al., 1997; Osman et al., 1986). When the race is tied, mean go RT = mean stop-signal delay + SSRT. SSRT can be calculated by subtracting mean stop-signal delay from mean go RT on no-stop-signal trials. This method can estimate SSRT accurately, but it is susceptible to distortion from skew and strategic slowing, and so it should be interpreted with caution (Verbruggen et al., 2013).

The *integration method* is more general than the mean method. It can be used no matter how stop-signal delays are set, and it is more robust to skew and strategic slowing (Verbruggen et al., 2013). It assumes SSRT is a constant, so any go RTs that finish before stop-signal delay + SSRT will be executed, and any go RTs that finish after it will be inhibited. Thus, the probability of responding on a stop-signal trial equals the proportion of the go RT distribution that is faster than stop-signal delay + SSRT:

$$P_r(t_d) = \int_0^{SSRT+t_d} f_{go}(t)dt, \qquad (3)$$

where t_d is stop-signal delay. The integration method inverts the relationship in Equation 3, using the go RT distribution, the probability of responding given a stop signal, and stop-signal delay to identify SSRT. The go RT distribution is integrated until the integral equals the probability of responding given a stop signal. At that point, $t = t_d + $ SSRT, and SSRT is estimated by subtracting t_d from t (see Logan, 1994; Logan & Cowan, 1984). In practice, calculating SSRT with the integration method involves rank-ordering the N go RTs in the no-stop-signal distribution, finding the Mth go RT, where $M = N \times P$(respond|signal), and subtracting stop-signal delay from the Mth go RT (see Logan, 1994).

The *distribution method* calculates the unobserved distribution of SSRTs from the observed distributions of no-stop-signal RTs and signal-respond RTs by rearranging Equation 2:

$$F_{stop}(t - t_d) = 1 - f_{sr}(t|t_d)P_r(t_d)/f_{go}(t) \qquad (4)$$

(Colonius, 1990; De Jong et al., 1990). Estimating SSRT distributions with Equation 4 is impractical because it depends heavily on accurate estimates of the tails of the distributions, which require a lot of observations (Matzke et al., 2013). Our model allows us to estimate the SSRT distribution more efficiently.

Independence Assumptions

The independent race model assumes two kinds of independence: stochastic independence and context independence (Colonius, 1990; Logan & Cowan, 1984). *Stochastic independence* means that the finishing times of the stop process and the go process are independent on a given trial. More precisely, it means that
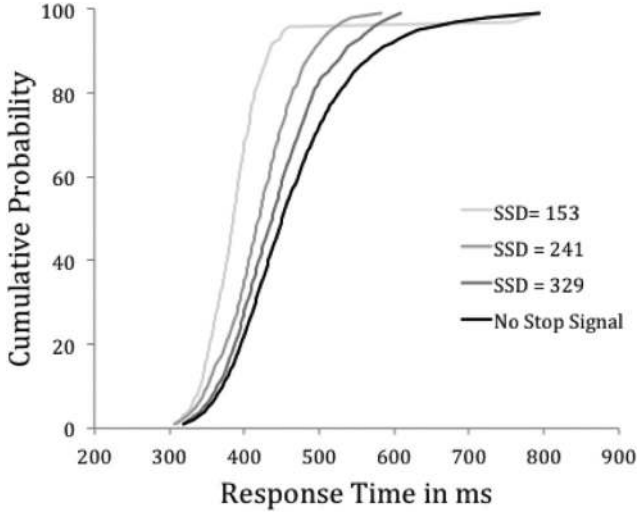

Figure 2. Quantile average response time distributions for no-stop-signal trials and signal-respond trials with stop-signal delays (SSDs) of 153, 241, and 329 ms. Data are taken from the two-choice condition of the multiple-choice experiment reported later in the article.

$$P_r(t_d) = \int_0^\infty f_{go}(t)(1 - F_{stop}(t - t_d))dt, \qquad (1)$$

where $F_{stop}(t - t_d)$ is the cumulative distribution function of finishing times for the stop process at delay t_d.

Equation 1 predicts the inhibition function through the term $1 - F_{stop}(t - t_d)$. Increasing stop-signal delay decreases $F_{stop}(t - t_d)$, which results in a higher probability of responding. Equation 1 also explains the effect of go RT on the inhibition function. Increasing mean go RT will decrease the probability that the go process will beat the stop process, so the inhibition function will shift to the right, as observed (see Figure 1). The model also predicts that changes in stop-signal delay can compensate for changes in go RT to align inhibition functions from different conditions, strategies, tasks, and subjects (Logan, 1981; Logan & Cowan, 1984). This prediction justifies the common practice of adjusting stop-signal delay with a tracking procedure to produce a desired probability of inhibition (Logan et al., 1997; Osman et al., 1986).

The independent race model also provides the distribution of signal-respond RTs at a given stop-signal delay, $f_{sr}(t|t_d)$, as

$$f_{sr}(t \mid t_d) = f_{go}(t)\left(1 - F_{stop}(t - t_d)\right)/P_r(t_d). \qquad (2)$$

The model explains why signal-respond and no-stop-signal RT distributions share a common minimum and fan out at higher quantiles with a steeper rise for shorter stop-signal delays. The term $1 - F_{stop}(t - t_d)$ decreases monotonically as t increases and acts as a filter that compresses the upper tail of the go RT distribution, $f_{go}(t)$. The longer the stop-signal delay, the less the go RT distribution is compressed at a given value of t, so the shallower the rise of the cumulative distribution. In the limit, the signal-respond RT distribution will approach the no-stop-signal RT distribution.

$$P(T_{go} < t_{go} \cap T_{stop} < t_{stop}) = P(T_{go} < t_{go}) \cdot P(T_{stop} < t_{stop}) \quad (5)$$

for all t_{go} and t_{stop}. *Context independence* means that the distribution of finishing times for the go process is the same whether or not a stop signal is presented. More precisely, it means that

$$P(T_{go} < t | \text{no stop signal}) = P(T_{go} < t | t_d) \quad (6)$$

for all t and t_d.

The independent race model does not assume *functional independence* of the stop and go processes (Ashby & Townsend, 1986). Functional independence means that factors that affect the distribution of finishing times for the go process do not affect the distribution of finishing times for the stop process and vice versa. More precisely, functional independence means that

$$[f_{go}(t|A) \neq f_{go}(t|B)] \wedge [f_{stop}(t|A) = f_{stop}(t|B)] = True \quad (7a)$$

or

$$[f_{go}(t|C) = f_{go}(t|D)] \wedge [f_{stop}(t|C) \neq f_{stop}(t|D)] = True$$
$$(7b)$$

or both Equations 7a and 7b are true (where $\wedge$ denotes logical conjunction). *A* and *B* are different conditions that affect the go task but do not affect the stop task, and *C* and *D* are different conditions that affect the stop task but not the go task. It is important to note that violations of functional independence do not imply violations of stochastic or context independence. Equations 5 and 6 could hold when Equation 7 is violated.

This article explores the functional independence of stop and go processes by examining the hypothesis that they share capacity. Capacity sharing is a common explanation of dual-task interference (Kahneman, 1973; Pashler, 1994; Posner & Boies, 1971), and several studies have asked whether the stop process and the go process share capacity (Logan, 1981; Logan & Burkell, 1986; Yamaguchi, Logan, & Bissett, 2012). We ask the question more rigorously, using parameters of our stop-signal models to measure capacity and assess functional independence (Townsend & Ashby, 1983).

Benefits and Costs of Generality

The independent race model is very general. This generality allows the race model to apply to any stop-signal task in any response modality, including keypresses (Logan, 1981), hand squeezes (De Jong et al., 1990), wrist and arm movements (Brunamonti, Ferraina, & Paré, 2012), eye movements (Logan & Irwin, 2000), typewriting (Logan, 1982), and speech (Xue, Aron, & Poldrack, 2008), or in any subject population, including children (Schachar & Logan, 1990), the elderly (Kramer et al., 1994), psychiatric and neurological patients (Aron et al., 2003; Thakkar, Schall, Boucher, Logan, & Park, 2011), monkeys (Hanes et al., 1998), and rodents (Eagle & Robbins, 2003). No parameters have to be estimated.

The independent race model does not specify the mechanisms that produce the finishing-time distributions, so it cannot explain the effects of structural and strategic manipulations on stopping performance and go RT. We address this limitation by proposing special race models that address such effects. The independent race model does not specify the mechanism that inhibits the response

after the stop process wins the race, so it cannot explain recent investigations of the neural interactions that cause stopping (Aron & Poldrack, 2006; Hanes et al., 1998; Paré & Hanes, 2003). We do not address this limitation in this article. The interaction occurs in a stage subsequent to the race, and the duration of that stage is very brief (Boucher et al., 2007). Moreover, our current modeling of the interactive stage suggests several viable mechanisms that are difficult to distinguish in behavioral and neural data, so we defer questions about the nature of the interaction until we learn more about it.

Alternative Models

The independent race model went unchallenged for 20 years and is still used to estimate SSRT in virtually every published stop-signal study. In the last few years, several alternative models have been proposed. All of these models assume a race between stop and go processes, so they predict inhibition functions and signal-respond RTs, like the independent race model. The alternative models focus more directly on mechanism, asking what the stop process does to stop the go process.

Boucher et al. (2007) proposed an *interactive race* model, in which the stop process has two stages: an *afferent* stage that detects the stop signal and apprehends its significance and an *interactive* stage that inhibits the go response. The go process is modeled as a single diffusion to a threshold, and the interactive stage of the stop process is modeled as a single diffusion that inhibits the growth of activation in the go process. Responses are inhibited if the interactive stage prevents the go process from reaching threshold (cf. Salinas & Stanford, 2013). In fits of the model to data from two monkeys, the afferent stage occupied most of SSRT. The interactive stage was very brief, and the inhibition from the stop process on the go process was very strong. Thus, the race was independent for most of its duration, and the interaction was brief and potent. Lo, Boucher, Paré, Schall, and Wang (2009) developed a spiking-neuron version of the interactive race model, and Wong-Lin, Eckhoff, Holmes, and Cohen (2010) developed a version that explains the optimization of reward rate.

These models are important because they specify the underlying mechanisms and connect mathematical models to underlying physiology (also see Forstmann, Wagenmakers, Eichele, Brown, & Serences, 2011; Gold & Shadlen, 2007; Purcell et al., 2010). They all assume that responses are stopped by inhibiting go activation (but see Salinas & Stanford, 2013). They all find that the afferent stage of the stop process is much longer than the interactive stage, so they all approximate the independent race model.

All of these models share a common limitation: They do not deal with choice. The go task is represented by a single accumulator that is guaranteed to reach threshold on every trial. There are no errors of choice. This is an important limitation because go tasks that involve choice RT are pervasive in the stop-signal literature (Logan, 1994; Logan & Cowan, 1984; Verbruggen & Logan, 2008b), and models of RT that address choice are pervasive in the mathematical modeling literature (Logan, 2004; Ratcliff & Smith, 2004; Teodorescu & Usher, 2013).

The goal of this article is to develop and test general and special independent race models that deal with choice in the stop-signal paradigm. The general model assumes that choice involves a race between all possible responses, including the go alternatives and

the stop response. The special models assume that each runner is a single diffusion process, whose duration depends on drift rate, threshold, and nondecision time parameters. We test special models by fitting them to data from a new experiment that manipulated the number of choices in the go task and a previous experiment that manipulated strategic slowing in anticipation of stop signals (Verbruggen & Logan, 2009c).

General Independent Race Model

The general independent race model extends the original independent race model (Logan & Cowan, 1984), describing the go task as a race between alternative responses. The general model assumes a race between a set A of possible responses that includes the stop response as well as each of the possible responses in the go task. It assumes stochastic independence for all runners (see Equation 5) and context independence for the stop process (see Equation 6). Functional independence (see Equation 7) is an empirical question, which we address through our tests of special race models. The general independent race model includes the original independent race model as a special case, in which the set A contains just two members: the stop process and the go process. The general independent race model addresses distributions of finishing times without specifying the form of the distributions or the processes that generate them. It provides a cognitive architecture in which more specific models can be articulated.

Our decision to represent the go process as a race between alternative responses represents a stronger commitment to cognitive architecture than the original independent race model made, and consequently, it is more controversial. On the one hand, several successful models of RT assume a race between independent runners (e.g., Brown & Heathcote, 2005, 2008; Logan, 1988; P. L. Smith & Van Zandt, 2000; Van Zandt, 2000b; Van Zandt, Colonius, & Proctor, 2000). On the other hand, other successful models assume competition between alternative responses, including random walk (e.g., Nosofsky & Palmeri, 1997), diffusion (e.g., Ratcliff, Van Zandt, & McKoon, 1999), and competitive leaky accumulator models (Usher & McClelland, 2001). In direct comparisons, some specific race models have not accounted for behavioral data as well as some specific competitive models (Ratcliff & Smith, 2004; Teodorescu & Usher, 2013), but in other contexts, race models sometimes do a better job of accounting for behavioral (Leite & Ratcliff, 2010) and physiological data (Ratcliff, Cherian, & Segraves, 2003; but see Purcell et al., 2010). An important virtue of race models for our present purposes is their mathematical simplicity and the transparent way they allow us to formulate and test mathematical models that assume specific forms of stochastic accumulation.

The general independent race model assumes that each runner in the race is a stochastic accumulator (Ratcliff & Smith, 2004; Teodorescu & Usher, 2013). The model assumes that stochastic accumulation proceeds independently for each response in the response set A and that the chosen response and the time at which it is chosen are determined by the accumulator that reaches its threshold first. Let $f_i(t)$ be the probability density function of the time t that accumulator i reaches its threshold. Let $f_{stop}(t - t_d)$ be the probability density function of the time t that the stop accumulator reaches its threshold given that stop-signal delay is t_d. Both $f_i(t)$ and $f_{stop}(t - t_d)$ are zero for values of their arguments

that are less than zero, and both depend on the specific stochastic accumulator model that is proposed (diffusion, Poisson counter, etc.). The general race model does not commit to any particular stochastic accumulator model.

Given these assumptions, the probability that go response i will occur is

$$P(\text{resp } i) = \int_0^\infty f_i(u) \prod_{j \in A, j \neq i} (1 - F_j(u))(1 - F_{stop}(u - t_d)) \, du, \quad (8a)$$

where $F_j(t)$ is the cumulative distribution function for go response j. On no-stop-signal trials, $t_d = \infty$, so the stop process has no chance of winning the race. On stop-signal trials, $t_d \ll \infty$, and the probability that the stop process wins the race is

$$P(\text{stop}) = \int_0^\infty f_{stop}(u - t_d) \prod_{i \in A} (1 - F_i(u)) \, du. \quad (8b)$$

Thus, for the general race model, the inhibition function is

$$P_r(t_d) = 1 - P(\text{stop}) = 1 - \int_0^\infty f_{stop}(u - t_d) \prod_{i \in A} (1 - F_i(u)) du. \quad (9)$$

The general race model assumptions also allow us to specify the joint probability density function of RT given response i, which is

$$f(t|i) = \left[f_i(t)(1 - F_{stop}(t - t_d)) \prod_{j \in A, j \neq i} (1 - F_j(t)) \right] / P_r(t_d). \quad (10)$$

If there is no stop signal, then $t_d = \infty$, and Equation 10 gives the distribution of RTs for response i. If there is a stop signal, then $t_d \ll \infty$, and Equation 10 gives the distribution of signal-respond RTs for response i. Signal-respond RTs will necessarily be faster than no-stop-signal RTs because the term $1 - F_{stop}(t - t_d)$ will compress the upper tail of the go distribution. The cumulative distribution of SSRTs can be calculated with the Colonius-De Jong method (see Equation 4), using

$$f_{sr}(t) = \left[\sum_{i \in A, j \neq stop} f_i(t) \prod_{j \in A, j \neq i, j \neq stop} (1 - F_j(t)) \right.$$
$$\left. \times (1 - F_{stop}(t - t_d)) \right] / P_r(t_d) \quad (11)$$

for the distribution $f_{sr}(t)$ of signal-respond RTs and

$$f_{go}(t) = \sum_{i \in A, i \neq stop} f_i(t) \prod_{j \in A, j \neq i, j \neq stop} (1 - F_j(t)) \quad (12)$$

for the distribution $f_{go}(t)$ of go RTs.

Special Independent Race Model:
The Diffusion Race Model

We developed special independent race models that describe the racing processes as stochastic accumulators. Each model specifies the finishing-time distribution for each runner in the race in terms of three parameters that capture most of the important effects in the RT literature: rate, threshold, and nondecision time (Ratcliff & Smith, 2004). We use the best fitting parameter values to assess the nature of capacity limitations and strategies in the stop-signal task. Rates address capacity limitations, and thresholds address strategies.

We investigated three special independent race models: a diffusion model (Ratcliff et al., 1999), a Poisson counter model (Van Zandt et al., 2000), and the linear ballistic accumulator model (Brown & Heathcote, 2008). We fit them to the multiple-choice stop-signal task described below. All three models fit the data well, and the model fits led to the same conclusions regarding processing capacity and number of choices (see the online supplemental materials). This suggests that the constraints in the general independent race model were doing most of the work. We chose to focus our investigation on the diffusion model. It fit better than the other models in the majority of cases, and it has been investigated more extensively.

Diffusion models are popular models of choice RT (e.g., Ratcliff & McKoon, 2008; Ratcliff et al., 1999), accounting for RT and error data in a wide variety of tasks from attention (P. L. Smith & Ratcliff, 2009) and intelligence (van Ravenzwaaij, Brown, & Wagenmakers, 2011) to lexical decision (Wagenmakers, Ratcliff, Gomez, & McKoon, 2008) and recognition memory (Ratcliff, 1978). Our *diffusion race model* assumes a race between N independent diffusion processes, each of which has a single boundary (Usher, Olami, & McClelland, 2002). The finishing-time distribution for each runner is simply the Wald distribution (see below). The finishing-time distribution for the winner of the race is not the Wald but instead is the distribution of the minima of the Wald distributions for all of the runners in the race.

The diffusion race model assumes that each stochastic accumulator is a Wiener diffusion process with a drift rate ξ, a starting point at 0, and a threshold (absorbing boundary) at z. We assume a drift coefficient equal to 1. The accumulator for the stop process does not begin until stop-signal delay expires. Under these assumptions, the finishing-time distributions $f_i(t)$ are given by inverse normal (Wald) densities with parameters determined by the drift rate and threshold. Thus, for the go process,

$$f_i(t) = z(2\pi t^3)^{-\frac{1}{2}} \exp[-\frac{1}{2t}(\xi t - z)^2], \qquad (13)$$

and for the stop process,

$$f_{stop}(t) = z(2\pi(t - t_d)^3)^{-\frac{1}{2}} \exp[-\frac{1}{2(t - t_d)}(\xi(t - t_d) - z)^2], \qquad (14)$$

if $t > t_d$ and 0 otherwise.

The model expressed in Equations 13 and 14 assumes no variability in threshold across trials, although threshold variability is important in accounting for fast error RT distributions (Ratcliff & Smith, 2004). We investigated a diffusion race model with threshold variability, to capture fast error RTs. We assumed threshold was a uniform random variable ranging from $z - a$ to $z + a$, with a mean of z and a variance of $a^2/3$. The finishing time of a runner unconditioned over the variable threshold z is found by computing

$$g_i(t|z, \xi) = (2a)^{-1} \int_{z-a}^{z+a} f_i(t|z, x_i) \, dz, \qquad (15)$$

where $(2a)^{-1}$ is the probability density function of the uniform threshold. The probability density function for Equation 15 can be computing analytically. Noting that $\varphi(x)$ and $\Phi(x)$ are the density and cumulative distribution functions of the standard normal dis-

tribution, respectively, and letting $\alpha = -(z - a - t\xi)/\sqrt{t}$ and $\beta = (z + a - t\xi)/\sqrt{t}$, then

$$g_i(t|z, \xi, a) = (2a)^{-1}[\phi(\alpha) - \phi(\beta) - \xi(\Phi(\alpha) - \Phi(\beta))] \qquad (16)$$

for $\xi > 0$ and $a > 0$. If $a = 0$, then $g_i(t|z,\xi) = f_i(t)$ (see Equation 13) or $f_{stop}(t)$ (see Equation 14). If $\xi = 0$, then

$$g_i(t|z, a) = (2a)^{-1}[\phi(\alpha) - \phi(\beta)]. \qquad (17)$$

We substituted Equations 16 and 17 for the generic distributions in Equations 8 and 10 to generate likelihood functions to fit the diffusion race model to the data. We used the best fitting parameters to estimate the distribution of SSRTs. This parametric method for estimating SSRT distributions is less susceptible to noise in the tails of the distributions than the nonparametric methods of Colonius (1990) and De Jong et al. (1990).

Our model fits yield estimates of drift rate, threshold, and nondecision time parameters in each condition of the experiments we fit. We assume that thresholds are determined mostly by strategic factors, like expectancies of events and rewards (Ratcliff, 2006; Ratcliff & Smith, 2004). We assume that drift rates are determined partly by structural factors, like capacity limitations, the quality of stimulus information, and the quality of memory representations (Nosofsky, Little, Donkin, & Fific, 2011; Ratcliff et al., 1999), and partly by strategic factors, like division of attention among stimuli (Logan, 1996; Logan & Gordon, 2001; P. L. Smith & Ratcliff, 2009) or stimulus dimensions (Logan & Gordon, 2001; Nosofsky & Palmeri, 1997). When there is no competition for attention, we predict selective influence of experimental manipulations on model parameters: Structural factors should affect drift rates, and strategic factors should affect thresholds. There is usually no competition for attention in the stop-signal paradigm. The go stimulus is presented by itself without any conflicting information from irrelevant distractors or stimulus dimensions to filter out. Thus, the predicted selective influence should be observed.

Capacity Limitations in Stop and Go Processes

The concept of processing capacity has had a long history in cognitive psychology. From Posner and Boies (1971) and Kahneman (1973) onward, researchers have proposed that central processes share capacity, such that one process performs less effectively when a concurrent process is active. There are many demonstrations of dual-task interference in the literature that are consistent with this proposal (for a review, see Pashler, 1994). Stop and go processes do not seem to share capacity in this way. SSRT is as fast as simple RT in many experiments and does not seem to be affected much by the demands of the go task (e.g., Logan, 1981; Logan & Burkell, 1986). Yamaguchi et al. (2012) measured SSRT in a dual-task experiment, in which subjects had to stop one of two go tasks. They found that SSRT was unaffected by the temporal overlap of the two tasks and no different in single- and dual-task conditions.

Previous stop-signal studies used estimates of mean SSRT to test the hypothesis that stop and go processes share capacity. Here, we test the hypothesis more rigorously, using the diffusion race

model to measure capacity in terms of parameters of the underlying stochastic accumulation processes.

Modeling Capacity Limitations

Townsend and colleagues formalized the concept of processing capacity as a measure of the rate of processing, distinguishing between unlimited, limited, and fixed capacity (Townsend & Ashby, 1983). Processing capacity for an individual process, like the ith runner in a race, can be measured as the rate v_i at which the process operates, and processing capacity for a set of N processes, like a processing stage or a set of runners in a race, can be measured as the sum of the rates of the component processes, $\Sigma_{i=1}^N v_i$. Unlimited capacity, limited capacity, and fixed capacity are defined in terms of the rates for individual processes and the sum of the rates over all processes.

A process has *unlimited capacity* if its rate is unchanged when another process enters the race. Thus, the rate of processing for the ith process is the same whether there are N or $N + 1$ runners. That is,

$$v_{i|N} = v_{i|N+1}. \tag{18}$$

A set of processes has unlimited capacity if the sum of the rates of the components increases without limit as more components are added to the race. Thus,

$$\sum_{i=1}^N v_i < \sum_{i=1}^N v_i + v_{i|N+1}. \tag{19}$$

A set of processes has *fixed* capacity if the sum of the rates of the runners is fixed at a constant value, C, regardless of the number of runners. Thus,

$$\sum_{i=1}^N v_i = \sum_{i=1}^{N+1} v_i = C. \tag{20}$$

A process has fixed capacity if its rate decreases when another runner is added to the race but the sum of the rates for all the runners remains the same. If capacity is divided equally among all runners in the race, then

$$v_{i|N} = \frac{C}{N} > v_{i|N+1} = \frac{C}{N+1}. \tag{21}$$

Limited-capacity processes fall between unlimited- and fixed-capacity models. The rate of processing for an individual process decreases as more runners enter the race, but the decrease is not as great as it would be if capacity were fixed. The sum of the rates of processing over all runners in the race increases as more runners enter the race, but the increase is not as great as it would be if capacity were fixed (also see Bundesen, 1990; Eidels, Donkin, Brown, & Heathcote, 2010). Since RT depends on processing rates, fixed capacity and limited capacity are essentially violations of context independence.

Capacity Limitations in the Diffusion Race Model

Equations 18–21 apply to situations in which we can identify the processing rate associated with a particular stage of processing. This is not always possible. Townsend and colleagues (Townsend & Altieri, 2012; Townsend, Houpt, & Silbert, 2012; Townsend & Wenger, 2004; Wenger & Townsend, 2000) developed techniques

for assessing capacity limitations from RT distributions, which reflect the sum of the durations of all stages of processing. We do not need to use such general techniques. Our diffusion race model allows us to assess the rate of processing in the perceptual and conceptual stages of the stop and go processes, which are the processes whose capacity limitations are at issue, so we can apply Equations 18–21 to estimated rate parameters from fits of the models to data. Changes in the rate parameter with number of choices tell us whether capacity is unlimited, limited, or fixed. These assessments of capacity limitations allow us to assess the functional independence of stop and go processes.

Capacity Limitations in Multiple-Choice RT Tasks

It has been known since the 19th century that RT increases monotonically with the number of alternative responses (Merkel, 1885). The increase is important because it means that RT depends not only on the stimulus that is actually presented but also on the set of alternative stimuli that could have been presented (see Garner, 1962). The increase is linear with the logarithm of the number of choice alternatives, which led Hick (1952) and Hyman (1953) to formulate a law that describes this increase, couched in terms of information theory. They interpreted the slope of the linear increase with the logarithm of the number of choice alternatives as the maximum rate at which humans could process information—the capacity for processing information. The link to capacity is supported by converging operations: Dual-task interference increases with the number of alternative responses in a choice RT task (Logan, 1979; M. C. Smith, 1969; Van Selst & Jolicoeur, 1997).

The idea that multiple-choice RT reflects capacity limitations is supported by modeling: Schneider and Anderson (2011) accounted for multiple-choice RT in terms of interference from memory retrieval, which increased with the number of alternatives and lowered the rate of processing for the chosen alternative. This is consistent with fixed capacity. Leite and Ratcliff (2010) fitted a large family of stochastic accumulator models to multiple-choice RTs and found that the models that fit the best allowed processing rate for the chosen alternative to decrease as the number of alternatives increased, consistent with limited capacity. However, the best fitting models also allowed nondecision time and response threshold to vary with number of alternatives. Usher and McClelland (2001) accounted for multiple-choice RT by holding input constant for the chosen alternative and varying the number of competing alternatives, which reduced the effective processing rate for the chosen alternative (also see Bogacz, Usher, Zhang, & McClelland, 2007). This is consistent with limited capacity. Usher et al. (2002) showed that multiple-choice RT could be accounted for by changes in response threshold with no changes in processing rate. This is consistent with unlimited capacity.

Our interpretation of rate parameters as measures of processing capacity allows a precise test of the hypothesis that stop and go processes share capacity. If they do, then the stop process is another runner in the same race that competes for capacity with the runners for each go response. The rate parameter for the stop process should decrease with the number of alternative responses, just as the rate parameters for the go processes do. If the stop process does not share capacity with the go process, then the rate of stop processing should not vary with the number of alternative responses.

The hypothesis that stop and go processes share capacity should be distinguished from the hypothesis that the stop process is limited in capacity. The stop process may have its own capacity limitations even if it does not share them with the go process. The stop process rate parameters must depend on the discriminability and intensity of the stop signal (Cavina-Pratesi, Bricolo, Prior, & Marzi, 2001; Salinas & Stanford, 2013), and that could be interpreted as a capacity limitation (see Bundesen, 1990; Logan, 2002).

To evaluate the diffusion race model and test the hypothesis that stop and go processes share capacity, we conducted an experiment in which six subjects each performed a multiple-choice RT task combined with a stop-signal task for 12 sessions. The multiple-choice task required subjects to identify a single visually presented five-letter word by pressing a key on a computer keyboard. Each session, subjects performed three blocks of 240 trials, one with two choice alternatives, one with four choice alternatives, and one with six choice alternatives, for a total of 8,640 trials per subject. The stop signal was a tone that was presented on 25% of the trials at stop-signal delays that were set separately for each subject and each choice condition to correspond to the 15th, 35th, 55th, 75th, and 95th percentiles of the subject's go RT distribution for that condition. The stop-signal delays were based on a practice block with no stop signals in the first session and remained the same throughout all 12 sessions. New words were used each session to keep go RT relatively constant (Logan, 1979). Further details of the procedure are presented in Appendix A.

We manipulated the number of choice alternatives to vary the load on capacity in the go task. If the stop task shares capacity with the go task, then SSRT should increase as the number of choice alternatives increases. If the stop task does not share capacity with the go task, then SSRT should not vary as the number of choice alternatives increases. We tested this hypothesis more rigorously by fitting the diffusion race model to the data.

If the number of choice alternatives affects the load on capacity, then the rate parameters for the go process in the diffusion race model should decrease as the number of choice alternatives increases. This would imply that the go process has limited capacity.

If the rate parameters decrease such that their sum remains constant over the number of choice alternatives (see Equations 20–21), then the go process has fixed capacity.

If the go process has limited or fixed capacity, we can ask whether it shares capacity with the stop process. If the go process and the stop process share capacity, then the rate parameters for the stop process should decrease as the number of choice alternatives in the go task increases. If the go process and the stop process share a fixed capacity, then the sum of the rate parameters for the stop process and the go processes should remain constant over the number of choice alternatives (see Equations 20–21).

Alternatively, if the go process and the stop process do not share capacity, then the rate parameter for the stop process should not be affected by the number of choice alternatives in the go task, and the sum of the rate parameters for the stop process and the go processes need not be constant over choice alternatives.

Results: Behavioral Data

Mean RTs for correct responses, collapsed across subjects and sessions, are plotted as a function of number of choice alternatives in Figure 3. Mean no-stop-signal RT increased with number of
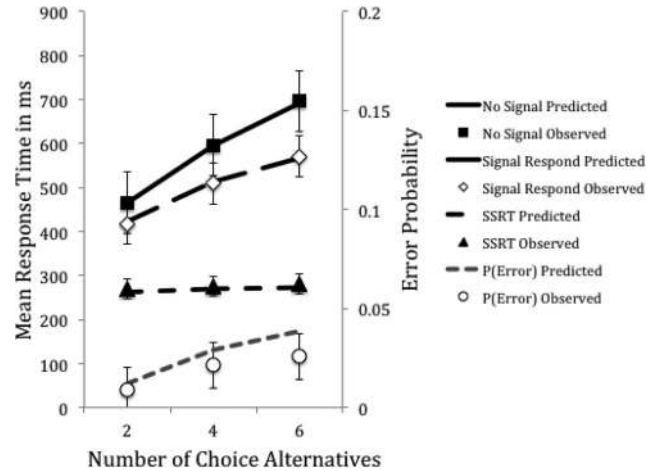

Figure 3. Mean observed (points) and predicted (lines) go response time, signal-respond response time, stop-signal response time (SSRT), and error rate—*P*(Error)—for diffusion race model as a function of number of choice alternatives in the multiple-choice experiment. Error bars are 95% confidence intervals.

choice alternatives, as is typical in multiple-choice tasks, increasing by 132 ms from two to four alternatives and by 101 ms from four to six alternatives (Hick, 1952; Hyman, 1953). Error rate increased with the number of alternatives as well (Brown, Steyvers, & Wagenmakers, 2009). These effects were stable across sessions (see Appendix A) because we introduced a new set of words each session to reduce item-specific learning (Logan, 1988), and they were stable within sessions because we include 72 practice trials before collecting data. These effects were stable across subjects: Go RT increased with number of alternatives for each subject. The increase in go RT with number of alternatives is important because it suggests that the go process has limited or fixed capacity. Demonstrating a limited- or fixed-capacity go process is the first step in asking whether the stop process shares capacity with the go process. However, increases in go RT with the number of alternatives can also occur if the go process is unlimited in capacity (i.e., by increasing threshold; see Usher et al., 2002), so we need to analyze the processing rates in the underlying stochastic accumulators to reach firm conclusions.

The data from stop-signal trials were typical of stop-signal experiments. Inhibition functions across subjects and sessions are plotted as a function of the number of choice alternatives in Figure 4. The probability of responding given a stop signal increased with stop-signal delay in each choice condition (Logan, 1981; Logan & Burkell, 1986; Logan & Cowan, 1984; Logan et al., 1984). Mean signal-respond RTs were faster than no-stop-signal RTs (Logan & Cowan, 1984). Their distributions had similar lower tails and differed primarily in their upper tails (see Figure 5; Osman, Kornblum, & Meyer, 1990). Mean signal-respond RTs also increased with the number of choice alternatives, by 94 ms from two to four alternatives and by 64 ms from four to six alternatives.

The effect of number of choice alternatives on SSRT is important theoretically because it addresses whether the stop process shares capacity with the go process. We calculated mean SSRT for each subject using the integration method (see Equation 3; Logan, 1994; Logan & Cowan, 1984) and plotted the means across sub-
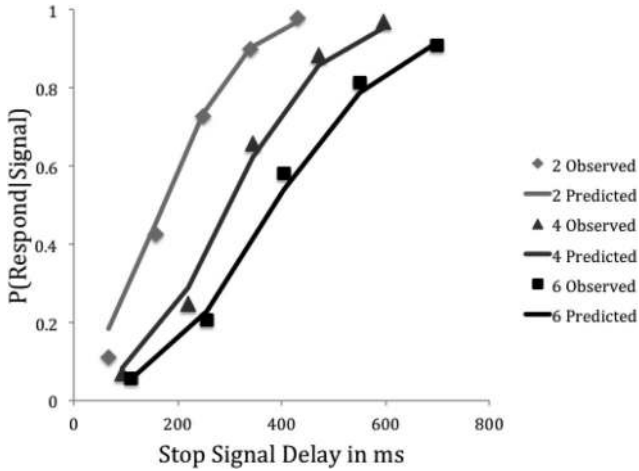
Figure 4. Observed (points) and predicted (lines) inhibition functions for diffusion race model averaged over subjects as a function of number of choice alternatives (2, 4, 6) and stop-signal delay in the multiple-choice experiment.

jects as a function of number of choice alternatives in Figure 3. Mean SSRT was not affected much by the number of choice alternatives, increasing by 7 ms from two to four choices and by 5 ms from four to six choices. These differences are small compared to the differences in no-stop-signal and signal-respond RT, suggesting that the stop task and the go task do not share capacity.

Analysis of individual-subject data showed that SSRT increased with number of choice alternatives in two of the six subjects (see Appendix A). For Subject 1, SSRT increased by 77 ms from two to six choices. For Subject 2, SSRT increased by 20 ms from two to six choices. These differences suggest the stop process was limited in capacity (but see Usher et al., 2002). Analysis of the rates of stochastic accumulation will be necessary to rule out unlimited-capacity processing. For the remaining subjects, SSRT did not increase with number of alternatives (differences between two and six choices were 0, −11, 7, and −33 ms for Subjects 3–6, respectively). Their results are consistent with an unlimited-capacity stop process. Analysis of the rates of stochastic accumulation will provide stronger evidence.

Results: Diffusion Race Model Fits

To evaluate the nature of capacity limitations in the stop and go processes, we fit eight versions of the diffusion race model to the multiple-choice data for correct and error responses (see Table 1). To fit the diffusion race models, we used Equations 16 and 17 to generate likelihood functions. For each model, we assumed there was one runner in the race for each choice alternative in the go task and one runner for the stop task. Each go runner was characterized by a rate and a threshold parameter. Within a set of choice alternatives, the correct response had a rate parameter of ξ_i, and each incorrect response had a rate parameter of ε_i. The rates were the same for each incorrect response in the set. There was one threshold z_i for all responses in a set of choice alternatives, though in different models the threshold could vary between sets of choice alternatives. The stop process had one rate parameter, ξ_{stop}, and one threshold, z_{stop}. There were two separate nondecision time
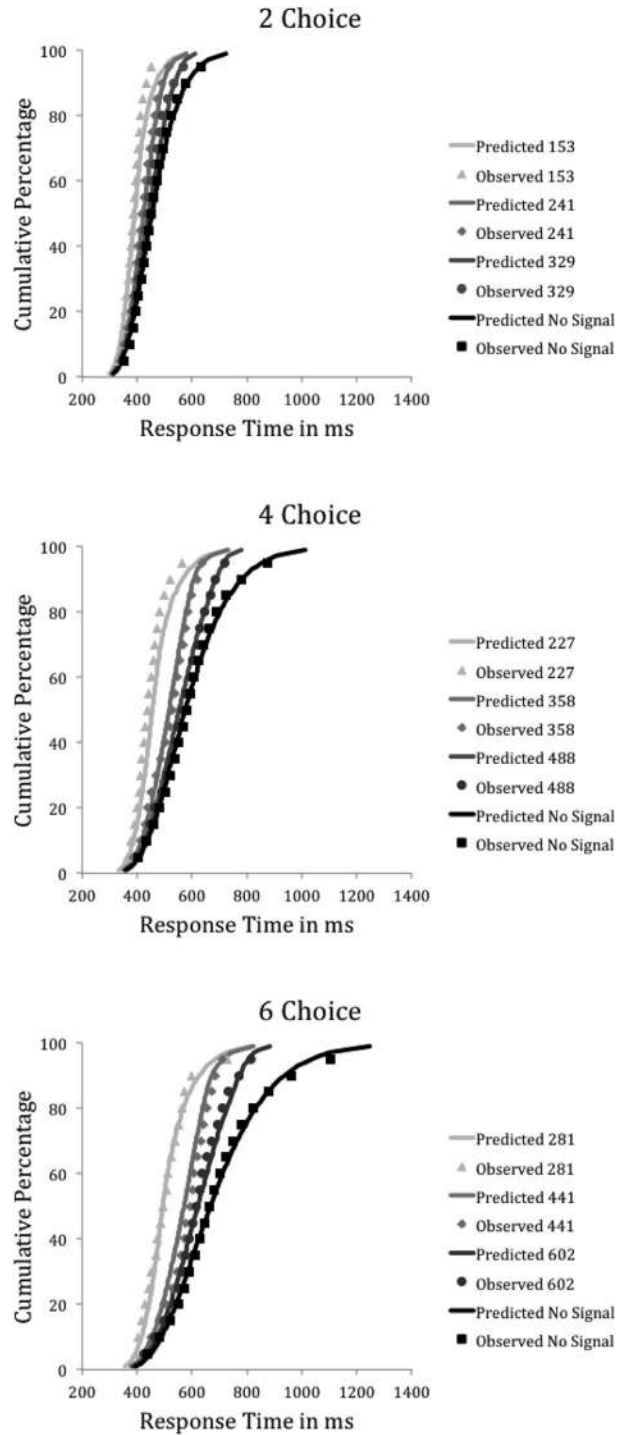

Figure 5. Diffusion race model fits to response time distributions for correct trials from the multiple-choice experiment. Quantile average response time distributions for no-stop-signal and signal-respond trials for the three middle stop-signal delays (153, 241, and 329 ms for two choices; 227, 358, and 488 ms for four choices; 281, 441, and 602 ms for six choices). The points represent the observed data. The lines represent predictions from the best diffusion race model, which assumed a limited-capacity go process and an unshared-capacity stop process. A: Two choice alternatives. B: Four choice alternatives. C: Six choice alternatives.

Table 1
Models Fitted to the Multiple-Choice Data

Model	Go threshold	Go rate	Stop threshold	Stop rate	Go capacity	Stop capacity	Param	Aggregate BIC	N fit best
1	Varied	Varied	Varied	Varied	Limited	Shared	17	590,976	0
2	*Varied*	*Varied*	*Fixed*	*Fixed*	*Limited*	*Unshared*	*13*	*590,736*	*2*
3	Varied	Varied	Varied	Fixed	Limited	Unshared	15	590,964	0
4	Varied	Varied	Fixed	Varied	Limited	Shared	15	590,881	1
5	Varied	Fixed	Varied	Varied	Unlimited	Shared	13	594,388	0
6	Fixed	Fixed	Varied	Varied	Limited	Shared	15	591,510	0
7	Varied	Fixed	Fixed	Fixed	Unlimited	Unshared	9	593,518	2
8	Fixed	Varied	Fixed	Fixed	Limited	Unshared	11	591,367	1

Note. The best fitting model in the aggregate fits is in bold italic font. Varied = parameter is allowed to vary with number of choice alternatives; Fixed = parameter is held constant across number of choice alternatives; Param = number of parameters; BIC = Bayesian information criterion.

parameters, one for the stop process and one for the go process, because stop and go stimuli were presented in different modalities. Within each process, nondecision time was not allowed to vary with number of choice alternatives (cf. Leite & Ratcliff, 2010). The threshold of the stop and go diffusions was allowed to vary uniformly between $z - a$ and $z + a$ to capture error RT distributions.

The eight versions of the diffusion race model differed in their assumptions about whether rates and thresholds for the go and stop processes were fixed or varied with the number of choice alternatives. We tested hypotheses about capacity limitations in the go process by comparing models in which go rates were fixed (unlimited capacity) or varied (limited capacity) with the number of choice alternatives. We tested hypotheses about shared capacity limitations in stop and go processes by comparing models in which stop rates were fixed (unshared capacity) or varied (shared capacity) with the number of choice alternatives. The models and their assumptions are presented in Table 1.

The fitting procedure found the parameter values that maximized the likelihood of the data (Myung, 2003; Van Zandt, 2000a). Each version of the model was fitted to the distributions of correct and error no-stop-signal RTs, the distributions of signal-respond RTs at each stop-signal delay, and the inhibition function. Each subject's data were fitted separately. Details of the fitting process are presented in Appendix B. Model fits were evaluated with the Bayesian information criterion (BIC; Raftery, 1995; Schwarz, 1978; Wagenmakers, 2007),

$$BIC = -2 \log L_i + k_i \log N, \tag{22}$$

where L_i is the maximized likelihood for model i, k_i is the number of parameters in model i, and N is the number of data points. Models with more parameters have greater flexibility and generally produce lower (i.e., better) negative maximum-likelihood values ($-2 \log L_i$). The last term in Equation 22 penalizes models with greater flexibility, adding $k_i \log N$ to the negative maximum-likelihood value. Models with lower BIC scores are preferred over models with higher BIC scores. We calculated aggregate BIC values over subjects by summing likelihoods, summing parameters, and summing numbers of observations and then applying Equation 22. Our use of aggregate BIC values assumes that the same model fit best for all subjects and that all subjects are independent of one another. We also calculated separate BIC values for each individual subject. The individual-subject BIC values allow us to evaluate consistency in the model fits across subjects.

The aggregate BIC values for each model are presented in Table 1. The model with the lowest aggregate BIC score assumed a limited-capacity go process and a stop process that did not share capacity with the go process (i.e., varied go rate and threshold, fixed stop rate and threshold; see Row 2 in Table 1). The predicted mean go RTs, mean SSRTs, and error rates (lines) are plotted with the observed values (points) in Figure 3. The model predicts no-stop-signal RT and signal-respond RT well, but it overpredicts error rate for six-alternative choices. The predicted inhibition functions (lines) are plotted with the observed inhibition functions (points) in Figure 4, showing close agreement. The predicted RT distributions (lines) are plotted with the observed values (points) in Figure 5. Like the observed values, the predicted values fanned out from a common minimum, with longer upper tails for longer stop-signal delays. This follows from the race model: The faster go RTs are fast enough to win the race regardless of the stop-signal delay, but the longer go RTs can be fast enough to win only when stop-signal delay is longer.

The values of the best fitting rate and threshold parameters for the model with the best aggregate fit (limited-capacity go, unshared-capacity stop), averaged across subjects, are presented in Table 2 and Figure 6. The rate parameters for the go process decreased as number of choice alternatives increased, indicating limited capacity (see Equation 17). The sum of the rates decreased from two choices ($\xi_{2go} + \varepsilon_{2go} = 0.252$) to four choices ($\xi_{4go} + 3\varepsilon_{4go} = 0.205$) to six choices ($\xi_{6go} + 5\varepsilon_{6go} = 0.153$), indicating stronger capacity limitations than a fixed-capacity model would predict (see Equation 18). One interpretation of this hyperlimited capacity is that it takes capacity to share capacity: Preparing and coordinating several response alternatives consume capacity that could be used for processing information (e.g., Logan, 1978, 1979; Pashler, 1994). Whatever the interpretation, the model fits suggest the go process is limited in capacity, and that allows us to ask whether the stop process shares the same capacity. In the best fitting model, the rate parameters for the stop process were constrained to be the same for each number of choice alternatives, suggesting that the stop process and the go process do not share capacity.

The fits to the individual-subject data were consistent with the aggregate fits. The number of subjects fit best by each model is presented in Table 1. The models of the go task differed between subjects: Some required rate changes, some required threshold changes, and some required both. The models of the stop task were more consistent: Five of the six subjects were fit best by an

Table 2

Mean Values Across Subjects for the Best Fitting Diffusion Race Model Parameters for the Multiple-Choice Data in the Aggregate Fits and for the Best Fitting Models for Individual Subjects

Choice	Go threshold	Threshold variability	Correct rate	Incorrect rate	Nondecision time	Stop threshold	Threshold variability	Stop rate	Nondecision time
				Best aggregate fit					
2	64.265	5.735	0.210	0.042	164	5.107	0.837	0.068	241
4	66.848	6.403	0.157	0.016	164	5.107	0.837	0.068	241
6	68.606	6.335	0.134	0.004	164	5.107	0.837	0.068	241
				Best fit for individual subjects					
2	61.890	2.139	0.200	0.039	160	4.882	8.837	0.089	241
4	66.375	3.086	0.154	0.016	160	4.882	8.837	0.082	241
6	68.173	3.155	0.134	0.004	160	4.882	8.837	0.077	241

unshared-capacity stop model. A shared-capacity stop model fit best for Subject 2, whose "observed" SSRTs (estimated from the data with the integration method; Equation 3) increased with number of choice alternatives, indicating a shared-capacity stop process.

The aggregate and individual-subject fits required similar changes in parameters to account for the data. We selected the best fitting model for each subject and averaged the best fitting parameter values across subjects. The average rate and threshold parameters for the stop and go processes are presented in Figure 6. As with the aggregate fits, the rate parameters for the go task change the most with the number of choice alternatives. Thus, go task capacity is limited. The rate parameters for the stop task do not change much with the number of choice alternatives. Thus, the stop task does not share capacity with the go task (Logan & Burkell, 1986; Yamaguchi et al., 2012).

The threshold parameter for the go task increases slightly with number of alternatives in both the aggregate fits and the average of the best fits to individual subjects. We interpret this as an adjustment subjects make to keep error rate low. The variability in stochastic evidence increases as drift rate decreases, and subjects may adjust threshold strategically to compensate for the increased noise.

SSRT distributions. The diffusion race model assumes that SSRT is a random variable, so the model fits allow us to estimate the distribution of SSRT. Two SSRT distributions are relevant. One is the *parent* distribution of SSRT from which runners in the race are sampled. The other is the *winning* distribution of SSRTs that are faster than the go process on individual runs of the race (i.e., on individual trials). The rate, threshold, and nondecision times for the stop process give the parent distributions. We calculated the winning distributions for the middle three stop-signal delays, using parameters from model with the best aggregate fit and Equations 10, 13, and 14, and we plotted them as cumulative distribution functions in Figure 7. The winning SSRT distributions shift to the left as stop-signal delay increases; the race is more biased against the stop process the longer the delay, so only the faster SSRTs win. This trend is opposite to the rightward shift in signal-respond RT distributions (see Figures 5 and 6), where increasing stop-signal delay biases the race in favor of the go process, allowing progressively slower go RTs to win the race.

The ability to estimate SSRT distributions is an important advance. Previously, Colonius (1990) and De Jong et al. (1990) proposed a nonparametric method for estimating SSRT distributions (see Equation 4), but their method is not practical. It requires very large amounts of data to produce stable estimates because calculations depend on the tails of the observed distributions, which are noisy. Matzke et al. (2013) developed a parametric method, which assumes that the SSRT distribution is ex-Gaussian. They provided powerful Bayesian methods that allow accurate estimation of ex-Gaussian parameters even with small amounts of data. However, their approach is descriptive, aimed at characterizing SSRT distributions and not the processes that generate them. Our method assumes a parametric form for the SSRT distributions (see Equations 13–17), and the parameters (rate and threshold) are readily interpretable as psychological processes.

Error RTs. The ability to fit error probabilities and RTs has become an important criterion for evaluating models of RT. The diffusion race model predicted error probabilities relatively well but could not capture the distribution of error RTs (see Figure 8, top panel). In developing the current model, we first fit a diffusion race model that assumed no variability in threshold. The diffusion model with no threshold variability predicted error probability relatively well but predicted error RTs that were much longer than observed error RTs. Then we fit the current diffusion race model that assumes variability in threshold, which often allows models to capture error RT distributions (Ratcliff & Smith, 2004). However, the diffusion race model with threshold variability did not do much better than the model with no threshold variability. Predictions from the current diffusion race model, which assumes threshold variability, are shown in the top panel of Figure 8. The estimated threshold variability was small—less than 10% of the threshold in the best fit to the aggregate and less than 5% of the threshold in the average of the individual-subject fits (see Table 2).

We suggest two interpretations of this failure to fit error RT distributions. One interpretation is that the diffusion race model does not fit the data well and should be rejected. Another interpretation is that there were not enough error data to allow the model to fit the errors. The overall error rate was less than 2%, so the contribution of error likelihood to the overall likelihood was very small. The fitting routine will be dominated by the vast majority of correct responses. To illustrate, we plotted the correct and error RT data as *defective distributions* in the bottom panel of Figure 8. Defective distributions are analogous to cumulative frequency distributions but describe the probabilities of observing an RT less than some value and the response is correct or incorrect.
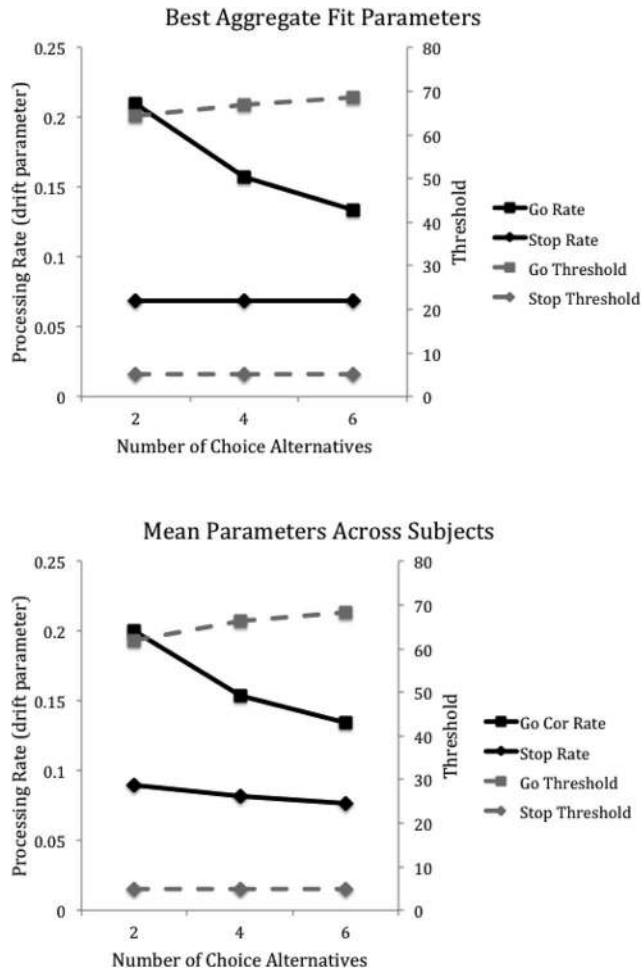
Best Aggregate Fit Parameters

Mean Parameters Across Subjects

Figure 6. Values of the best fitting rate and threshold parameters for the stop process and the go process in diffusion race model fits to the multiple-choice data. The top panel contains parameters from the model that fit the aggregate data best (limited-capacity go, unshared-capacity stop). The bottom panel contains the average of the parameters for the best fitting model for each subject.

Rather than ranging from 0 to 1, they range from 0 to the probability of the response being correct or incorrect. The defective distributions for correct response rise nearly to 1.0, while the defective distributions for error responses barely reach 0.02. At this scale, the failure to predict error RT distributions does not look very substantial.

We also tried fitting the data with a version of the linear ballistic accumulator model, which has fit error RT distributions successfully (Brown & Heathcote, 2008). The model assumed each runner involved a linear increase to a threshold, and the slope of the linear increase for each runner was drawn from a normal distribution with a mean of ξ_i and a standard deviation of 1.0. The model fit the data set about as well as the diffusion race model but also failed to fit the error RT distributions. Predicted error RTs were much longer than observed error RTs. The linear ballistic accumulator model has fit error data well in other contexts (Brown & Heathcote, 2008), so we interpret its failure to capture error RTs as a limitation of our data set rather than the model.

We ran simulations of the diffusion race model and found that it could predict fast errors when threshold variability was high. In Appendix C, we report the results of one simulation in which we reduced the difference between correct and error drift rates to produce a mean error rate of 36% and set threshold variability equal to 75% of the threshold. Error RTs were faster than correct RTs. We then fit the model to the simulated data and found that it accounted for fast errors and recovered the parameters accurately. Thus, we conclude that the diffusion race model is capable of
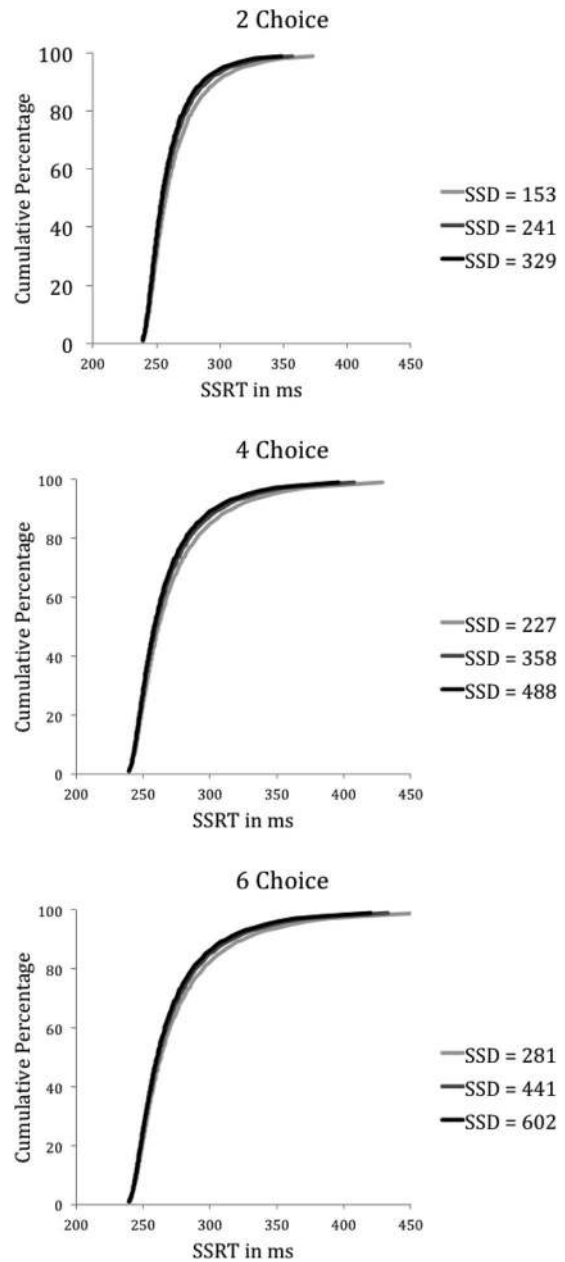
2 Choice

4 Choice

6 Choice

Figure 7. Distributions of stop-signal response times (SSRTs) that won the race against the go process estimated from the limited-capacity go, unshared-capacity stop diffusion race model in the multiple-choice experiment. Top panel: Two-choice go task. Middle panel: Four-choice go task. Bottom panel: six-choice go task. SSD = stop-signal delay.
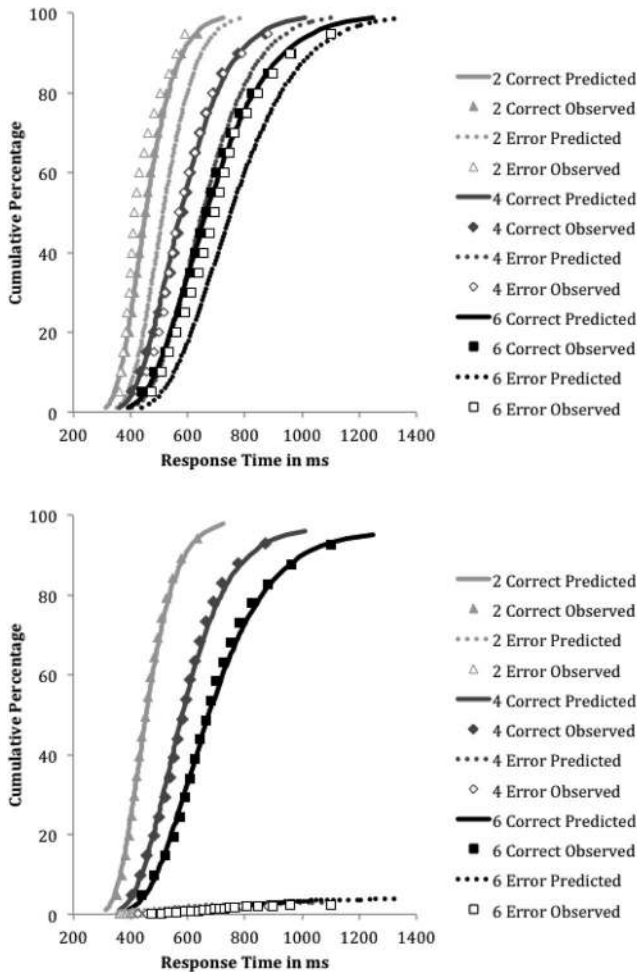
Figure 8. Distributions of correct and error response times for no-stop-signal trials for two-, four-, and six-choice alternatives in the multiple-choice experiment. The top panel presents cumulative distribution functions that asymptote at 100%. The bottom panel presents defective distribution functions that asymptote at the response probability.

producing fast errors, and we interpret the poor fits to the current error RTs as indicating that error rate was too low to allow us to evaluate error RTs properly.

Discussion

The analysis of the data and the analysis of the models suggest that the stop process does not share capacity with the go process. Go RT increased dramatically as the number of choice alternatives increased, suggestive of limited capacity (Townsend, 1971, 1990; Townsend & Ashby, 1983), but mean SSRT increased only slightly. The fits of the diffusion race model allowed us to interpret the increase in go RT with the number of choice alternatives as evidence of capacity limitations because limited-capacity models of the go process were necessary to fit the data. The same fits allowed us to conclude that the stop process did not share capacity with the go process. The best fitting models assumed that the stop processing rates were constant across numbers of choice alternatives, indicating that the stop process did not share capacity with

the go process. This conclusion is consistent with functional independence between go and stop processes (see Equation 6).

Our finding that the stop process does not share capacity with the go process is remarkable because many processes share capacity (Pashler, 1994), especially control processes (Logan, 1978, 1979; Shiffrin & Schneider, 1977). It is worth speculating on reasons why the stop process may escape the limitations that apply to other processes. One possibility is that the stop task is given higher priority than other tasks, and high priority tasks are given first access to limited processing capacity (Meyer & Keiras, 1997). In everyday life, the act of control underlying stop-signal inhibition is recruited to compensate for errors, for sudden changes in input, or for sudden changes in goals that make the current course of action inappropriate or irrelevant. Compensating for these changes should have a higher priority than continuing an inappropriate or irrelevant course of action. This suggests that stop signals might produce dual-task interference in processing subsequent go stimuli (e.g., Horstmann, 2003), which might explain post-stop-signal slowing (Bissett & Logan, 2011; Rieger & Gauggel, 1999). Further research is required to explore this possibility.

Another possibility is that control is hierarchical and stop-signal inhibition is recruited by a higher level system that is not subject to the same capacity limitations as the lower level system that chooses responses and executes them (Logan & Cowan, 1984; Logan & Crump, 2011). However, the evidence for hierarchical control is controversial despite the widespread appeal of the idea (see Cooper & Shallice, 2000, 2006, vs. Botvinick & Plaut, 2004, 2006), and it is not clear that higher level processes rely on different capacities than lower level ones (see Logan, 1979). More research is required to evaluate this possibility.

A third possibility is that the versions of the stop and go tasks that we have investigated are not sufficiently demanding to show evidence of capacity limitations. The stop task involves a single response to a single tone, like a simple RT task, and so may not demand much processing capacity. More difficult stop tasks that require discrimination among stop signals produce longer SSRTs (Bedard et al., 2002) and sometime produce violations of the race model (Bissett & Logan, 2013). *Change* tasks that require an overt response to the stop signal in addition to inhibiting the go response also increase SSRT (Logan & Burkell, 1986). Go tasks that require inhibition of competing alternatives also produce longer SSRTs, suggesting capacity limitations (Kramer et al., 1994; Ridderinkhof, Band, & Logan, 1999; Verbruggen, Liefooghe, & Vandierendonck, 2004; but see Verbruggen, Liefooghe, Szmalec, & Vandierendonck, 2005; Verbruggen, Liefooghe, & Vandierendonck, 2005). Thus, the conditions under which SSRT is affected by complexity in the stop task and go task require further research. The special race models developed in this article may be useful in determining whether the increased SSRTs reflect capacity limitations or strategies.

Control Strategies in the Stop-Signal Task

The stop-signal task presents subjects with diametrically opposing demands. The faster they perform the go task, the less likely they are to succeed at stopping; the slower they perform the go task, the more likely they are to succeed at stopping. Subjects often cope with these demands by strategically slowing performance on the go task to increase their likelihood of stopping: Go RT is slower when stop signals occur more frequently (Bissett & Logan,

2011; Logan, 1981; Logan & Burkell, 1986; Ramautar et al., 2004) and when subjects are given signals that indicate that stop signals are likely (Chikazoe et al., 2009; Verbruggen & Logan, 2009c). This strategic slowing may be severe enough to subvert the experiment (Leotti & Wager, 2010), invalidating estimates of SSRT (Verbruggen et al., 2013). Researchers try to eliminate strategic slowing by introducing the go task before the stop task so subjects can learn how quickly they can respond to it without competing demands, by instructing subjects to avoid slowing, or by rewarding fast go responses. However, subjects often slow strategically when stop signals are introduced despite these precautions (Ramautar et al., 2004; Verbruggen, Liefooghe, & Vandierendonck, 2004, 2006).

Researchers have assumed that slowing in anticipation of a stop signal is strategic because it is proactive and occurs rapidly in response to changes in stop-signal probability and cues. Our diffusion race model addresses the mechanism underlying strategic slowing, attributing the slowing to parameters of the model that the executive system can adjust strategically. Thus, we expect proactive slowing to be explained by threshold adjustment and not by changes in drift rate. Proactive slowing may also be explained by delaying the onset of stochastic accumulation (Pouget et al., 2011), measured as an increase in nondecision time.

Explicitly Cuing Stop-Signal Relevance

We fit the diffusion race model to an experiment by Verbruggen and Logan (2009c) that showed strategic slowing in response to explicit cues indicating whether subjects had to stop when stop signals occurred. Eighteen subjects were given the explicit cues *all* or *none* that told them how many of the stop signals required stopping in the next four to eight trials. The go task involved classifying characters (discriminating *Z* and */*). The stop signal was a tone. Tones occurred on 33% of all trials. Stop-signal delay was adjusted by a tracking algorithm that produced successful stopping on 50% of stop-signal trials. The cue was presented for 1,000 ms before the first trial in a run and remained on the screen throughout the run. Each subject performed 576 trials, divided evenly between *all* and *none* conditions.

Verbruggen and Logan (2009c) designed their experiment to evoke strategic slowing. The blocks of *all* and *none* trials were cued explicitly and were very short (four to eight trials). This led to very rapid changes in RT on the first trial after a cue change. Thus, we should expect differences in diffusion race model parameters that executive processes control strategically, like threshold and maybe onset. We should expect no differences in diffusion race parameters that reflect structural and informational limitations, such as drift rate.

Results. The observed go RTs, signal-respond RTs, SSRTs, and error rates from Verbruggen and Logan (2009c) were averaged across subjects and plotted in Figure 9. The data show that subjects responded in accord with the cues. They inhibited responses on 53% of the trials when the stop signal was relevant (*all* cues) but only on 2% of the trials when the stop signal was irrelevant (*none* cues). Subjects also slowed strategically following the cues: Go RT was 501 ms when stop signals were relevant and 408 ms when stop signals were irrelevant. The slowing was accompanied by an increase in accuracy, suggestive of a threshold adjustment: Accuracy was 97% when stop signals were irrelevant and 99% when stop signals were relevant. Further analysis showed that slowing
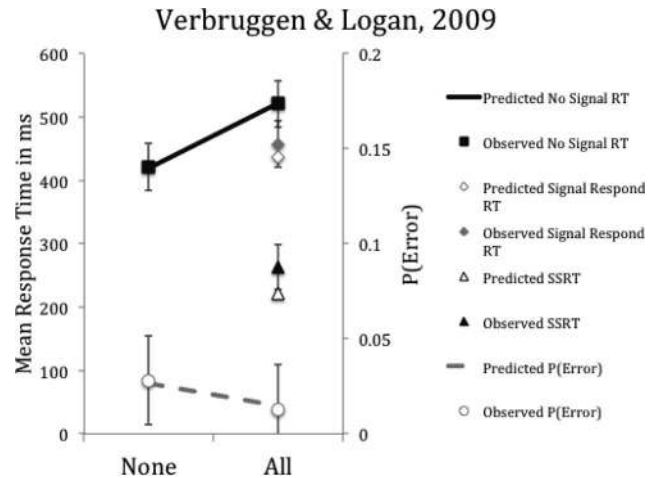

Figure 9. Mean observed (points) and predicted (lines) go response time (RT), signal-respond response time, stop-signal response time (SSRT), and error rate—*P*(Error)—for diffusion race model for conditions in which none of the stop signals were relevant and all of the stop signals were relevant in Verbruggen and Logan (2009c). Error bars are 95% confidence intervals.

occurred on the very first trial in a run after the cue changed, supporting the idea that the slowing was strategic. When stop signals were relevant, SSRT was 263 ms.

Verbruggen and Logan (2009c) fit the Fast-dm version of the two-choice diffusion model (Voss & Voss, 2007) to the no-stop-signal go data and found that threshold and nondecision time changed with stop-signal relevance but drift rate did not. These results are consistent with our hypothesis that strategies in the stop-signal task affect threshold but not drift rate. Our modeling addresses more of the data.

Diffusion race model fits. We fit a set of eight diffusion race models to the data (see Table 3), fitting no-stop-signal RT distributions for correct and error responses when stop signals were and were not relevant and signal-respond distributions for correct responses as a function of stop-signal delay. The set of models was generated from the factorial combination of fixing versus varying threshold, rate, and nondecision time between the relevant stop-signal and irrelevant stop-signal conditions (see Table 3). We also fixed and varied the variability in a uniform distribution of thresholds, in an attempt to capture error RTs. We fixed variability in threshold between conditions whenever we fixed threshold between conditions and let it vary between conditions whenever we let threshold vary between conditions. We did not see much theoretical value in allowing threshold variability to vary while threshold was fixed between conditions, so we excluded those conditions to reduce the number of models we fit from 16 to eight.

For each of the eight models, we assumed there was one runner in the race for each of the two choice alternatives in the go task and one runner for the stop task, and each runner was characterized by a rate and a threshold parameter. The correct go response had a rate parameter of ξ_i, and the incorrect response had a rate parameter of ε_i. There was one mean threshold z_i for both go responses. The stop process had one rate parameter, ξ_{stop}, and one threshold, z_{stop}. Both stop and go thresholds were allowed to vary uniformly between

Table 3
Models Fitted to the Verbruggen and Logan (2009c) Data

Model	Go threshold	Go rate	Go threshold variability	Go nondecision time	Param	Aggregate BIC	N fit best
1	Fixed	Fixed	Fixed	Fixed	9	97,954	0
2	Fixed	Fixed	Fixed	Varied	10	96,920	0
3	*Varied*	*Fixed*	*Varied*	*Fixed*	*11*	*96,438*	*9*
4	Varied	Fixed	Varied	Varied	12	96,595	5
5	Fixed	Varied	Fixed	Fixed	11	96,712	2
6	Fixed	Varied	Fixed	Varied	12	96,674	1
7	Varied	Varied	Varied	Fixed	13	96,622	1
8	Varied	Varied	Varied	Varied	14	96,785	0

Note. The best fitting model in the aggregate fits is in bold italic font. Param = number of parameters; BIC = Bayesian information criterion.

$z_i - a_i$ and $z_i + a_i$ to capture error RT distributions. The stop and go processes had separate nondecision times. The eight models differed in whether these parameters were fixed or were allowed to vary between *all* and *none* trials.

We fit the models to the data by maximizing the likelihood using the methods described in Appendix B. We used BIC to evaluate goodness of fit, calculating an aggregate BIC and BICs for individual subjects.

The BIC values for the aggregate fits are presented in Table 3. The best fitting model was one in which threshold varied but rate and nondecision time stayed constant as stop-signal relevance was manipulated (see Row 3 in Table 3). This model fit best in nine of the 18 individual-subject fits. The next most popular model, which fit best in five of the 18 subject fits, was one in which threshold and nondecision time varied but rate remained constant as stop-signal relevance was manipulated (see Row 4 in Table 3). This model is similar to the one that Verbruggen and Logan (2009c) fit to their data. Overall, 15 of 18 subjects were fit best by a model with variable threshold, 14 of 18 were fit best by a model with fixed rate, and 12 of 18 were fit best by a model with fixed nondecision time.

The mean go RTs, signal-respond RTs, SSRTs, and error rates predicted from the best fitting model (threshold varied, rate, and nondecision time fixed) are plotted along with the observed values in Figure 9. There is good agreement between predicted and observed values for all measures, except for SSRT. All predicted values fell within the 95% confidence intervals of the observed values, except for predicted SSRT, which was 41 ms faster than observed SSRT (estimated from the data with the integration method).

The predicted and observed RT distributions for correct no-stop-signal RTs in the *all* and *none* conditions and for signal-respond RTs in the *all* condition are plotted in Figure 10. There is good agreement between predicted and observed values.

The values of the best fitting parameters for the go task are plotted as a function of stop-signal relevance in Figure 11. The top panels present the parameters from the model that fit the aggregate best, which assumed varied threshold, constant rate, and constant nondecision time. The bottom panels present the average parameter values for the best fitting model for each subject as a function of condition. The patterns are similar: Threshold and threshold variability were larger in the *all* condition than in the *none* condition in both aggregate and individual-subject fits. Drift rate and nondecision time were the same for *all* and *none*

in the aggregate fits. They both decreased from *none* to *all* in the individual-subject fits, trading off their effects on RT: The reduced drift rate would slow RT but the reduced nondecision time would speed it. Altogether, these results are consistent with selective influence and our hypothesis that strategies affect thresholds more than rates.

The predicted distribution of SSRT is plotted in Figure 12. For the best aggregate fit, the mean stop rate and threshold were 0.160 and 0.684, respectively. For the individual-subject fits, the mean stop rate and threshold were 0.153 and 0.998, respectively.

The predicted and observed distributions for correct and error RTs are presented in Figure 13. The top panel presents cumulative distribution functions, which asymptote at 1.0. The bottom panel presents defective distributions, which asymptote at response probability. As before, the model captured correct RT distributions and error probabilities well but overpredicted error RTs. Again, we note that error rate was very low (2%), so there may not have been enough errors to allow the model to fit the error RTs (see Appendix C).
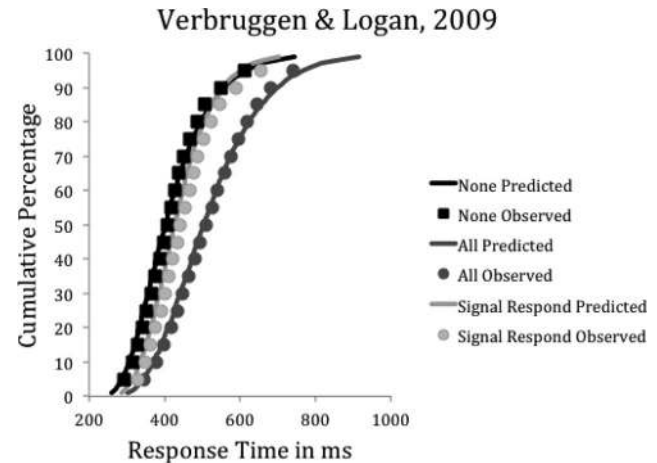

Figure 10. Diffusion race model fits to response time distributions for correct trials from Verbruggen and Logan (2009c). Quantile average response time distributions for no-stop-signal trials in the condition in which none of the stop signals were relevant and for no-stop-signal and signal-respond trials in the condition in which all of the stop signals were relevant. The points represent the observed data. The lines represent predictions from the best diffusion race model.
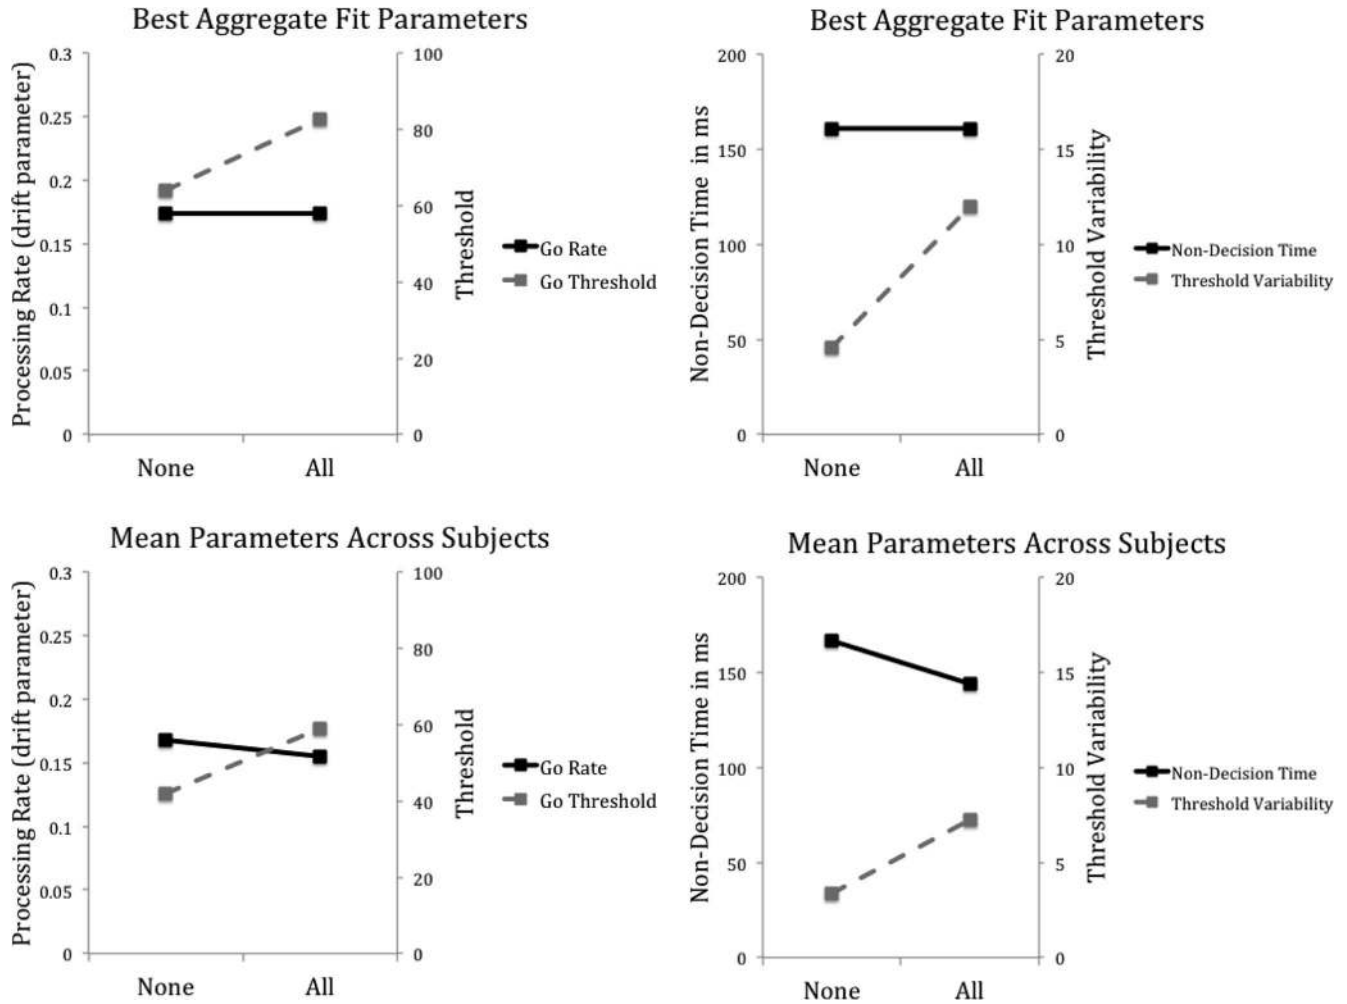
Figure 11.   Values of the best fitting rate and threshold (left panels) and threshold variability and nondecision
time (right panels) parameters for the diffusion race model fits to the Verbruggen and Logan (2009c) data. The
top panels contain parameters from the model that fit the aggregate data best. The bottom panels contain the
average of the parameters for the best fitting model for each subject.

Discussion.   The model fits suggest that the strategic slowing
reported by Verbruggen and Logan (2009c) is best described by a
model in which threshold and threshold variability change when
stop signals become relevant but rate and nondecision time stay
constant. This model fit is consistent with our hypothesis that
executive processes adjust threshold and perhaps onset of accu-
mulation strategically but do not adjust drift rate. The present
results, together with the multiple-choice results, show a predicted
selective influence of experimental manipulations on diffusion
model parameters. Manipulations that loaded capacity changed
drift rate but not threshold. Manipulations that shifted strategies
changed threshold but not drift rate.

Percentage of Stop Signals

Bissett and Logan (2011, Experiment 1) conducted an experi-
ment in which 24 subjects performed a stop-signal task in which
stop signals occurred on 20% or 40% of the trials, pitting the
priority of the go task against the priority of stopping. RT became

slower as percentage of stop signals increased (also see Logan,
1981; Logan & Burkell, 1986; Ramautar et al., 2004; Verbruggen
& Logan, 2009c). We modeled this slowing by fitting a set of 16
diffusion race models to the data. The set of models was created
from the factorial combination of fixing versus varying go rate, go
threshold, go nondecision time, and stop process parameters be-
tween the 20% and 40% stop-signal conditions. Summaries of the
results and model fits are presented in the online supplemental
materials.

The fits were good, but there was not much consensus on the
best fitting model. Ten of the 16 models fit best for at least one
subject; no model fit best for more than four subjects. Collapsing
across models, there was consistency for the stop process: Twenty
subjects were fit best by models with stop rates and thresholds
fixed for 20% and 40% stop signals. There was less consistency for
the go process: Nine subjects were fit best with go rates fixed, and
11 were fit best with go thresholds fixed. This inconsistency
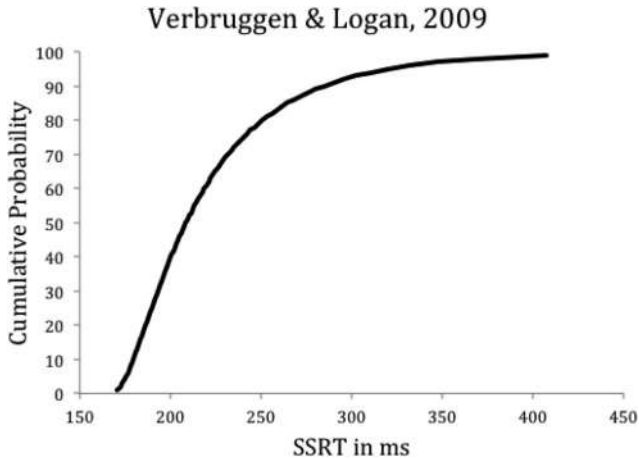challenges the selective-influence hypothesis: For some subjects,

Figure 12. Distributions of stop-signal response times (SSRTs) that won the race against the go process estimated from the diffusion race model with the best aggregate fit to the Verbruggen and Logan (2009c) data.

strategies were mediated by changes in drift rate. We discuss possible top-down influences on drift rate in the online supplemental materials.

General Discussion

We developed a general independent race model that construed the go process as a race between alternative responses and the stop process as another runner in the race (see Equations 8–10). We developed special race models in which each runner in the race is a single diffusion process (see Equations 13–17) whose parameters are selectively influenced by structural and strategic factors.

We tested the selective influence of capacity limitations on drift rate in a new multiple-choice stop-signal task. Increasing the number of choice alternatives decreased drift rate but had little effect on threshold, consistent with selective influence. Analysis of drift rates showed that the stop process and the go process did not share capacity (Logan & Burkell, 1986; Yamaguchi et al., 2012). The stop process may have had its own capacity limitations, but it did not share them with the go process.

We tested selective influence of strategy on threshold by fitting data from an experiment by Verbruggen and Logan (2009c) in which the relevance of stop signals was manipulated. Increasing stop-signal relevance increased threshold but had little effect on drift rate or nondecision time, consistent with selective influence. We note that selective influence was not strongly evident in the fits to Bissett and Logan (2011; see the online supplemental materials).

More generally, the studies of selective influence illustrate the use of the special model to generate specific hypotheses about the stop and go processes and test them rigorously by competitive model fitting. The independent race model (Logan & Cowan, 1984) described stop-signal performance in terms of relations among observed RT distributions. The special diffusion race model explains stop-signal performance in terms of changes in its parameters between experimental conditions. This is an important advance over previous theories.

An important strength of both general and special race models is that they deal with choice in the go task. No previous model of

stop-signal performance has dealt with choice in the go task (Boucher et al., 2007; Lo et al., 2009; Logan & Cowan, 1984; Wong-Lin et al., 2010). Some of these models dealt with stopping saccadic eye movements, where choice errors almost never happen. However, most stop-signal tasks use keypress responses, where most subjects make choice errors. Thus, the ability to deal with choice is an important step forward, and our manipulation of the number of choice alternatives was a significant challenge. The general model characterizes choice as a race and defines the race architecture. The special model characterizes each runner as a diffusion process and specifies its parameters.

SSRT Distributions

The diffusion race model provides an estimate of the distribution of SSRT and relates it to the parameters of stochastic accumulation of evidence about the stop signal. This is an important advance. Previous models provided point estimates of SSRT (Lo-
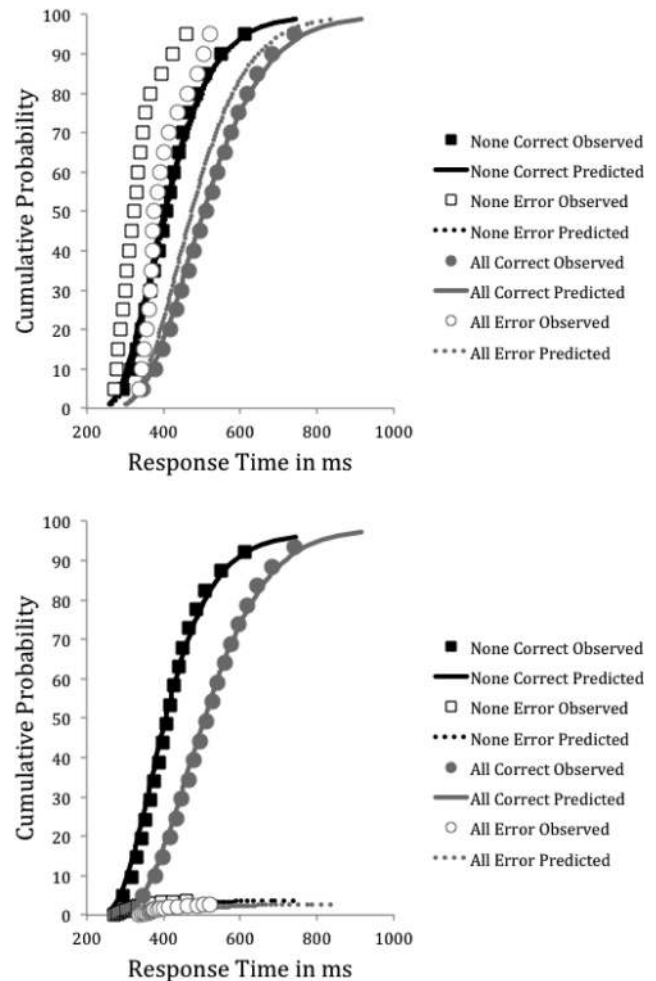

Figure 13. Distributions of correct and error response times for no-stop-signal trials for the none and all conditions from Verbruggen and Logan (2009c). The top panel presents cumulative distribution functions that asymptote at 100%. The bottom panel presents defective distribution functions that asymptote at the response probability.

gan & Cowan, 1984) and nonparametric methods for estimating SSRT distributions that required impractical numbers of trials to overcome noise in the tails of the distributions (Colonius, 1990; De Jong et al., 1990). The ability to account for RT distributions is an important desideratum in developing models of cognitive processes (Balota & Yap, 2011; Ratcliff & Smith, 2004; Van Zandt, 2000a).

The diffusion race model estimates of SSRT distributions complement recent Bayesian hierarchical methods for estimating SSRT distributions developed by Matzke et al. (2013). Their methods also parameterize the stop and go processes, characterizing their finishing-time distributions as ex-Gaussian (i.e., the convolution of an exponential distribution and a normal, or Gaussian, distribution). The ex-Gauss describes RT distributions accurately and is used widely in the literature (Balota & Yap, 2011; Ratcliff & Murdock, 1976). However, its parameters do not map directly onto psychological processes, like rates and thresholds in stochastic accumulators (Matzke & Wagenmakers, 2009), so the ex-Gauss fits provide limited insight into the underlying stop and go processes. The diffusion race model provides greater insight.

An advantage of Matzke et al.'s (2013) Bayesian hierarchical method is that it can be applied to rather small data sets provided there are lots of subjects. The diffusion race model did well with large (multiple-choice) and small (Verbruggen & Logan, 2009c) data sets. It would be interesting to see how few data it requires. It should be possible to implement the diffusion race in Bayesian hierarchical modeling to gain the same advantages as the Matzke et al. model.

Go RT Distributions

The diffusion race model accurately predicted response probabilities and go RT distributions for correct responses on no-stop-signal trials and signal-respond trials. It accurately predicted error probability, but it overpredicted error RT distributions. The fit to the correct RT distributions and response probability is an important advance (cf. Boucher et al., 2007; Lo et al., 2009; Logan & Cowan, 1984; Salinas & Stanford, 2013; Wong-Lin et al., 2010). Moreover, the model captured the relation between no-stop-signal and signal-respond RT distributions accurately in most of the fits. Most likely, this follows more from the general race architecture than the special diffusion race model. It is a general property of an independent race (Logan & Cowan, 1984).

The overprediction of error RT is troublesome because the ability to predict error RT distributions is an important criterion for evaluating stochastic accumulator models (Ratcliff & Smith, 2004; Teodorescu & Usher, 2013). We interpret the overprediction as a result of error rates that were too low to influence goodness of fit. They were 2% in the multiple-choice experiment and 2% in Verbruggen and Logan (2009c). The model can simulate and fit fast errors in data with a higher error rate (36%) and larger threshold variability (see Appendix C).

Overprediction of error RT is not unique to the diffusion race model. We found it when we fit the Poisson race model (Van Zandt et al., 2000) and the linear ballistic accumulator model (Brown & Heathcote, 2008) to the multiple-choice data. The linear ballistic accumulator model accounts for error RT distributions in conditions with higher error rates, so its failure to account for error

RT distributions here should not be viewed as a limitation of the model.

Our analysis in Appendix C suggests the overestimation can be overcome by running stop-signal experiments with higher error rates, for example, by manipulating the discriminability of the go stimuli (Logan, 1981). However, the stop-signal task is used with many special populations whose cognitive capacities are impaired or have not yet developed, and these populations work best with go stimuli that are easy to discriminate and so produce low error rates. Hence, it may be better to use the standard tasks to get high-quality data and set aside concerns about predicting error RT distributions.

Other Special Race Models?

The general independent race model places two main constraints on the special race models implemented within it: The stop process must race with the go process, and the go process must also be a race among choice alternatives. Many special race models are possible within these constraints, with different assumptions about the stochastic accumulators for each runner. We tried a Poisson counter model (Van Zandt, 2000b) and the linear ballistic accumulator model (Brown & Heathcote, 2008), but we focused on a diffusion to a single threshold. We hope to explore other alternatives. The leaky competing accumulator model (Usher & McClelland, 2001) is promising because it applies to multiple-choice tasks. It assumes interaction between alternative responses instead of an independent race, so the stop-signal task would have to be modeled as a race between two runners: the stop process and the go response that wins the competition in the leaky competing accumulator model. Townsend and Wenger (2004) provided a useful analysis of channel interactions in this model. It would also be tempting to model the go process as the standard two-choice diffusion model that pervades much of the literature (Ratcliff, 1978; Ratcliff & McKoon, 2008). However, the standard diffusion model cannot deal with more than two choices, so it could not account for our multiple-choice data (but see Leite & Ratcliff, 2010). It assumes competition between the two alternative responses, so the stop process would race against the go response that wins the competition.

We are most strongly committed to the assumption that the stop process races with the go process. It follows from the original independent race model (Logan & Cowan, 1984). We are less strongly committed to the assumption that the go process is a race. We prefer race models, but we have no problem with special models in which the go process is interactive (Ratcliff & McKoon, 2008; Usher & McClelland, 2001). We are least strongly committed to the assumption that each runner in the race is a diffusion process, and we are open to the possibility that other stochastic accumulator models may fit better. Our theory addresses the parameters of stochastic accumulators—thresholds, rates, and non-decision times—that are common to all models. We are committed to the assumption that the runners in the race are stochastic accumulators, but we are not strongly committed to any specific model.

Acts of Control

We believe we can understand cognitive control by understanding the acts of control that implement it (Logan, 1985; Logan &

Cowan, 1984; Logan & Gordon, 2001). We can characterize the acts of control in the executive's repertoire by listing all of them or by describing features that are common to all of them. We see merit in both approaches.

Acts of control are like other acts a person performs. They are like the perceptual-motor acts in typical RT experiments. They are instigated by a condition of the external or internal environment; they have a duration, which we assume depends primarily on stochastic accumulation; and they end with an action that changes the state of a subordinate process (Logan, 1985). Acts of control are different from other acts in that some of their inputs and all of their outputs are changes in states of subordinate processes. Thus, acts of control must be grounded in a theory of the subordinate processes they control (Logan & Gordon, 2001).

Our theory of executive processing is grounded in our theory that subordinate processes are stochastic accumulators. The parameters of the stochastic accumulators are the points of contact that allow the executive to control the subordinates. Some parameters are set by the executive, and some are set by the environment and the person's history. A task set is a set of drift rate, threshold, and nondecision time parameters that program the stochastic accumulator to make a task-relevant choice. Changes in task set are accomplished by acts of control that change these parameters (Logan & Gordon, 2001).

It is tempting to suggest a taxonomy of acts of control, depending on the immediacy and extremity of their effects (cf. Norman, 1981). Stopping is immediate and extreme. SSRT is short, and the response is either inhibited or executed. Shifting attention is immediate but less extreme. Attention shifts quickly, but it modulates ongoing responses rather than preventing or enabling them. Proactive slowing is less immediate and less extreme. It occurs before or between trials, and it modulates performance on the next trial.

However, commonalities among acts of control may be more important than the differences. In our theory, all acts of control have an instigating condition, a duration, and an effect, and all effects are the same: They change the parameters of the stochastic accumulators. The change may be more or less extreme, but parameters always change. The act of control that stops a response requires a change in drift rate that is large enough to keep the accumulators from reaching threshold (Boucher et al., 2007; Salinas & Stanford, 2013). The act of control that shifts attention requires smaller a change in drift rate to facilitate processing at the cued location (Logan & Gordon, 2001). The act of control that produces proactive slowing requires a change in threshold or nondecision time (Dutilh et al., 2012; Pouget et al., 2011).

Our theory suggests further commonalities between acts of control and the controlled subordinate actions they govern. They are built on the same substrate—stochastic accumulator models. They differ primarily in the conditions that trigger them and the actions they take—in content rather than form. Acts of control represent and change the state of the cognitive system, whereas controlled acts represent and change the state of the world. They both act in the same way, by stochastic accumulation to a threshold.

In this respect, our theory is similar to production system theories of cognitive control, in which performance is modeled as the application of "if . . . then" rules to perceptual, memory, and motor systems (Anderson et al., 2004; Meyer & Keiras, 1997; Newell, 1990). All productions have the same structure. There is a condi-

tion (if) that triggers the production and an action (then) that is taken when the production is triggered. The difference between productions representing acts of control and productions representing controlled acts is in content, not form.

It may be profitable to think of stochastic accumulator models and production rules as different perspectives on the same mental operations. The condition term in the production rule specifies the input that drives the stochastic accumulator. The action term in the production rule specifies the action that is taken when the stochastic accumulator hits threshold. The two perspectives provide a more complete description that may be useful in identifying, measuring, and cataloging acts of control (and controlled actions).

Our theory that acts of control and controlled acts are both implemented as stochastic accumulators suggests they should be affected similarly by habit and strategy. Controlled actions show Stroop (1935) and Stroop-like effects that are modulated by strategy (Logan & Zbrodoff, 1979; Tzelgov, Henik, & Leiser, 1990). We see similar effects in acts of control. SSRT is faster when stop signals repeat (Bissett & Logan, 2012), and SSRT is slower when the stop signal is "GO" than when it is "STOP" (Verbruggen & Logan, 2009a). Acts of control can be associated with stimuli and processing episodes, just as controlled actions are (Logan & Etherton, 1994), suggesting short-term and long-term priming effects. Subjects associate stopping with go stimuli that are paired with stop signals and show longer go RTs when the go stimuli repeat, even at long retention intervals (Lenartowicz, Verbruggen, Logan, & Poldrack, 2011; Verbruggen & Logan, 2008b, 2009a; Verbruggen, Logan, Liefooghe, & Vandierendonck, 2008). Subjects may even automatize control (Verbruggen & Logan, 2008a) and transfer it to other situations (Verbruggen, Adams, & Chambers, 2012). All of these effects are readily interpretable as changes in drift rates, thresholds, and nondecision times in stochastic accumulators.

Our theory suggests that cognition and cognitive control are driven by events. Some events trigger acts of control. Other events trigger controlled acts. Our theory suggests that cognition and cognitive control involve discrete acts that begin with some event (like a go stimulus) and end with another (like a keypress). The reaction time experiment is its paradigm case: a stimulus followed by a response (Donders, 1969/1868; Sternberg, 1969). We do not address more continuous control or hierarchical control (Logan & Crump, 2011). We do not believe that all cognitive control is event driven and discrete, but it may be a useful alternative hypothesis in searching for more continuous hierarchical control. At present, we know the act of control triggered by a stop signal is event driven and discrete. Future research will address other acts and other kinds of control.

Limitations

There are two major limitations to our theory: We have no theory of drift rate, and we have no theory of how top-down parameters should be set. A theory of drift rate would allow us to separate structural factors, like similarity and discriminability (Logan, 2002; Nosofsky & Palmeri, 1997), from strategic factors, like attention and task set (Logan & Gordon, 2001; P. L. Smith &

Ratcliff, 2009). For now, drift rate is a free parameter. We allow it to vary to optimize the fit to the data. We do not constrain it with a theory that says how it is calculated. A theory of drift rate is an important step in the future development of special race models.

A theory that explains how top-down parameters are set would be an important step toward a theory of executive control. The diffusion race model tells us how much the threshold increases when stop signals become relevant, but that quantity reflects what is required to optimize the fit. It does not come from a theory that says what information the executive considers in deciding the value of the threshold and how the executive chooses a value based on that information. We need a theory that explains these phenomena. That is an important goal for future research.

Conclusions

The general and special independent race models extend the original independent race model in important directions. They account for choice, which is an important advance over previous stop-signal models (Boucher et al., 2007; Lo et al., 2009; Logan & Cowan, 1984; Salinas & Stanford, 2013; Wong-Lin et al., 2010). They predict RT distributions and response probability, so they can be compared with models of RT (Ratcliff & Smith, 2004; Teodorescu & Usher, 2013). They provide estimates of SSRT distributions, which have been elusive until recently (Matzke et al., 2013). They allow precise tests of hypotheses because of their commitment to stochastic accumulators. They frame hypotheses in terms of drift rates, thresholds, and nondecision times, and they test hypotheses with RT distributions and response probabilities. We implemented a special race model that assumed each runner was a diffusion process, and we fit it to two data sets to test hypotheses about selective influence of capacity limitations on rate parameters and strategies on threshold parameters. The model fit well, and the parameters behaved as expected.

References

Åkerfelt, A., Colonius, H., & Diederich, A. (2006). Visual-tactile saccadic inhibition. *Experimental Brain Research, 169,* 554–563. doi:10.1007/s00221-005-0168-x

Anderson, J. R., Bothell, D., Byrne, M. D., Douglass, S., Lebiere, C., & Qin, Y. (2004). An integrated theory of the mind. *Psychological Review, 111,* 1036–1060. doi:10.1037/0033-295X.111.4.1036

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,* 115–116. doi:10.1038/nn1003

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. doi:10.1523/JNEUROSCI.4682-05.2006

Ashby, F. G., & Townsend, J. T. (1986). Varieties of perceptual independence. *Psychological Review, 93,* 154–179. doi:10.1037/0033-295X.93.2.154

Balota, D. A., & Yap, M. J. (2011). Moving beyond the mean in studies of mental chronometry: The power of response time distributional analyses. *Current Directions in Psychological Science, 20,* 160–166. doi:10.1177/0963721411408885

Bedard, A. C., Nichols, S., Barbosa, J. A., Schachar, R., Logan, G. D., & Tannock, R. (2002). The development of selective inhibitory control across the life span. *Developmental Neuropsychology, 21,* 93–111. doi:10.1207/S15326942DN2101_5

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. doi:10.1037/a0021800

Bissett, P. G., & Logan, G. D. (2012). Post-stop-signal adjustments: Inhibition improves subsequent inhibition. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 38,* 955–966. doi:10.1037/a0026778

Bissett, P. G., & Logan, G. D. (2013). Selective stopping? Maybe not. *Journal of Experimental Psychology: General*. Advance online publication. doi:10.1037/a0032122

Bogacz, R., Usher, M., Zhang, J., & McClelland, J. L. (2007). Extending a biologically-inspired model of choice: Multi-alternatives, nonlinearity, and value-based multidimensional choice. *Philosophical Transactions of the Royal Society B, 362,* 1655–1670. doi:10.1098/rstb.2007.2059

Botvinick, M. M., Braver, T. S., Barch, D. M., Carter, C. S., & Cohen, J. D. (2001). Conflict monitoring and cognitive control. *Psychological Review, 108,* 624–652. doi:10.1037/0033-295X.108.3.624

Botvinick, M., & Plaut, D. C. (2004). Doing without schema hierarchies: A recurrent connectionist approach to normal and routine sequential action. *Psychological Review, 111,* 395–429. doi:10.1037/0033-295X.111.2.395

Botvinick, M., & Plaut, D. C. (2006). Such stuff as habits are made on: A reply to Cooper and Shallice (2006). *Psychological Review, 113,* 917–927. doi:10.1037/0033-295X.113.4.917

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. doi:10.1037/0033-295X.114.2.376

Brown, S. D., & Heathcote, A. (2005). A ballistic model of choice response time. *Psychological Review, 112,* 117–128. doi:10.1037/0033-295X.112.1.117

Brown, S. D., & Heathcote, A. (2008). The simplest complete model of choice reaction time: Linear ballistic accumulation. *Cognitive Psychology, 57,* 153–178. doi:10.1016/j.cogpsych.2007.12.002

Brown, S. D., Steyvers, M., & Wagenmakers, E.-J. (2009). Observing evidence accumulation during multi-alternative decisions. *Journal of Mathematical Psychology, 53,* 453–462. doi:10.1016/j.jmp.2009.09.002

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. doi:10.1016/j.neuroscience.2012.04.051

Bundesen, C. (1990). A theory of visual attention. *Psychological Review, 97,* 523–547.

Cavina-Pratesi, C., Bricolo, E., Prior, M., & Marzi, C. A. (2001). Redundancy gain in the stop-signal paradigm: Implications for the locus of coactivation in simple reaction time. *Journal of Experimental Psychology: Human Perception and Performance, 27,* 932–941. doi:10.1037/0096-1523.27.4.932

Chamberlain, S. R., & Sahakian, B. J. (2007). The neuropsychiatry of impulsivity. *Current Opinion in Psychiatry, 20,* 255–261.

Chambers, C. D., Garavan, H., & Belgrove, M. A. (2009). Insights into the neural basis of response inhibition from cognitive and clinical neuroscience. *Neuroscience and Biobehavioral Reviews, 33,* 631–646. doi:10.1016/j.neubiorev.2008.08.016

Chikazoe, J., Jimura, K., Hirose, S., Yamashita, K., Miyashita, Y., & Konishi, S. (2009). Preparation to inhibit a response complements response inhibition during performance of a stop-signal task. *Journal of Neuroscience, 29,* 15870–15877. doi:10.1523/JNEUROSCI.3645-09.2009

Cohen, J. D., Dunbar, K., & McClelland, J. L. (1990). On the control of automatic processes: A parallel distributed processing account of the

Stroop effect. *Psychological Review, 97,* 332–361. doi:10.1037/0033-295X.97.3.332

Colonius, H. (1990). A note on the stop-signal paradigm, or how to observe the unobservable. *Psychological Review, 97,* 309–312. doi:10.1037/0033-295X.97.2.309

Cooper, R., & Shallice, T. (2000). Contention scheduling and the control of routine activities. *Cognitive Neuropsychology, 17,* 297–338. doi:10.1080/026432900380427

Cooper, R. P., & Shallice, T. (2006). Hierarchical schemas and goals in the control of sequential behavior. *Psychological Review, 113,* 887–916. doi:10.1037/0033-295X.113.4.887

De Jong, R., Coles, M. G. H., Logan, G. D., & Gratton, G. (1990). In searching the point of no return: The control of response processes. *Journal of Experimental Psychology: Human Perception and Performance, 16,* 164–182. doi:10.1037/0096-1523.16.1.164

Dimitrov, M., Nakic, M., Elpern-Waxman, J., Granetz, J., O'Grady, J., Phipps, M., . . . Grafman, J. (2003). Inhibitory attentional control in patients with frontal lobe damage. *Brain and Cognition, 52,* 258–270. doi:10.1016/S0278-2626(03)00080-0

Donders, F. C. (1969). On the speed of mental processes. In W. G. Koster (Ed.), *Attention and performance II* (pp. 412–431). Amsterdam, the Netherlands: North-Holland. (Original work published 1868)

Dutilh, G., Vandekerckhove, J., Forstmann, B. U., Keuleers, E., Brysbaert, M., & Wagenmakers, E.-J. (2012). Testing theories of post-error slowing. *Attention, Perception, & Psychophysics, 74,* 454–465. doi:10.3758/s13414-011-0243-2

Eagle, D. M., & Robbins, T. W. (2003). Inhibitory control in rats performing a stop-signal reaction-time task: Effects of lesions of the medial striatum and *d*-amphetamine. *Behavioral Neuroscience, 117,* 1302–1317. doi:10.1037/0735-7044.117.6.1302

Eidels, A., Donkin, C., Brown, S. D., & Heathcote, A. (2010). Converging measures of workload capacity. *Psychonomic Bulletin & Review, 17,* 763–771. doi:10.3758/PBR.17.6.763

Emeric, E. E., Brown, J. W., Boucher, L., Carpenter, R. H. S., Hanes, D. P., Harris, R., . . . Schall, J. D. (2007). Influence of history on countermanding saccade performance in humans and macaque monkeys. *Vision Research, 47,* 35–49. doi:10.1016/j.visres.2006.08.032

Ersche, K. D., Jones, P. S., Williams, G. B., Turton, A. J., Robbins, T. W., & Bullmore, E. T. (2012, February 3). Abnormal brain structure implicated in stimulant drug addiction. *Science, 335,* 601–604. doi:10.1126/science.1214463

Forstmann, B. U., Anwander, A., Schäfer, A., Neumann, J., Brown, S., Wagenmakers, E.-J., . . . Turner, R. (2010). Cortico-striatal connections predict control over speed and accuracy in perceptual decision making. *PNAS: Proceedings of the National Academy of Sciences, USA, 107,* 15916–15920. doi:10.1073/pnas.1004932107

Forstmann, B. U., Dutilh, G., Brown, S., Neumann, J., von Cramon, D. Y., Ridderinkhof, K. R., & Wagenmakers, E.-J. (2008). Striatum and pre-SMA facilitate decision-making under time pressure. *PNAS: Proceedings of the National Academy of Sciences, USA, 105,* 17538–17542. doi:10.1073/pnas.0805903105

Forstmann, B. U., Wagenmakers, E.-J., Eichele, T., Brown, S., & Serences, J. T. (2011). Reciprocal relations between cognitive neuroscience and formal cognitive models: Opposites attract? *Trends in Cognitive Sciences, 15,* 272–279. doi:10.1016/j.tics.2011.04.002

Friedman, N. P., Miyake, A., Young, S. E., Defries, J. C., Corley, R. P., & Hewitt, J. K. (2008). Individual differences in executive functions are almost entirely genetic in origin. *Journal of Experimental Psychology: General, 137,* 201–225. doi:10.1037/0096-3445.137.2.201

Garner, W. R. (1962). *Uncertainty and structure as psychological concepts.* Oxford, England: Wiley.

Gold, J. I., & Shadlen, M. N. (2007). The neural basis of decision making. *Annual Review of Neuroscience, 30,* 535–574. doi:10.1146/annurev.neuro.29.051605.113038

Hanes, D. P., Patterson, W. F., II, & Schall, J. D. (1998). Role of frontal eye fields in countermanding saccades: Visual, movement, and fixation activity. *Journal of Neurophysiology, 79,* 817–834.

Hick, W. E. (1952). On the rate of gain of information. *Quarterly Journal of Experimental Psychology, 4,* 11–26. doi:10.1080/17470215208416600

Holroyd, C. B., & Coles, M. G. H. (2002). The neural basis of human error processing: Reinforcement learning, dopamine, and the error-related negativity. *Psychological Review, 109,* 679–709. doi:10.1037/0033-295X.109.4.679

Horstmann, G. (2003). The psychological refractory period of stopping. *Journal of Experimental Psychology: Human Perception and Performance, 29,* 965–981. doi:10.1037/0096-1523.29.5.965

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. doi:10.1016/j.neuropsychologia.2006.01.010

Hyman, R. (1953). Stimulus information as a determinant of reaction time. *Journal of Experimental Psychology, 45,* 188–196. doi:10.1037/h0056940

Kahneman, D. (1973). *Attention and effort.* Englewood Cliffs, NJ: Prentice-Hall.

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. doi:10.1037/0882-7974.9.4.491

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. doi:10.1037/h0021266

Leite, F. P., & Ratcliff, R. (2010). Modeling reaction time and accuracy of multiple-alternative decisions. *Attention, Perception, & Psychophysics, 72,* 246–273. doi:10.3758/APP.72.1.246

Lenartowicz, A., Verbruggen, F., Logan, G. D., & Poldrack, R. A. (2011). Inhibition-related activation in the right inferior frontal gyrus in the absence of inhibitory cues. *Journal of Cognitive Neuroscience, 23,* 3388–3399. doi:10.1162/jocn_a_00031

Leotti, L. A., & Wager, T. D. (2010). Motivational influences on response inhibition measures. *Journal of Experimental Psychology: Human Perception and Performance, 36,* 430–447. doi:10.1037/a0016802

Lo, C.-C., Boucher, L., Paré, M., Schall, J. D., & Wang, X.-J. (2009). Proactive inhibitory control and attractor dynamics in countermanding action: A spiking neural circuit model. *Journal of Neuroscience, 29,* 9059–9071. doi:10.1523/JNEUROSCI.6164-08.2009

Logan, G. D. (1978). Attention in character classification: Evidence for the automaticity of component stages. *Journal of Experimental Psychology: General, 107,* 32–63. doi:10.1037/0096-3445.107.1.32

Logan, G. D. (1979). On the use of a concurrent memory load to measure attention and automaticity. *Journal of Experimental Psychology: Human Perception and Performance, 5,* 189–207. doi:10.1037/0096-1523.5.2.189

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* (pp. 205–222). Hillsdale, NJ: Erlbaum.

Logan, G. D. (1982). On the ability to inhibit complex actions: A stop-signal study of typewriting. *Journal of Experimental Psychology: Human Perception and Performance, 8,* 778–792. doi:10.1037/0096-1523.8.6.778

Logan, G. D. (1985). Executive control of thought and action. *Acta Psychologica, 60,* 193–210.

Logan, G. D. (1988). Toward an instance theory of automatization. *Psychological Review, 95,* 492–527. doi:10.1037/0033-295X.95.4.492

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, CA: Academic Press.

Logan, G. D. (1996). The CODE theory of visual attention: An integration of space-based and object-based attention. *Psychological Review, 103,* 603–649. doi:10.1037/0033-295X.103.4.603

Logan, G. D. (2002). An instance theory of attention and memory. *Psychological Review, 109,* 376–400.

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. doi:10.1037/0096-1523.12.4.549

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. doi:10.1037/0033-295X.91.3.295

Logan, G. D., Cowan, W. B., & Davis, K. A. (1984). On the ability to inhibit responses in simple and choice reaction time tasks: A model and a method. *Journal of Experimental Psychology: Human Perception and Performance, 10,* 276–291. doi:10.1037/0096-1523.10.2.276

Logan, G. D., & Crump, M. J. C. (2011). Hierarchical control of cognitive processes: The case for skilled typewriting. In B. H. Ross (Ed.), *The psychology of learning and motivation* (Vol. 54, pp. 1–27). Burlington, VT: Academic Press.

Logan, G. D., & Etherton, J. L. (1994). What is learned in automatization? The role of attention in constructing an instance. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 20,* 1022–1050.

Logan, G. D., & Gordon, R. D. (2001). Executive control of visual attention in dual-task situations. *Psychological Review, 108,* 393–434. doi:10.1037/0033-295X.108.2.393

Logan, G. D., & Irwin, D. E. (2000). Don't look! Don't touch! Inhibitory control of eye and hand movements. *Psychonomic Bulletin & Review, 7,* 107–112.

Logan, G. D., Schachar, R. J., & Tannock, R. (1997). Impulsivity and inhibitory control. *Psychological Science, 8,* 60–64. doi:10.1111/j.1467-9280.1997.tb00545.x

Logan, G. D., & Zbrodoff, N. J. (1979). When it helps to be misled: Facilitative effects of increasing the frequency of conflicting stimuli in a Stroop-like task. *Memory & Cognition, 7,* 166–174.

Matzke, D., Dolan, C. V., Logan, G. D., Brown, S. D., & Wagenmakers, E.-J. (2013). Bayesian parametric estimation of stop-signal reaction time distributions. *Journal of Experimental Psychology: General, 142,* 1047–1073. doi:10.1037/a0030543

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. doi:10.3758/PBR.16.5.798

Merkel, J. (1885). Die zeitlichen Verhältnisse der Willensthätigkeit [The temporal circumstances of the will]. *Philosophische Studien, 2,* 73–127.

Meyer, D. E., & Keiras, D. E. (1997). A computational theory of executive cognitive processes and multiple-task performance: Part 1. Basic mechanisms. *Psychological Review, 104,* 3–65. doi:10.1037/0033-295X.104.1.3

Miyake, A., Friedman, N. P., Emerson, M. J., Witzki, A. H., Howerter, A., & Wager, T. (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. doi:10.1006/cogp.1999.0734

Myung, I. J. (2003). Tutorial on maximum likelihood estimation. *Journal of Mathematical Psychology, 47,* 90–100. doi:10.1016/S0022-2496(02)00028-7

Newell, A. (1990). *Unified theories of cognition.* Cambridge, MA: Harvard University Press.

Norman, D. A. (1981). Categorization of action slips. *Psychological Review, 88,* 1–15.

Nosofsky, R. M., Little, D. R., Donkin, C., & Fific, M. (2011). Short-term memory scanning viewed as exemplar-based categorization. *Psychological Review, 118,* 280–315.

Nosofsky, R. M., & Palmeri, T. J. (1997). An exemplar-based random walk model of speeded classification. *Psychological Review, 104,* 266–300. doi:10.1037/0033-295X.104.2.266

Ollman, R. T. (1973). Simple reactions with random countermanding of the "go" signal. In S. Kornblum (Ed.), *Attention and performance IV* (pp. 571–581). New York, NY: Academic Press.

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. doi:10.1037/0096-1523.12.3.243

Osman, A., Kornblum, S., & Meyer, D. E. (1990). Does response programming necessitate response execution? *Journal of Experimental Psychology: Human Perception and Performance, 16,* 183–198.

Paré, M., & Hanes, D. P. (2003). Controlled movement processing: Superior colliculus activity associated with countermanded saccades. *Journal of Neuroscience, 23,* 6480–6489.

Pashler, H. (1994). Dual-task interference in simple tasks: Data and theory. *Psychological Bulletin, 116,* 220–244. doi:10.1037/0033-2909.116.2.220

Posner, M. I. (1980). Orienting of attention. *Quarterly Journal of Experimental Psychology, 32,* 3–25.

Posner, M. I., & Boies, S. J. (1971). Components of attention. *Psychological Review, 78,* 391–408. doi:10.1037/h0031333

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. *Journal of Neuroscience, 31,* 12604–12612. doi:10.1523/JNEUROSCI.1868-11.2011

Purcell, B. A., Heitz, R. P., Cohen, J. Y., Schall, J. D., Logan, G. D., & Palmeri, T. J. (2010). Neurally constrained modeling of perceptual decision making. *Psychological Review, 117,* 1113–1143. doi:10.1037/a0020311

Raftery, A. E. (1995). Bayesian model selection in social research (with discussion). *Sociological Methodology, 25,* 111–196. doi:10.2307/271063

Ramautar, J. R., Kok, A., & Ridderinkhof, K. R. (2004). Effects of stop-signal probability in the stop-signal paradigm: The N2/P3 complex further validated. *Brain and Cognition, 56,* 234–252. doi:10.1016/j.bandc.2004.07.002

Ratcliff, R. (1978). A theory of memory retrieval. *Psychological Review, 85,* 59–108. doi:10.1037/0033-295X.85.2.59

Ratcliff, R. (2006). Modeling response signal and response time data. *Cognitive Psychology, 53,* 195–237. doi:10.1016/j.cogpsych.2005.10.002

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. doi:10.1152/jn.01049.2002

Ratcliff, R., & McKoon, G. (2008). The diffusion decision model: Theory and data for two-choice decision tasks. *Neural Computation, 20,* 873–922. doi:10.1162/neco.2008.12-06-420

Ratcliff, R., & Murdock, B. B., Jr. (1976). Retrieval processes in recognition memory. *Psychological Review, 83,* 190–214. doi:10.1037/0033-295X.83.3.190

Ratcliff, R., & Smith, P. L. (2004). A comparison of sequential sampling models for two-choice reaction time. *Psychological Review, 111,* 333–367. doi:10.1037/0033-295X.111.2.333

Ratcliff, R., Van Zandt, T., & McKoon, G. (1999). Connectionist and diffusion models of reaction time. *Psychological Review, 106,* 261–300. doi:10.1037/0033-295X.106.2.261

Rieger, M., & Gauggel, S. (1999). Inhibitory after-effects in the stop signal paradigm. *British Journal of Psychology, 90,* 509–518.

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. doi:10.1016/S0001-6918(99)00010-4

Rieger, M., & Gauggel, S. (1999). Inhibitory after-effects in the stop signal paradigm. *British Journal of Psychology, 90,* 509–518. doi:10.1348/000712699161585

Salinas, E., & Stanford, T. R. (2013). The countermanding task revisited: Fast stimulus detection is a key determinant of psychophysical performance. *Journal of Neuroscience, 33,* 5668–5685. doi:10.1523/JNEUROSCI.3977-12.2013

Schachar, R. J., Crosbie, J., Barr, C. L., Ornstein, T. J., Kennedy, J., Malone, M., . . . Pathare, T. (2005). Inhibition of motor responses in siblings concordant and discordant for attention deficit hyperactivity disorder. *American Journal of Psychiatry, 162,* 1076–1082. doi:10.1176/appi.ajp.162.6.1076

Schachar, R. J., & Logan, G. D. (1990). Impulsivity and inhibitory control in normal development and childhood psychopathology. *Developmental Psychology, 26,* 710–720. doi:10.1037/0012-1649.26.5.710

Schneider, D. W., & Anderson, J. R. (2011). A memory-based model of Hick's law. *Cognitive Psychology, 62,* 193–222. doi:10.1016/j.cogpsych.2010.11.001

Schwarz, G. E. (1978). Estimating the dimension of a model. *Annals of Statistics, 6,* 461–464. doi:10.1214/aos/1176344136

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.

Slater-Hammel, A. T. (1960). Reliability, accuracy, and refractoriness of a transit reaction. *Research Quarterly, American Association for Health, Physical Education and Recreation, 31,* 217–228.

Smith, M. C. (1969). The effect of varying information on the psychological refractory period. *Acta Psychologica, 30,* 220–231. doi:10.1016/0001-6918(69)90052-3

Smith, P. L., & Ratcliff, R. (2009). An integrated theory of attention and decision making in visual signal detection. *Psychological Review, 116,* 283–317. doi:10.1037/a0015156

Smith, P. L., & Van Zandt, T. (2000). Time-dependent Poisson counter models of response latency in simple judgment. *British Journal of Mathematical and Statistical Psychology, 53,* 293–315. doi:10.1348/000711000159349

Sternberg, S. (1969). The discovery of processing stages: Extensions of Donders' method. *Acta Psychologica, 30,* 276–315.

Stevens, M., Lammertyn, J., Verbruggen, F., & Vandierendonck, A. (2006). Tscope: A C library for programming cognitive experiments on the MS Windows platform. *Behavior Research Methods, 38,* 280–286. doi:10.3758/BF03192779

Stroop, J. R. (1935). Studies of interference in serial verbal reactions. *Journal of Experimental Psychology, 18,* 643–662.

Tannock, R., Schachar, R. J., Carr, R. P., Chajczyk, D., & Logan, G. D. (1989). Effects of methylphenidate on inhibitory control in hyperactive children. *Journal of Abnormal Child Psychology, 17,* 473–491. doi:10.1007/BF00916508

Teodorescu, A. R., & Usher, M. (2013). Disentangling decision models: From independence to competition. *Psychological Review, 120,* 1–38. doi:10.1037/a0030776

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. doi:10.1016/j.biopsych.2010.08.016

Townsend, J. T. (1971). A note on the identifiability of parallel and serial processes. *Perception & Psychophysics, 10,* 161–163. doi:10.3758/BF03205778

Townsend, J. T. (1990). Serial vs. parallel processing: Sometimes they look like Tweedledum and Tweedledee but they can (and should) be distinguished. *Psychological Science, 1,* 46–54. doi:10.1111/j.1467-9280.1990.tb00067.x

Townsend, J. T., & Altieri, N. (2012). An accuracy–response time capacity assessment function that measures performance against standard parallel predictions. *Psychological Review, 119,* 500–516. doi:10.1037/a0028448

Townsend, J. T., & Ashby, F. G. (1983). *Stochastic modeling of elementary psychological processes.* Cambridge, England: Cambridge University Press.

Townsend, J. T., Houpt, J. W., & Silbert, N. H. (2012). General recognition theory extended to include response times: Predictions for a class of parallel systems. *Journal of Mathematical Psychology, 56,* 476–494. doi:10.1016/j.jmp.2012.09.001

Townsend, J. T., & Wenger, M. J. (2004). A theory of interactive parallel processing: New capacity measures and predictions for a response time inequality series. *Psychological Review, 111,* 1003–1035. doi:10.1037/0033-295X.111.4.1003

Tzelgov, J., Henik, A., & Leiser, D. (1990). Controlling Stroop interference: Evidence from a bilingual task. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 16,* 760–771.

Usher, M., & McClelland, J. L. (2001). The time course of perceptual choice: The leaky, competing accumulator model. *Psychological Review, 108,* 550–592. doi:10.1037/0033-295X.108.3.550

Usher, M., Olami, Z., & McClelland, J. L. (2002). Hick's law in a stochastic race model with speed–accuracy tradeoff. *Journal of Mathematical Psychology, 46,* 704–715. doi:10.1006/jmps.2002.1420

van Ravenzwaaij, D., Brown, S., & Wagenmakers, E.-J. (2011). An integrated perspective on the relation between response speed and intelligence. *Cognition, 119,* 381–393. doi:10.1016/j.cognition.2011.02.002

Van Selst, M., & Jolicoeur, P. (1997). Decision and response in dual-task interference. *Cognitive Psychology, 33,* 266–307. doi:10.1006/cogp.1997.0662

Van Zandt, T. (2000a). How to fit a response time distribution. *Psychonomic Bulletin & Review, 7,* 424–465. doi:10.3758/BF03214357

Van Zandt, T. (2000b). ROC curves and confidence judgments in recognition memory. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 26,* 582–600. doi:10.1037/0278-7393.26.3.582

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. doi:10.3758/BF03212980

Verbruggen, F., Adams, R., & Chambers, C. D. (2012). Proactive motor control reduces monetary risk taking in gambling. *Psychological Science, 23,* 805–815. doi:10.1177/0956797611434538

Verbruggen, F., Aron, A. R., Stevens, M. A., & Chambers, C. D. (2010). Theta burst stimulation dissociates attention and action updating in human inferior frontal cortex. *PNAS: Proceedings of the National Academy of Sciences, USA, 107,* 13966–13971. doi:10.1073/pnas.1001957107

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. doi:10.1177/0956797612457390

Verbruggen, F., Liefooghe, B., Szmalec, A., & Vandierendonck, A. (2005). Inhibiting responses when switching: Does it matter? *Experimental Psychology, 52,* 125–130. doi:10.1027/1618-3169.52.2.125

Verbruggen, F., Liefooghe, B., & Vandierendonck, A. (2004). The interaction between stop signal inhibition and distractor interference in the

flanker and Stroop task. *Acta Psychologica, 116,* 21–37. doi:10.1016/j.actpsy.2003.12.011

Verbruggen, F., Liefooghe, B., & Vandierendonck, A. (2005). On the difference between response inhibition and negative priming: Evidence from simple and selective stopping. *Psychological Research, 69,* 262–271. doi:10.1007/s00426-004-0177-6

Verbruggen, F., Liefooghe, 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. doi:10.1080/17470210500151386

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. doi:10.1037/a0013170

Verbruggen, F., & Logan, G. D. (2008b). Long-term aftereffects of response inhibition: Memory retrieval, task goals, and cognitive control. *Journal of Experimental Psychology: Human Perception and Performance, 34,* 1229–1235. doi:10.1037/0096-1523.34.5.1229

Verbruggen, F., & Logan, G. D. (2008c). Response inhibition in the stop-signal paradigm. *Trends in Cognitive Sciences, 12,* 418–424. doi:10.1016/j.tics.2008.07.005

Verbruggen, F., & Logan, G. D. (2009a). Automaticity of cognitive control: Goal priming in response-inhibition paradigms. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 35,* 1381–1388. doi:10.1037/a0016645

Verbruggen, F., & Logan, G. D. (2009b). Models of response inhibition in the stop-signal and stop-change paradigms. *Neuroscience & Biobehavioral Reviews, 33,* 647–661. doi:10.1016/j.neubiorev.2008.08.014

Verbruggen, F., & Logan, G. D. (2009c). Proactive adjustments of response strategies in the stop-signal paradigm. *Journal of Experimental Psychology: Human Perception and Performance, 35,* 835–854. doi:10.1037/a0012726

Verbruggen, F., Logan, G. D., Liefooghe, B., & Vandierendonck, A. (2008). Aftereffects of response inhibition: Repetition priming or between-trial control adjustments? *Journal of Experimental Psychology:*

Human Perception and Performance, 34, 413–426. doi:10.1037/0096-1523.34.2.413

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. doi:10.3758/BRM.40.2.479

Vince, M. A. (1948). The intermittency of control movements and the psychological refractory period. *British Journal of Psychology, 38,* 149–157.

Voss, A., & Voss, J. (2007). Fast-dm: A free program for efficient diffusion model analysis. *Behavior Research Methods, 39,* 767–775. doi:10.3758/BF03192967

Wagenmakers, E.-J. (2007). A practical solution to the pervasive problems of *p* values. *Psychonomic Bulletin & Review, 14,* 779–804. doi:10.3758/BF03194105

Wagenmakers, E.-J., Ratcliff, R., Gomez, P., & McKoon, G. (2008). A diffusion model account of criterion shifts in the lexical decision task. *Journal of Memory and Language, 58,* 140–159. doi:10.1016/j.jml.2007.04.006

Wenger, M. J., & Townsend, J. T. (2000). Basic response time tools for studying general processing capacity in attention, perception, and cognition. *Journal of General Psychology, 127,* 67–99. doi:10.1080/00221300009598571

Williams, B. R., Ponesse, J. S., Schachar, R. J., Logan, G. D., & Tannock, R. (1999). Development of inhibitory control across the life span. *Developmental Psychology, 35,* 205–213. doi:10.1037/0012-1649.35.1.205

Wong-Lin, K., Eckhoff, P., Holmes, P., & Cohen, J. D. (2010). Optimal performance in a countermanding saccade task. *Brain Research, 1318,* 178–187. doi:10.1016/j.brainres.2009.12.018

Xue, G., Aron, A. R., & Poldrack, R. A. (2008). Common neural substrates for inhibition of spoken and manual responses. *Cerebral Cortex, 18,* 1923–1932. doi:10.1093/cercor/bhm220

Yamaguchi, M., Logan, G. D., & Bissett, P. G. (2012). Stopping while going! Response inhibition does not suffer dual-task interference. *Journal of Experimental Psychology: Human Perception and Performance, 38,* 123–134. doi:10.1037/a0023918

Appendix A

Experiment 1 Method and Results

Method

Subjects

Six people from the Vanderbilt University (Nashville, TN) community participated for monetary compensation ($12/hour). All subjects had normal or corrected-to-normal vision, and all were naive as to the purpose of the experiment.

Apparatus and Stimuli

The experiment was run on a Pentium 4 PC running Tscope (Stevens, Lammertyn, Verbruggen, & Vandierendonck, 2006). The stimuli were presented on a 21-in. cathode ray tube monitor. In the go task, subjects indicated which word in a set of two, four, or six was presented by pressing one key on a QWERTY keyboard. The words were presented centrally in a white uppercase Courier font (size = 24) on a black background. In the two-choice condition, subjects pressed *F* for one word and *J* for the other word; in the four-choice condition, subjects pressed *D, F, J,* or *K*; in the six-choice condition, subjects pressed *S, D, F, J, K,* or *L*. We used different words in each session and in each condition to reduce practice effects. In all conditions, the mapping of words onto response keys was randomized. On stop-signal trials, a loud and clear auditory signal (80 dB, 100 ms, 500 Hz) was presented through closed headphones (Sennheiser eH150), using stop-signal presentation functions of STOP-IT (Verbruggen, Logan, & Stevens, 2008).

Procedure

The experiment consisted of 12 sessions. At the beginning of the first session, instructions were given orally by the experimenter. Subjects were instructed to respond as quickly and accurately as possible. In each session, there were three conditions: two-choice, four-choice, and six-choice. Each condition started with the presentation of the words that could occur in the block and the word-key mapping. The order of the conditions within each session was randomized.

Table A1
Summary Tables for One-Way Analyses of Variance Evaluating the Effects of Number of Choice Alternatives on Go Response Times, Signal-Respond Response Times, Stop-Signal Response Times, and Error Rates in Experiment 1

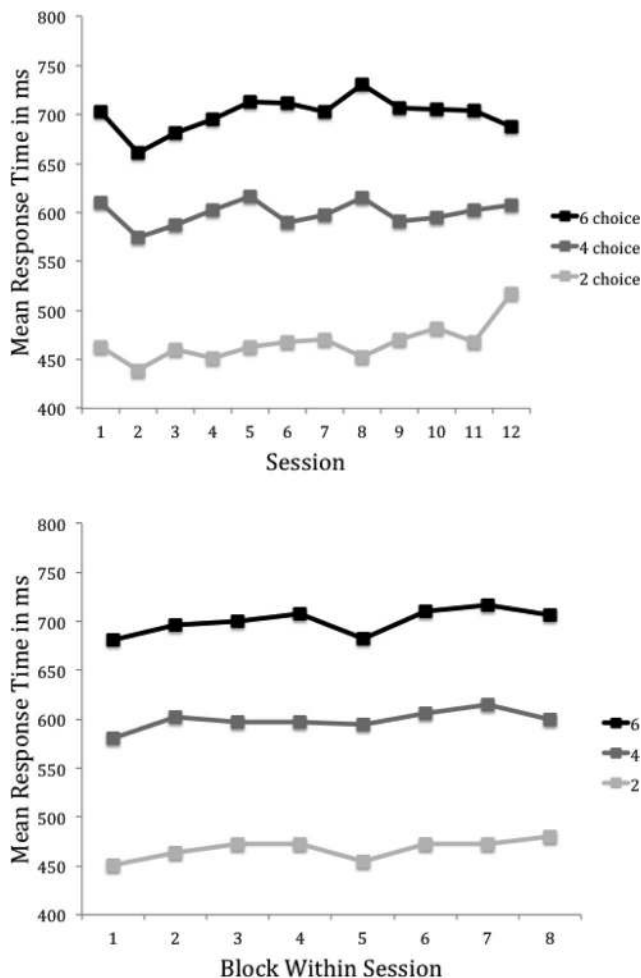| Dependent variable	F ratio	Degrees of freedom	*MSE*	p
Go response time	28.0	2, 10	2,920	<.001
Signal-respond response time	37.7	2, 10	1,344	<.001
Stop-signal response time	2.2	2, 10	305	>.16
Error rate	27.5	2, 10	0.00008	<.001

Figure A1. Mean response time for no-stop-signal trials as a function of number of choice alternatives and session (top panel) and block within session (bottom panel).

All trials started with the presentation of a fixation sign (+ + + + + +), which was replaced by a single word after 250 ms. The word remained on the screen for 2,000 ms, regardless of response time (RT). The following trial started 250 ms after the offset of the word. On 25% of the trials, a stop signal was presented at a variable delay after the onset of the go stimulus. The delay was based on performance in the signal practice block in Session 1 (see below). In each choice condition, stop-signal delay was 15%, 35%, 55%, 75%, or 95% of the corresponding mean RT in the practice phase. Delays were randomized and occurred with equal probability.

(*Appendices continue*)

Each choice condition started with a practice block of 36 trials without stop signals. This no-stop-signal block was followed by another practice block of 36 trials with stop signals. After the two practice blocks, there were two experimental blocks of 120 trials. At the end of each block, we presented the number of no-stop-signal errors, the mean RT, and the probability of stopping. Subjects had to pause for 10 s between each block.

Results

Mean no-stop-signal RTs, signal-respond RT, stop-signal RTs (SSRTs), and error rates, collapsed across subjects and sessions, appear in Figure 3 in the main text. SSRTs were estimated using the integration method (Logan, 1994; Logan & Cowan, 1984; Verbruggen & Logan, 2009b). For each number of choice alternatives, the no-stop-signal RTs were rank-ordered, and the *Mth* RT was selected, where *M* was obtained by multiplying the number of RTs in the distribution by *P*(respond|signal) at a given delay. To estimate SSRT, stop-signal delay (SSD) was subtracted from the *Mth* RT. This process was repeated for each SSD for each subject.

The results were then averaged across SSDs for which *P*(respond|signal) was larger than .05 but smaller than .95.

We subjected the go RTs, signal-respond RT, SSRTs, and error rates to separate one-way analyses of variance (ANOVAs) with number of choice alternatives (two, four, six) as within-subject factor. Summary tables for these ANOVAs appear in Table A1. Number of choice alternatives affected go RT, signal-respond RT, and error rate but did not affect SSRT.

The top panel of Figure A1 plots go RTs from no-stop-signal trials for two, four, and six choices as a function of session. The figure shows that performance was quite stable across sessions, most likely because we changed the words in each choice set each session. The bottom panel of Figure A1 plots go RTs from no-stop-signal trials for two, four, and six choices as a function of block within a session. The figure shows that performance was quite stable across blocks within a session, most likely because the 72 trials of practice at the beginning of each choice condition absorbed initial changes in performance due to learning.

Appendix B

Fitting Models to Data

To understand how the models were fit to the data, we must expand the simplified notation used in the text. Each subject provides response times (RTs) in each condition c (two, four, or six choices in the multiple-choice experiment; *none* vs. *all* conditions in Verbruggen & Logan, 2009c) under different stop-signal delays t_d. A trial can be either a go trial or a stop trial. To simplify exposition, we will set stop-signal delay t_d equal to infinity for go trials. On any trial, a response may be correct, incorrect, or inhibited, and the RT of that response will depend on a set of parameters θ_c that are appropriate for condition c and stop-signal delay t_d. We can then write the finishing-time distribution of response i under condition c as $f_i(t|\theta_c, t_d)$. We write an observed RT from an individual subject on trial j in Experiment 1 as T_{jlc}. We also define $R_{jlc} = 1$ if a response is made on trial j (if $T_{jlc} > 0$) and 0 otherwise. Let $A_{kjlc} = 1$ if the response k is correct on trial j and 0 otherwise, so that a number (two, four, or six) of A_{kjlc}s are defined for each trial j. An individual subject's data can then be written as the set of vectors $\{T_c, R_c, \{A_c\}\}$.

The models were fit to the data in the multiple-choice task and Verbruggen and Logan (2009c) using maximum likelihood (Myung, 2003; Van Zandt, 2000a). The likelihood is a function that reflects the probability of the data given a set of parameter values. We define the likelihood by way of Equations 8–12 in the main text, the probability of a response $P_r(t_d)$, and the joint probability $f(t; i)$ of response i at time t. Expanding the notation of Equations 8–12, we write the probability of a response on a stop trial as $P_r(t_d|\theta_{jlc})$ and the joint probability of response i at time t as

$f_{RT}(t; i|\theta_{jlc}, t_d)$, making explicit the dependence of these probabilities on the model parameters θ_{jlc} on trial j and the stop-signal delay t_d on trial j.

For the diffusion race model in the multiple-choice experiment and the Verbruggen and Logan (2009c) experiment, the parameters θ_{jlc} for condition c include the thresholds z_{glc} and z_{slc} for the go and stop processes, respectively; the correct drift rates ξ_{glc} and ξ_{slc} for the go and stop processes, respectively; the incorrect drift rate ε_c for the go process; and the nondecision times t_g and t_s for the go and stop processes, respectively. If response i is observed on trial j at time $T_{jlc} = t$, the joint likelihood of that response is

$$f(t|t_d, \theta_{jlc}) = f_i(t - t_g|z_{glc}, A_{ijlc}\xi_{glc} + (1 - A_{ijlc})\varepsilon_c)$$
$$\times (1 - F_{stop}(t - t_s - t_d|z_{slc}, \xi_{slc}))$$
$$\times \prod_{k \neq i} (1 - F_k(t - t_g|z_{glc}, A_{kjlc}\xi_{glc} + (1 - A_{kjlc})\varepsilon_c)).$$

The likelihood of an inhibited response on trial j is

$$1 - P_r(t_d) = \int_0^\infty f_{stop}(t - t_s - t_d|z_{slc}, \xi_{slc})$$
$$\times \prod_k (1 - F_k(t - t_g|z_{glc}, A_{kjlc}\xi_{glc} + (1 - A_{kjlc})\varepsilon_c))dt.$$

The likelihood for the model, for the data from condition c and stop-signal delay t_d, is

$$L(\theta_c|T_c, R_c, A_c, t_d) = \prod_{j \in c} (1 - P_r(t_d))^{1 - R_{jlc}} f(t_j|t_d, \theta_{jlc})^{R_{jlc}}.$$

(*Appendices continue*)

The total likelihood across all conditions in the experiment is

$$L(\theta|T, R, A) = \prod_c \prod_d L(\theta_c|T_c, R_c, A_c, t_d).$$

Model fits were obtained by minimizing the negative log total likelihood $-\log L(\theta|T, R, A)$ over all conditions and stop-signal delays simultaneously. We used the Nelder-Mead simplex algorithm as programmed in R's optim function. We obtained starting values by minimizing $-\log L(\theta_c|T_c, R, A_c)$ (multiple-choice experiment) or $-\log L(\theta_c|T_c, R, A_c)$ (Verbruggen & Lo-

gan, 2009c) separately for each condition c. We then used those values to minimize $-\log L(\theta|T, R, A)$ over all conditions, repeating until there was no further change in the final values of $-\log L(\theta|T, R, A)$ or $-\log L(\theta|T, R, A, P)$. Outliers faster than 150 ms were excluded but were very rare. Observations that returned likelihoods of 0 (log likelihoods of $-\infty$) were given log likelihoods of -750 (a value orders of magnitude larger than those obtained when the parameters assumed reasonable values). Nondecision time was bounded between 0 and the minimum RT.

Appendix C

Diffusion Race Model Can Produce Fast Errors

The best fitting diffusion race models predicted error response times (RTs) that were slower than observed error RTs and often slower than correct RTs. We attributed this to the low frequency of errors (2% or less), so error RTs had little influence on the fitting routine. Here, we report a simulation of the diffusion race model that produces fast errors, and we show that the fits of the diffusion race model to the simulated data predict fast errors and recover the parameters of the simulation.

We began with the parameters of the best fitting model in the aggregate fits, in which go threshold and rate varied with the number of choice alternatives but stop threshold and rate did not (see Table 2 in the main text). We increased error rate by reducing the difference in drift rate between correct and incorrect response alternatives to simulate the effect of reducing the discriminability of the choice alternatives in the go task. We produced fast errors by fixing threshold variability to equal 75% of the threshold. We simulated 8,640 trials to produce data similar to individual-subject data in the multiple-choice experiment. The simulated (observed) RT distributions for correct and error responses are plotted in Figure C1. The simulated mean RTs for correct and error responses and simulated error probabilities are presented in Table C1. The parameters used to generate the simulated data are presented in Table C2.

Overall, the simulated error RTs were 115 ms faster than the simulated correct RTs. Figure C1 shows that the distributions of

error RTs were consistently faster than the distributions of correct RTs. This simulation demonstrates that the diffusion race model can predict fast errors. The simulated error rate was 0.36.

We fit a diffusion race model, in which go rate and threshold varied with choice alternatives but stop rate and threshold did not, to the simulated data. The predicted distributions are plotted with the observed distributions in Figure C1. The means for correct RT, incorrect RT, and error rate appear in Table C1. The best fitting parameter values appear in Table C2. The model fit the simulated data well, recovering the parameters accurately and predicting error RTs that were faster than correct RTs. The fit demonstrates that the model can fit fast errors when error rate is high (0.36). Threshold variability was fixed in the simulation that generated the data, but it was a free parameter in the fits. The fits produced estimates of threshold variability that were close to those in the simulated data (see Table C1), indicating that our model fits can recover threshold variability well when there are enough errors to contribute significantly to the likelihood.

The simulation and the fit demonstrate in principle that the diffusion race model can generate and fit fast errors when error rate is high. Whether the diffusion race model can fit error RTs in real data with high error rate is a question we will address in future research.
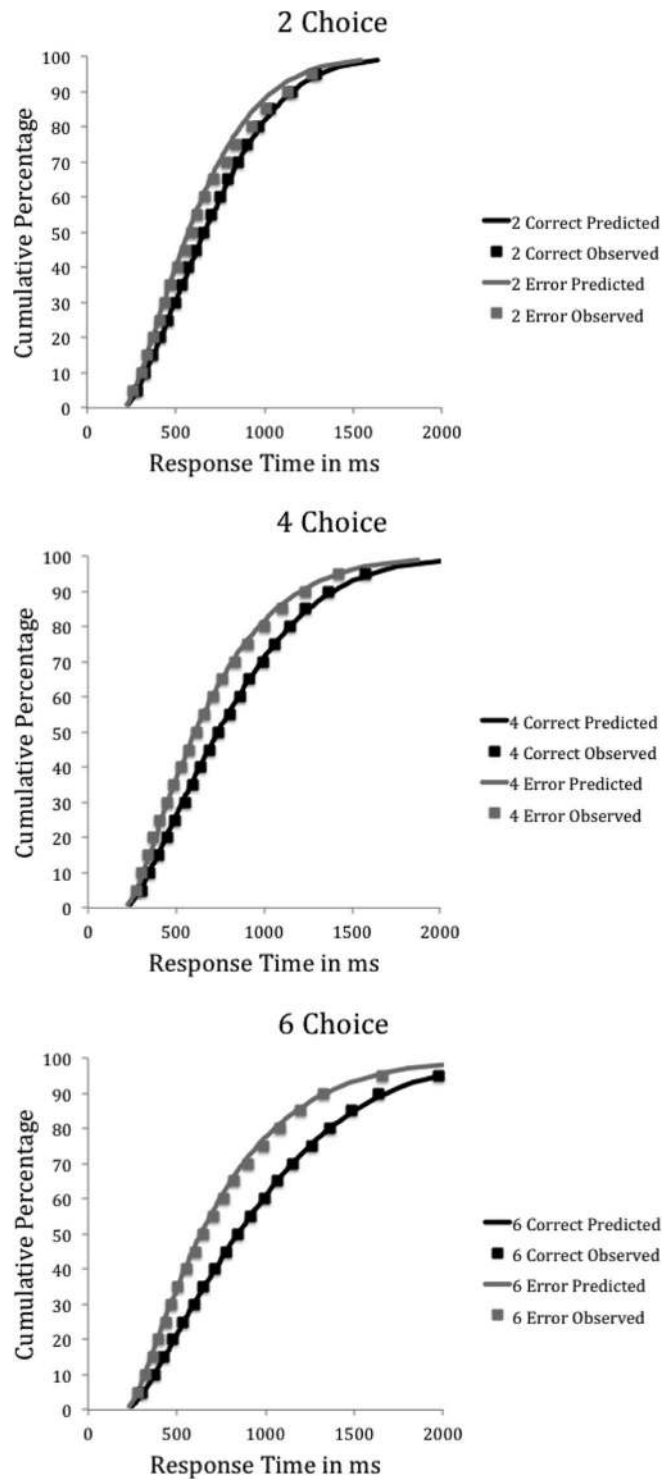
(Appendices continue)

Figure C1. Observed (simulated) and predicted response time distributions for correct and error responses from a diffusion race model with limited capacity in the go task and unshared capacity in the stop task for two-choice (top panel), four-choice (middle panel), and six-choice (bottom panel) alternatives.

(*Appendices continue*)

Table C1

Mean Correct and Incorrect Response Times, Error Rates, and Error Rates as a Function of Number of Choice Alternatives (2, 4, 6) for Observed (Simulated) and Predicted (Fitted) Data Generated From a Diffusion Race Model With a Limited Capacity Go Process and an Unshared Capacity Stop Process

Measure	2		4		6	
	Predicted	Observed	Predicted	Observed	Predicted	Observed
Correct response time	707	702	811	812	945	941
Error response time	627	652	692	701	742	758
P(Error)	0.206	0.204	0.411	0.410	0.450	0.447

Table C2

Parameters From a Diffusion Race Model With a Limited-Capacity Go Process and an Unlimited-Capacity Stop Process Used to Generate Simulated Data (Observed) and Parameters (Predicted) From Fitting a Model With the Same Structure to the Simulated Data

Choice	Go threshold	Threshold variability	Correct rate	Incorrect rate	Nondecision time	Stop threshold	Threshold variability	Stop rate	Nondecision time
				Observed (simulated)					
2	64.000	48.000	0.105	0.040	132	5.000	3.750	0.070	164
4	67.000	50.250	0.080	0.020	132	5.000	3.750	0.070	164
6	69.000	51.750	0.065	0.004	132	5.000	3.750	0.070	164
				Best fit					
2	61.363	45.750	0.101	0.036	156	9.816	0.000	0.0904	193
4	65.347	48.721	0.078	0.019	156	9.816	0.000	0.0904	193
6	68.646	51.181	0.066	0.004	156	9.816	0.000	0.0904	193