Paradox resolved: stop signal race model with negative dependence

Hans Colonius^{a,c,1} and Adele Diederich^{b,1,2}

^aOldenburg University; ^bJacobs University Bremen

This manuscript was compiled on February 18, 2018

The ability to inhibit our responses voluntarily is an important case of cognitive control. The stop-signal paradigm is a popular tool to study response inhibition. Participants perform a response time task (*go task*) and, occasionally, the go stimulus is followed by a stop signal after a variable delay, indicating subjects to withhold their response (*stop task*). The main interest of modeling is in estimating the unobservable stop-signal processing time, that is, the covert latency of the stopping process as a characterization of the response inhibition mechanism. In the *independent race model* the stop-signal task is represented as a race between stochastically independent go and stop processes. Without making any specific distributional assumptions about the processing times, the model allows to estimate the mean time to cancel a response. However, neurophysiological studies on countermanding saccadic eye movements have shown that neural correlates of go and stop processes consist of networks of mutually *interacting* gaze-shifting and gaze-holding neurons. This poses a major challenge in formulating linking propositions between the behavioral and neural findings. Here we propose a *dependent race model* that postulates perfect negative stochastic dependence between go and stop activations. The model is consistent with the concept of interacting processes while retaining the simplicity and elegance of the distribution-free independent race model. For mean data, the dependent model's predictions remain identical to those of the independent model. The resolution of this apparent paradox advances the understanding of mechanisms of response inhibition and paves the way for modeling more complex situations.

response inhibition | stop signal paradigm | independent race model | perfect negative dependence

A recurrent theme in cognitive modeling is the difficulty to uniquely identify the processes underlying the generation of response times and probabilities in behavioral paradigms like simple yes-no tasks or same-different judgments of stimulus pairs. A prime example is the serial vs. parallel processing issue (1–3) showing that even rigorous mathematical analyses of the underlying assumptions may not always resolve the difficulty completely. Recently, efforts in model-based cognitive neuroscience to constrain behavioral models by findings obtained from neuroscientific methods have intensified, from spike train analyses to EEG and fMRI recordings (4). One area where important advances have emerged from this joint endeavor is the modeling of cognitive control and, in particular, response inhibition (5, 6). Response inhibition refers to the ability to suppress responses that are no longer required or have become inappropriate; it is important for survival, such as stopping yourself from crossing the street when a car comes around the corner without noticing you. Deficits of response inhibition have been linked to several disorders like attention-deficit/hyperactivity, obsessive-compulsive behavior and substance abuse. Probing response inhibition has been used widely to study executive control and flexibility in behavior; for reviews, see (6–9) and a recent theme issue of *Philosophical Transactions of the Royal Society B* (2017) on ‘Movement suppression: brain mechanisms for stopping and stillness’ (10).

In the laboratory, a very useful tool for the study of inhibition is the *stop-signal paradigm* where participants perform a response time task (*go task*), such as moving their gaze to the location of a pre-defined target. Occasionally, the go stimulus is followed by a stop signal after a variable time delay, indicating subjects to withhold the response (*stop task*). Performance in the stop-signal paradigm has been modeled as a race between a “go process”, triggered by the presentation of the go stimulus, and a “stop process” triggered by the presentation of the stop signal (6, 8). When the stop process finishes before the go process, the response is inhibited; otherwise, it is executed. The main interest of the modeler is in estimating the unobservable stop-signal reaction time (SSRT), that is, the latency of the stopping process as a characterization of the response inhibition mechanism. In the *independent race model* (IND model, for short) the stop-signal task is represented as a race between *stochastically independent* go and stop processes (11). Under certain simplifying assumptions, mean SSRT can be estimated efficiently without making any specific assumptions about the distribution of the processing times (6). In several hundreds of studies, the IND model has been applied in virtually every stop-signal experiment providing important measures of cognitive control like SSRT. Although the model makes no commitment to the underlying computational or neural processes that generate the go and stop processing times, the IND model is considered as defining constraints that any model of response inhibition must follow (12).

However, neurophysiological studies in the frontal eye fields (FEF) and superior colliculus (SC) of macaque monkeys performing a countermanding task with saccadic eye movements have shown that the neural correlates of go and stop processes produce eye movement behavior through a network of *interacting* gaze-shifting and gaze-holding neurons (13–16). This discrepancy between the behavioral and neural data is widely perceived as a paradox (9, 12, 17, 18): how can interacting circuits of mutually inhibitory neurons instantiate stop and go processes with stochastically independent finishing times?

Here we propose a variant of the race model which, instead of independence, assumes *perfect negative stochastic dependence* between go and stop processes (PND model, for short). It resolves the apparent paradox and nonetheless retains the distribution-free property of the independent race model. Moreover, the PND model's predictions, considered at the level of mean SSRT, are shown to be necessarily identical to those of the independent model. Notably, we argue that it is very difficult to empirically

distinguish between the two race model versions without introducing further, parametric assumptions. Thus, there is no reason to uphold the stochastic independence assumption of the race model whenever a distribution-free version is to be considered.

Neurally inspired modeling

Investigating the neural underpinnings of response inhibition in saccades, Hanes and Schall (13) first showed that macaque monkey behavior in saccade countermanding corresponded in detail to that of human performance in manual stop-signal tasks. Then, recording from the frontal eye fields they isolated neurons involved in gaze-shifting and gaze-holding that represent a larger circuit of such neurons that extends from cortex through basal ganglia and superior colliculus to brainstem (14). Importantly, that result was based on their postulate that for neurons to participate in controlling movement initiation two criteria must be met: first, neurons must discharge differently when movements are initiated or withheld; if neurons still discharge when movements are canceled, their activity was not affected by the stop process. Second, the differential modulation on canceled trials must occur before SSRT; otherwise, the neural modulation happens after the movement has already been canceled (18, p.1012). From these findings, Boucher and colleagues (17) developed an *interactive race model* linking the interacting circuits of mutually inhibitory gaze-holding and gaze-shifting neurons with stochastic accumulation in the go and stop processes of the race model (see also 19). A response is observed only if the go process reaches a certain threshold of activation. Stopping occurs if the stop process interferes with the go process by inhibiting activation in the go accumulator to prevent it from reaching the threshold. Alternative models assume that response inhibition results from blocking the input to the go unit (“blocked-input models”) (20) or postulate a spiking neural network of hundreds of units representing populations of movement neurons, fixation neurons and inhibitory interneurons, and a control unit that turns the fixation neurons on and off (21).

These neural models are computationally explicit: they are based on systems of stochastic differential equations; estimating parameters to best fit the behavioral data, they are able to predict the distribution of cancel times, i.e., the times at which neural activity modulates on trials on which subjects stop successfully, relative to SSRT (12). The models fit the behavioral data just as well as the independent race model. Note that this is just another instantiation of the paradox. An attempt to resolve the contradiction between stochastic independence for behavioral data and interdependence at the neural level has been to postulate that the stop process is independent of the go process for much of its duration, followed by a late and potent interaction between stopping and going that reverses the trajectory of go activation (12, 17). Strictly speaking, this would require introducing two subsequent processing stages in the IND model, independence followed by strong interaction. No formal modeling of this extension has been undertaken, however, to the best of our knowledge.

Background: General race model

Next, we present a formal framework for the general race model making no assumption at all about the dependency between stop and go processes. One should distinguish between two different experimental conditions termed context $\mathcal{GO}$ where only a go signal is presented, and context $\mathcal{STOP}$ where a stop signal is presented in addition. A race between processes triggered by the go and stop signal is represented by two random variables: T_{go} and T_{stop} (referred to as SSRT above) denote the random processing times for the go and stop signal, respectively, in context $\mathcal{STOP}$ with a bivariate distribution function denoted H ,

$$H(s, t) = \Pr[T_{go} \leq s, T_{stop} \leq t], \quad [1]$$

defined for all real numbers s and t , with $s, t \geq 0$. The marginal distributions of $H(s, t)$ are

$$F_{go}(s) = \Pr[T_{go} \leq s, T_{stop} < \infty] \text{ and} \\ F_{stop}(t) = \Pr[T_{go} < \infty, T_{stop} \leq t].$$

In context $\mathcal{STOP}$ the go signal triggers a realization of random variable T_{go} and the stop signal triggers a realization of random variable T_{stop} . In context $\mathcal{GO}$, however, only processing of the go signal occurs. In general, the latter one may be different from the marginal distribution $F_{go}(s)$ in context $\mathcal{STOP}$. However, the *general race model* rules this out by adding the important assumption of “context invariance”, also known as “context independence” (e.g. 9):

Significance Statement

The ability to suppress a response that is no longer required or has become inappropriate is an important act of cognitive control. Deficits in inhibiting a response to a stop signal have been linked to disorders like attention-deficit/hyperactivity. A widely accepted model for the stop signal paradigm assumes that stop signal processing occurs independently from go signal processing and permits to estimate the covert latency of the stop process. However, neurophysiological data suggest interactive neural networks of going and stopping raising a contradiction between the behavioral and neural findings. Introducing a model with strong negative stochastic dependency between go and stop processing resolves the paradox while retaining the computational simplicity of the original model.

Authors declare no conflict of interest.

¹H.C. and A.D. contributed equally to this work.

²To whom correspondence should be addressed. E-mail: hans.colonius@uni-oldenburg.de

249 **Context invariance (CI).** In context $\mathcal{GO}$, the distribution of go signal processing time is assumed to be 311
250
$$F_{go}(s) = \Pr[T_{go} \leq s, T_{stop} < \infty],$$
 312
251 313
252 i.e., it is identical to the marginal distribution $F_{go}(s)$ in context $\mathcal{STOP}$. 314
253 From these assumptions, the probability of observing a response to the go signal, given a stop signal is presented with delay 315
254 t_d [ms] ($t_d \geq 0$) after the go signal, is defined by 316
255 317
256
$$p_r(t_d) = \Pr[T_{go} < T_{stop} + t_d].$$
 318
257 319
258 According to the model, the probability of observing a response to the go signal no later than time t , given the stop signal was 320
259 presented with delay t_d , is given by (conditional) distribution function 321
260
$$F_{sr}(t | t_d) = \Pr[T_{go} \leq t | T_{go} < T_{stop} + t_d],$$
 322
261 323
262 known as *signal-response RT* distribution. 324
263 The main interest in race modeling is to obtain information about the distribution of the unobservable stop signal processing 325
264 time T_{stop} , or about some of its parameters, given sample estimates of $F_{go}(t)$, $F_{sr}(t | t_d)$, and $p_r(t_d)$. As observed in (11), 326
265 letting stop signal delay t_d vary, the *inhibition function* $p_r(t_d)$ can be formally considered as the distribution function of a 327
266 random variable, $T_d \equiv T_{go} - T_{stop}$, say. Then, 328
267 329
268
$$E[T_d] = E[T_{go}] - E[T_{stop}],$$
 330
269 331
270 with $E[\cdot]$ denoting expected value (mean) of a random variable. Solving for the estimate of $E[T_{stop}]$ immediately yields an 332
271 estimate of the mean of the unobservable distribution of T_{stop} . It is well known that the reliability of this estimation method, 333
272 known as the *mean method*, depends on how precise the inhibition function $p_r(t_d)$ and the mean of T_{go} are estimated (6, 9). 334
273 Obtaining estimates of higher moments, or the entire distribution of T_{stop} , requires further non-parametric assumptions about 335
274 the bivariate distribution $H(s, t)$. 336
275 **Independent vs. negatively dependent race model** 337
276 338
277 The most common version of the race model, as introduced by Logan & Cowan (11), postulates stochastic independence 339
278 between T_{go} and T_{stop} : 340
279 **Stochastic independence:** 341
280
$$H(s, t) = \Pr[T_{go} \leq s] \times \Pr[T_{stop} \leq t] = F_{go}(s) \times F_{stop}(t),$$
 342
281 343
282 for all s, t ($s, t \geq 0$). 344
283 In addition to estimating the mean via Equation [5], an estimate of the variance of stop signal processing time is obtainable 345
284 in the IND model from 346
285
$$\text{Var}[T_{stop}] = \text{Var}[T_d] - \text{Var}[T_{go}],$$
 347
286 348
287 due to assuming stochastic independence. Finally, it can be shown that the unobservable stop signal distribution function 349
288 $F_{stop}(t)$ is expressible as a function of the observables $p_r(t_d)$, $f_{go}(t)$ and $F_{sr}(t | t_d)$ (see Methods section). 350
289 In order to resolve the paradox described above, a race model with negative dependency between go and stop signal 351
290 processing times is proposed next. We define a bivariate distribution function for T_{go} and T_{stop} exhibiting *perfect negative* 352
291 *dependence* (see, e.g., 22) as follows : 353
292 **Perfect negative stochastic dependence:** 354
293
$$H^-(s, t) = \max\{F_{go}(s) + F_{stop}(t) - 1, 0\}.$$
 355
294 356
295 for all s, t ($s, t \geq 0$). The marginal distributions of $H^-(s, t)$ are the same as before, that is, $F_{go}(s)$ and $F_{stop}(t)$. Note that this 357
296 perfect negative stochastic dependence (PND) model is parameter-free just like the IND race model, that is, we do not assume 358
297 some specific parametric distribution. It can be shown (see Methods) that then 359
298
$$F_{stop}(T_{stop}) = 1 - F_{go}(T_{go})$$
 360
299 361
300 holds “almost surely”, that is, with probability 1. Thus, for any F_{go} percentile we immediately obtain the corresponding F_{stop} 362
301 percentile as complementary probability and vice versa, expressing perfect negative dependence between T_{go} and T_{stop} . The 363
302 relation in Equation [8] is also interpretable as “ T_{stop} is (almost surely) a decreasing function of T_{go} ”. 364
303 The PND model arguably constitutes the most direct implementation of the notion of “mutual inhibition” observed in 365
304 neural data: any increase of inhibitory activity (speed-up of T_{stop}) elicits a corresponding decrease in “go” activity (slow-down 366
305 of T_{go}) and vice versa. 367
306 **Predictions from the PND race model.** We list some predictions from the PND race model (for further details, see Methods 368
307 section): 369
308 370
309 **Probability of stopping.** Under perfect negative dependence between T_{go} and T_{stop} , the probability of stopping $p_r(t_d)$ is an 371
310 increasing function of t_d , as it should. 372

Signal-response RT distribution. Signal-response RT distribution

$$F_{sr}(t|t_d) = \Pr[T_{go} \leq t | T_{go} < T_{stop} + t_d], \quad [4]$$

approaches the go signal distribution $F_{go}(t)$ with increasing stop signal delay t_d . Moreover, for varying values of t_d , $F_{sr}(t|t_d)$ exhibits the “fan effect” typically found in empirical data and also predicted by the IND race model (6, 11).

To illustrate, Figure 1 presents the go signal distribution and signal-response RT distributions with different delay values t_d for both the IND (dashed curves) and the PND race model (solid lines) assuming exponential distributions for T_{go} and T_{stop} . While the exponential distribution lacks empirical support, it was chosen here just to illustrate quantitative differences between the model versions.

An important feature of the PND model is suggested by Figure 1: each signal-response distribution $F_{sr}(t|t_d)$ crosses 1 at a certain finite point t . It is shown below that this point is predictable from the observables $p_r(t_d)$ and $F_{go}(t)$. While this “crossing” property is not shared by the IND model (because the “coupling” of T_{go} and T_{stop} as expressed in Equation (8) is absent), the strength of this test will of course depend on the precision of estimates of $F_{sr}(t|t_d)$.

Estimating moments of the T_{stop} distribution. Given that the marginal distributions F_{go} and F_{stop} are the same under both models, any estimates for $E[T_{stop}]$ based on the mean method (Eq. [5]) are necessarily identical for the IND and PND model.

However, for the variance we obtain

$$\text{Var}[T_{stop}] = \text{Var}[T_d] - \text{Var}[T_{go}] + 2 \text{Cov}[T_{go}, T_{stop}]. \quad [9]$$

As the covariance in the above equation is unknown, an estimate for the stop signal variance is no longer available under the PND model. Nevertheless, given that $\text{Cov}[T_{go}, T_{stop}]$ is the most extreme negative covariance under any bivariate distribution for T_{go} and T_{stop} (23), the PND model stop signal variance can never be larger than the one under independence.

Discussion

Testing IND vs. PND race models. Because stop signal processing times are not observable, empirical testing of non-parametric stop signal race models is severely limited in general. In particular, both the IND and PND version of the race model can only be tested in conjunction with context invariance (CI). Presuming CI is valid, the following predictions can be tested: (1) mean signal-response RT should be faster than mean go signal RT; (2) mean signal-response RT should increase with stop signal delay t_d ; and (3) both these tests are implied by the “fan” structure of the distribution functions permitting an additional qualitative test of IND and PND race models. Because data not consistent with this “fan” structure would be evidence against both the IND and PND model, obviously, none of these tests would allow us to distinguish between IND and PND.

By construction, the PND model has the same distribution for stop signal processing as the IND model; therefore, estimates for mean SSRT, i.e., the expected value of stop signal processing, will be identical for either model as long as estimates are based on the mean method. In a sense, this is good news because it implies that adopting perfect negative dependency between go and stop signal processing does not invalidate previous SSRT estimates of all empirical studies employing the mean method. Note that this holds as well for estimating SSRT via the *integration method*, another way to estimate SSRT. However, because the integration method presumes stop signal processing time to be constant, the distinction between IND and PND model becomes meaningless.

This leaves us with the shape of the family of signal-response distributions, that is, $F_{sr}(t|t_d)$ for varying values of t_d , as the only potential means of distinguishing between the two models. Whether or not this “crossing” test that -as outlined above- consists of checking the way in which the signal-response distributions approach their upper bound of 1, proves to be as visible in real data compared to the exponential toy example (Figure 1) will depend on the specific, but unobservable, stop signal distribution.

Conclusion. Several authors have stressed that the level of description provided by the race model is quite different from that of neural models (17, 20): the independent race model is mute about the possible underlying mechanisms of response inhibition, and its primary purpose is to provide a measure of stop signal processing time, which the model allows without having to assume a specific distribution or estimate parameters. Nevertheless, it has been claimed that the independent race model “captures the essence of computation and ...that it formulates the constraints that any model of response inhibition must follow” (cf. 20, p.3). The race model with perfect negative dependence suggested here operates at the same level of generality and yields the same measure of (mean) stop signal processing time but, importantly, resolves the paradox the independent model is facing due to the neurophysiological findings. We conclude that the race model with perfect negative stochastic dependence is a natural way to unify the observation of interacting circuits of mutually inhibitory gaze-holding and gaze-shifting neurons with data on the behavioral level.

Beyond the simple stop signal task, several other forms of acts of control have been studied that pose an even greater challenge for efforts to identify the underlying cognitive processes, such as switching tasks or shifting of attention (5). One promising area is *selective stopping* where response inhibition is required for some stimuli (e.g. red light) but not others (green light) (e.g., 24). Representing cognitive processes for such more complex tasks may call for types of graded, instead of perfect, dependency. Appropriate race models can easily be introduced by appealing to more general forms of stochastic dependency via copulas (22).

Methods

This section presents some computational details for both the IND and PND model, in general and for the special case of exponential distributions (for producing Figure 1).

We assume that distribution functions F_{go} and F_{stop} possess densities, f_{go} and f_{stop} , and are increasing, where “increasing” always refers to “strictly increasing” here. Note that although it is usually taken for granted in the race model that densities exist, strictly speaking, this is not required neither by the independent nor the dependent race model.

Independent race model. Under stochastic independence of T_{go} and T_{stop} ,

$$\begin{aligned} p_r(t_d) &= \Pr[T_{go} < T_{stop} + t_d] \\ &= \int_0^\infty f_{go}(t)[1 - F_{stop}(t - t_d)] dt. \end{aligned} \quad [10]$$

Moreover,

$$\text{Var}[T_d] = \text{Var}[T_{stop}] + \text{Var}[T_{go}], \quad [11]$$

implying Equation [6].

Writing $f_{sr}(t | t_d)$ for the density function of signal-response time distribution $F_{sr}(t | t_d)$, it has been shown in (25) that

$$f_{sr}(t | t_d) = f_{go}(t) [1 - F_{stop}(t - t_d)] / p_r(t_d). \quad [12]$$

From that, an explicit expression of the distribution of the unobservable stop signal processing time T_{stop} is given by:

$$F_{stop}(t - t_d) = 1 - \frac{f_{sr}(t | t_d) p_r(t_d)}{f_{go}(t)}. \quad [13]$$

Unfortunately, as investigated in (26, 27), gaining reliable estimates for the stop signal distribution using Equation [13] requires unrealistically large numbers of observations.

Perfect negative dependent race model. Defining the bivariate distribution by

$$H^-(s, t) = \max\{F_{go}(s) + F_{stop}(t) - 1, 0\} \quad [7]$$

for all s, t ($s, t \geq 0$), the marginal distributions of $H^-(s, t)$ are again $F_{go}(s)$ and $F_{stop}(t)$. A classic result from the *theory of copulas* (see SI) asserts that

(i) H^- is the lower bound of all bivariate distributions H of (T_{go}, T_{stop}) , i.e., $H^-(s, t) \leq H(s, t)$ for all (s, t) ;

(ii) $\Pr[F_{go}(T_{go}) + F_{stop}(T_{stop}) = 1] = 1$.

From (i) the covariance between the random variables can be shown to be the smallest possible across all bivariate distributions $H(s, t)$ (23). Statement (ii) is equivalent to

$$F_{stop}(T_{stop}) = 1 - F_{go}(T_{go}). \quad [8]$$

holding almost surely.

Probability of stopping. Under perfect negative dependence between T_{go} and T_{stop} ,

$$\begin{aligned} p_r(t_d) &= \Pr[T_{go} - T_{stop} < t_d] \\ &= \Pr[T_{go} - F_{stop}^{-1}[1 - F_{go}(T_{go})] < t_d] \\ &= \Pr[g(T_{go}) < t_d] \\ &= \Pr[T_{go} < g^{-1}(t_d)], \end{aligned} \quad [14]$$

where function g , defined as $g(t) = t - F_{stop}^{-1}[1 - F_{go}(t)]$, is increasing and so its inverse g^{-1} is increasing as well. Thus, $p_r(t_d)$ is increasing in t_d .

Signal-response RT distribution. For stop signal delay t_d ,

$$\begin{aligned} F_{sr}(t | t_d) &= \Pr[T_{go} \leq t | T_{go} < T_{stop} + t_d] \\ &= \Pr[T_{go} \leq t \cap T_{go} < T_{stop} + t_d] / p_r(t_d) \\ &= \Pr[T_{go} \leq t \cap T_{go} < g^{-1}(t_d)] / \Pr[T_{go} < g^{-1}(t_d)] \\ &= F_{go}(\min\{t, g^{-1}(t_d)\}) / F_{go}(g^{-1}(t_d)). \end{aligned}$$

To see that $F_{sr}(t | t_d)$ obeys the “fan effect” consider two different delays, t_d and t_d^* , say, with $t_d < t_d^*$; for a fixed value of t , we first assume $t < g^{-1}(t_d) < g^{-1}(t_d^*)$. Then

$$\begin{aligned} F_{sr}(t | t_d^*) &= F_{go}(t) / F_{go}(g^{-1}(t_d^*)) \\ &< F_{go}(t) / F_{go}(g^{-1}(t_d)) \\ &= F_{sr}(t | t_d). \end{aligned}$$

The cases $g^{-1}(t_d) < t < g^{-1}(t_d^*)$ and $g^{-1}(t_d) < g^{-1}(t_d^*) < t$ can be shown similarly.

Expected stop signal processing time. A formal expression for mean (expected) stop signal processing time is obtained as follows. We have

$$\begin{aligned} E[T_{stop}] &= \int_0^\infty [1 - F_{stop}(s)] ds \\ &= \int_\infty^0 F_{go}(t) \frac{ds}{dt} dt \quad (\text{a.s.}), \end{aligned} \quad [15]$$

using Equation [8] and the fact that $1 - F_{stop}(s)$ and $F_{go}(t)$ have opposite limit values. With

$$s = F_{stop}^{-1}[1 - F_{go}(t)]$$

we have

$$\frac{ds}{dt} = -f_{go}(t)/f_{stop} [F_{stop}^{-1}[1 - F_{go}(t)]] .$$

Inserting into [15] yields

$$E[T_{stop}] = \int_0^\infty \frac{F_{go}(t) f_{go}(t)}{f_{stop} [F_{stop}^{-1}[1 - F_{go}(t)]]} dt. \quad [16]$$

IND model: exponential case. Define independent, exponential distributions for T_{go} and T_{stop} with parameters $\lambda_{go} > 0$ and $\lambda_{stop} > 0$ for context *STOP* by

$$\begin{aligned} H(s, t) &= \Pr[T_{go} \leq s] \times \Pr[T_{stop} \leq t] \\ &= (1 - \exp[-\lambda_{go} s]) \times (1 - \exp[-\lambda_{stop} t]), \end{aligned}$$

for all $s, t \geq 0$. Then

$$\begin{aligned} p_r(t_d) &= \int_0^{t_d} f_{go}(t) dt + \int_{t_d}^\infty f_{go}(t) [1 - F_{stop}(t - t_d)] \\ &= 1 - \frac{\lambda_{stop}}{\lambda_{stop} + \lambda_{go}} \exp[-\lambda_{go} t_d]. \end{aligned} \quad [17]$$

For $t > t_d$, the density of the signal-response distribution is given by,

$$\begin{aligned} f_{sr}(t | t_d) &= f_{go}(t) [1 - F_{stop}(t - t_d)] / p_r(t_d) \\ &= \frac{\lambda_{go} \exp[-\lambda_{go} t] \exp[-\lambda_{stop}(t - t_d)]}{\left(1 - \frac{\lambda_{stop}}{\lambda_{stop} + \lambda_{go}} \exp[-\lambda_{go} t_d]\right)} \\ &= \frac{1}{K} (\lambda_{go} + \lambda_{stop}) \exp[-(\lambda_{go} + \lambda_{stop})(t - t_d)], \end{aligned} \quad [18]$$

with $K = \exp[\lambda_{go} t_d] (1 + \lambda_{stop}/\lambda_{go}) - \lambda_{stop}/\lambda_{go}$. For $t_d = 0$, we have $K = 1$ and the signal-response density is identical to an exponential density for an independent race between T_{stop} and T_{go} , with parameter $\lambda_{go} + \lambda_{stop}$ and $p_r(t_d) = \lambda_{go}/(\lambda_{go} + \lambda_{stop})$. For $t \leq t_d$, the density simplifies to

$$\begin{aligned} f_{sr}(t | t_d) &= f_{go}(t) / p_r(t_d) \\ &= \frac{1}{\lambda_{stop} + \lambda_{go}} \exp[-\lambda_{go}(t - t_d)]. \end{aligned} \quad [19]$$

Computation of the expected value of signal-response RTs yields:

$$\begin{aligned} E[T_{go} | T_{go} < T_{stop} + t_d] &= \int_0^\infty t f_{sr}(t | t_d) dt \\ &= \frac{\lambda_{go} [1 + (\lambda_{go} + \lambda_{stop}) t_d]}{(\lambda_{go} + \lambda_{stop}) \{ \exp[\lambda_{go} t_d] (\lambda_{go} + \lambda_{stop}) - \lambda_{stop} \}}. \end{aligned} \quad [20]$$

In particular, for $t_d = 0$, we obtain $E[T_{go} | T_{go} < T_{stop} + t_d] = 1/(\lambda_{go} + \lambda_{stop})$, consistent with the density we mentioned above for this value of the stop signal delay.

PND model: exponential example. Inserting exponential margins into the bivariate distribution,

$$\begin{aligned} H^-(s, t) &= \max\{F_{go}(s) + F_{stop}(t) - 1, 0\} \\ &= \max\{1 - \exp[-\lambda_{go} s] - \exp[-\lambda_{stop} t], 0\}, \end{aligned} \quad [21]$$

for all $s, t \geq 0$. From [14],

$$\begin{aligned} p_r(t_d) &= \Pr[T_{go} - T_{stop} < t_d] \\ &= \Pr[T_{go} - F_{stop}^{-1}[1 - F_{go}(T_{go})] < t_d] \\ &= \Pr[T_{go} - F_{stop}^{-1}[\exp(-\lambda_{go} T_{go})] < t_d] \\ &= \Pr[T_{go} + 1/\lambda_{stop} \log[1 - \exp(-\lambda_{go} T_{go})] < t_d] \\ &= \Pr[g(T_{go}) < t_d] \\ &= \Pr[T_{go} < g^{-1}(t_d)]. \end{aligned} \quad [22]$$

Note that function

$$g(T_{go}) \equiv T_{go} + 1/\lambda_{stop} \log[1 - \exp(-\lambda_{go} T_{go})]$$

cannot be solved explicitly for T_{go} . Therefore, in order to compute $p_r(t_d)$ and plot signal-response time distributions

$$F_{sr}(t | t_d) = F_{go}(\min\{t, g^{-1}(t_d)\}) / F_{go}(g^{-1}(t_d)),$$

we sampled ($n = 100,000$) from the bivariate distribution function $H^-(s, t)$ using function **simCop** (based on the conditional simulation method, see (22)) from the **copBasic** package of the open source software R (<http://www.r-project.org>).

Table 1 lists the crossing points for the signal-response time distributions obtained from the simulation. They correspond to the vertical lines in Figure 1.

Table 1. Predictions for crossing points $F_{sr}(t | t_d) = 1.0$ in PND model

t_d [ms]	10	50	100	150
$g^{-1}(t_d)$	53.6	79.9	117.7	161.4

Compare with Figure 1

Supporting Information (SI)

Fréchet-Hoeffding bounds and perfect dependency. The dependence between two random variables of a random vector (X, Y) is completely described by its probability distribution. Let $G(x, y)$ be the bivariate distribution function of some random vector (X, Y) :

$$G(x, y) = P(X \leq x, Y \leq y)$$

with marginal distributions F_X and F_Y . Then, it always holds that

$$\begin{aligned} G^-(x, y) &= \max\{F_X(x) + F_Y(y) - 1, 0\} \leq G(x, y) \\ &\leq \min\{F_X(x), F_Y(y)\} = G^+(x, y), \end{aligned}$$

for all x, y for which G is defined (the support of G). Both $G^-(x, y)$ and $G^+(x, y)$ are known as *Fréchet-Hoeffding bounds* and are themselves distribution functions for (X, Y) : G^- corresponds to “perfect” negative dependence between X and Y , while G^+ corresponds to “perfect” positive dependence (see (22) for proofs of this and other claims in this section). “Perfect” dependency can be formulated in various ways. For simplicity, we assume X and Y have continuous distribution functions, the only case we will need below. Here are three equivalent characterizations of *perfect negative dependence* (analogous results exist for perfect positive dependence, but they are not of concern here):

(i) Either $\Pr(X > x, Y > y) = 0$ for all x, y or $\Pr(X \leq x, Y \leq y) = 0$ for all x, y ;

(ii) $\Pr[F_X(X) + F_Y(Y) = 1] = 1$;

(iii) X is almost surely a decreasing function of Y .

Consider the equation inside the “Pr” expression in (ii): Solving for X by taking the inverse function of $F_X(X)$ we get

$$X = F_X^{-1}[1 - F_Y(Y)],$$

from which (iii) follows. Random variables with perfect negative dependence are also known as *antithetic variates* in simulation studies.

1. D. Algom, A. Eidels, R.X.D. Hawkins, B. Jefferson, and J.T. Townsend. Features of response times: identification of cognitive mechanisms through mathematical modeling. In J.R. Busemeyer, Z. Wang, J.T. Townsend, and A. Eidels, editors, *The Oxford Handbook of Mathematical and Computational Psychology*, chapter 4, pages 63–98. Oxford University Press, New York, NY 10016, 2015. ISBN 9780199957996.
2. J.T. Townsend. Serial vs. parallel processing: Sometimes they look like tweedledum and tweedledee but they can (and should be) distinguished. *Psychological Science*, 1:46–54, 1990.
3. J.T. Townsend and M.J. Wenger. A theory of interactive parallel processing: new capacity measures and predictions for a response time inequality series. *Psychological Review*, 111(4):1003–1035, 2004.
4. B.U. Forstmann and E.J. Wagenmakers, editors. *An introduction to model-based cognitive neuroscience*. Springer-Verlag GmbH, New York Heidelberg Dordrecht London, 1st edition, 2015. ISBN 978-1-4939-2235-2. URL http://www.ebook.de/product/22860990/an_introduction_to_model_based_cognitive_neuroscience.html.
5. G.D. Logan. Taking control of cognition: an instance perspective on acts of control. *American Psychologist*, 72(9):875–884, 2017.
6. F. Verbruggen and G.D. Logan. Models of response inhibition in the stop-signal and stop-change paradigms. *Neuroscience & Biobehavioral Reviews*, 33:647–661, 2009. <http://doi.org/10.1016/j.neubiorev.2008.08.014>.
7. G.D. Logan. On the ability to inhibit thought and action: A users’ guide to the stop signal paradigm. In D. Dagenbach and T.H. Carr, editors, *Inhibitory processes in attention, memory, and language*, pages 189–239. Academic Press, San Diego, 1994.
8. F. Verbruggen and G.D. Logan. Response inhibition in the stop-signal paradigm. *Trends in Cognitive Sciences*, 12:418–424, 2008. <http://doi.org/10.1016/j.tics.2008.07.005>.
9. D. Matzke, F. Verbruggen, and G.D. Logan. The stop-signal paradigm. In E.J. Wagenmakers, editor, *Stevens’ Handbook of Experimental Psychology and Cognitive Neuroscience: Methodology*, volume 4. John Wiley & Sons, 4th edition, in press.
10. R. Carpenter and I. Noorani. Preface: Movement suppression: brain mechanisms for stopping and stillness. volume 372 of *Philosophical Transactions of the Royal Society B*. The Royal Society Publishing, 2017. 10.1098/rstb.2016.0542.
11. G.D. Logan and W.B. Cowan. On the ability to inhibit thought and action: a theory of an act of control. *Psychological Review*, 91(3):295–327, 1984.
12. J.D. Schall, T.J. Palmeri, and G.D. Logan. Models of inhibitory control. *Philosophical Transactions of the Royal Society of London B*, 372(20160193), 2017. <http://dx.doi.org/10.1098/rstb.2016.0193>.
13. D.P. Hanes and J.D. Schall. Countermanding saccades in macaque. *Visual Neuroscience*, 12:929–937, 1995.
14. D.P. Hanes, W.F. Patterson, and J.D. Schall. Role of frontal eye field in countermanding saccades: visual, movement and fixation activity. *Journal of Neurophysiology*, 79:817–834, 1998.
15. M. Paré and D.P. Hanes. Controlled movement processing: superior colliculus activity associated with countermanded saccades. *Journal of Neuroscience*, 23:6480–6489, 2003.
16. J.W. Brown, D.P. Hanes, J.D. Schall, and V. Stuphorn. Relation of frontal eye field activity to saccade initiation during a countermanding task. *Experimental Brain Research*, 190:135–151, 2008. DOI 10.1007/s00221-008-1455-0.
17. L. Boucher, T.J. Palmeri, G.D. Logan, and J.D. Schall. Inhibitory control in mind and brain: an interactive race model of countermanding saccades. *Psychological Review*, 114(2):376–397, 2007.
18. J.D. Schall and D.C. Godlove. Current advances and pressing problems in studies of stopping. *Current Opinion in Neurobiology*, 22:1012–1021, 2012.
19. D.P. Hanes and J.D. Schall. Neural control of voluntary movement initiation. *Science*, 274:427–430, 1996.
20. G.D. Logan, M. Yamaguchi, J.D. Schall, and T.J. Palmeri. Inhibitory control in mind and brain 2.0: Blocked-input models of saccadic countermanding on the ability to inhibit thought and action: General and special theories of an act of control. *Psychological Review*, 122:115–147, 2015.
21. C.C. Lo, L. Boucher, M. Paré, J.D. Schall, and X.J. Wang. Proactive inhibitory control and attractor dynamics in countermanding action: a spiking neural circuit model. *Journal of Neuroscience*, 29:9059–9071, 2009.
22. R. B. Nelsen. *An introduction to copulas*, volume 139 of *Lecture Notes in Statistics*. Springer-Verlag, New York, NY, 2nd edition, 2006.
23. M. Denuit and J. Dhaene. Simple characterizations of comonotonicity and countermonotonicity by extremal correlations. *Belgian Actuarial Bulletin*, 3(1):22–27, 2003.
24. P.G. Bissett and G.D. Logan. Selective stopping ? Maybe not. *Journal of Experimental Psychology: General*, 143(1):455–472, 2014.
25. H. Colonius. A note on the stop-signal paradigm, or how to observe the unobservable. *Psychological Review*, 97(2):309–312, 1990.
26. G.P.H. Band, M.W. van der Molen, and G.D. Logan. Horse-race model simulations of the stop-signal procedure. *Acta Psychologica*, 112:105–142, 2003.
27. D. Matzke, C.V. Dolan, G.D. Logan, S.D. Brown, and E.-J. Wagenmakers. Bayesian parametric estimation of stop-signal reaction time distributions. *Journal of Experimental Psychology: General*, 142(4):1047–1073, 2013.

Figure caption for Figure 1. The “fan” effect: all signal-response RT distributions $F_{sr}(t|t_d)$ are strictly ordered by delay (t_d) from left to right, for $t_d = 10, 50, 100, 150$ [ms], converging toward no-stop signal distribution $F_{go}(t)$ (rightmost curve) for $t_d \rightarrow \infty$, for exponential distributions with rate parameters $\lambda_{go} = .01$ and $\lambda_{stop} = .02$; PND model curves were obtained by simulation (see Methods); dashed: IND model; solid: PND model.

Fig. 1