Stopping Timed Actions

Dora Matzke^a, Luke Strickland^b

University of Amsterdam^a, Curtin University^b

Chandra Sripada^c, Alexander Weigard^c

University of Michigan^c

Rohan Puri^d, Jason L. He^e

University of Tasmania^d, King's College London^e

Rebecca J. Hirst^f, Andrew Heathcote^g

Trinity College Dublin^f, University of Newcastle^g

Correspondence concerning this article should be addressed to:

Dora Matzke

Department of Psychology

PO Box 15906, 1001 NK Amsterdam, The Netherlands

E-mail: d.matzke@uva.nl

Abstract

The ability to inhibit ongoing responses that suddenly become inappropriate is essential for safe and effective interaction with an ever-changing and unpredictable world. Response inhibition is quantified by the unobservable stop-signal reaction time (SSRT), the completion time of an inhibitory process triggered by a signal to stop responding. SSRTs can be inferred based on a model in which inhibitory and response processes race with each other to control behavior. Inhibition is usually studied in the context of choice responses, but there has been increasing interest in what is often a key component of skilled behavior, stopping a response that is timed to coincide with an anticipated event. We show that SSRT measurement via the standard race model fails for anticipated responses because the stop signal changes the perception of the passage of time, due to the long-known “filled-interval illusion”. We propose a computational model of anticipated response inhibition that takes account of this distortion of time perception and show that this model produces valid estimates of not only SSRT, but also another key process that determines inhibitory ability, lapses in attention. Our new model and accompanying Bayesian estimation procedures provide a solid basis for the burgeoning study of timed-action control.

Keywords: Anticipated Responses, Context Independence, Independent Race Model, Stop-Signal Paradigm, Stop-Signal Reaction Time, Response Inhibition

Introduction

It is often important to stop an ongoing response timed to coincide with an anticipated event. Driving provides commonplace examples, such as preparing to accelerate in order to pass through an anticipated gap in traffic, but then having to abort when another car takes the opening. The need for timed responding is particularly pervasive when interacting with machinery, because of mechanical delays between taking a controlling action and the desired effect occurring. Hence, the study of timed responding has broad applied interest and the ability to stop anticipatory responding is critical for safe and efficient interactions with the modern world. Problems caused by a failure to stop timed responses can occur not only because there is insufficient time to act on the signal indicating the need to abort an action, but also because of failures related to attention. In the driving example such “trigger failures” could occur when one does not notice a looming car, or because even though it is perceived the goal to avoid an accident by stopping is not activated, perhaps due to fatigue or distraction.

Response inhibition—broadly construed as the ability to stop any type of ongoing action that suddenly becomes inappropriate—has long been recognized as a key component of executive control (Aron, Robbins, & Poldrack, 2014; Miyake et al., 2000; Ridderinkhof, van den Wildenberg, Segalowitz, & Carter, 2004) that is considered to be central not only to everyday interactions with the world, but also to various neurological and psychiatric disorders (e.g., Badcock, Michie, Johnson, & Combrinck, 2002; Schachar, Mota, Logan, Tannock, & Klim, 2000). We first briefly review the dominant method used to investigate response inhibition, the choice-based stop-signal paradigm, the associated independent race model model that enables it to provide a quantitative index of inhibitory ability (Logan, 1994), and recent developments that have extended this model to deal with trigger failures (Matzke, Love, & Heathcote, 2017). We then consider emerging attempts to understand the inhibitory control of anticipated responding. In a broad array of anticipated response inhibition data sets we show that the standard independent race model, even when accounting for trigger failures, does not provide valid measures of the ability to stop timed

actions. To remedy this problem, we propose a new computational cognitive model that provides an accurate account of anticipated response inhibition performance and provides valid estimates of both inhibitory ability and trigger failures.

Traditional Approaches to Measuring Response Inhibition

Response inhibition is most often investigated using the stop-signal paradigm (Logan & Cowan, 1984; for a recent review, see Matzke, Verbruggen, & Logan, 2018). Traditional stop-signal studies rely on a two-choice response time (RT) task (e.g., pressing left button for a left arrow, or right button for a right arrow). On a minority of trials, this primary “go” task is interrupted by a stop signal, at a variable delay (“stop-signal delay” or SSD), indicating that participants should cancel their response. Response inhibition succeeds when the stop signal comes soon enough after the onset of the go stimulus but fails when it comes later – too close to the moment of response execution. The stop-signal paradigm has been used in a large number of studies, both with healthy and clinical populations, to examine the cognitive, developmental, and neural aspects of response inhibition (e.g., Aron & Poldrack, 2006; Badcock et al., 2002; Bissett & Logan, 2011; Fillmore, Rush, & Hays, 2002; Forstmann et al., 2012; Hughes, Fulham, Johnston, & Michie, 2012; Matzke, Hughes, Badcock, Michie, & Heathcote, 2017; Schachar et al., 2000; Schachar & Logan, 1990; Verbruggen, Stevens, & Chambers, 2014; Williams, Ponesse, Schachar, Logan, & Tannock, 1999).

The stop-signal paradigm is popular because, when combined with the independent race model, it enables the estimation of the unobservable latency of the stop process, known as the stop-signal reaction time (SSRT, Colonius, 1990; de Jong, Coles, Logan, & Gratton, 1990; Logan & Cowan, 1984; for an overview, see Matzke et al., 2018). In this model, go and stop processes race (with the stop process handicapped by SSD) to control behavior: if the go process finishes first, a response is made; if the stop process finishes first, no response is made. For the purpose of mathematical simplicity, the race model relies on two types of independence assumptions: stochastic and context independence.

Stochastic independence states that, on a given trial, the finishing times of the go process and the stop process are uncorrelated: for all $t_{go} \geq 0$ and $t_{stop} \geq 0$,

$$P(T_{go} < t_{go} \cap T_{stop} < t_{stop}) = P(T_{go} < t_{go}) \times P(T_{stop} < t_{stop}), \quad (1)$$

where T_{go} and T_{stop} are random variables representing the finishing times of the go and the stop processes with continuous probability density functions $f_{go}(t)$ and $f_{stop}(t)$, respectively. Context independence states that the distribution of the finishing times of the go process is the same on no-signal (“go”) trials and stop-signal trials: for all t_{go} and t_{SSD} ,

$$P(T_{go} < t_{go}) = P(T_{go} < t_{go} | t_{SSD}). \quad (2)$$

Both assumptions are essential for the validity of non-parametric SSRT estimation methods, such as the mean and the integration method (e.g., Logan, 1994; Matzke et al., 2018; Verbruggen et al., 2019).¹ Non-parametric methods provide computationally straightforward ways to compute summary measures of the speed of the stop process. The most popular version of the mean method entails subtracting the mean of the inhibition function (i.e., probability of responding on stop-signal trials as a function of SSD) from the mean of the go RTs, where the mean of the inhibition function is approximated using the mean of the SSDs. The integration method entails first selecting the n^{th} go RT, where n equals the number of go RTs multiplied by the overall probability of responding on stop-signal trials, and then subtracting mean SSD from the n^{th} go RT. These methods do not assume a specific form for the finishing-time distributions of the go and the stop process, but nonetheless do make distributional assumptions. The mean method assumes that the distribution of go RTs is symmetrical. Distribution symmetry is important because skewed go RT distributions typically result in asymmetrical inhibition functions, in which case mean SSD provides a poor estimate of the mean of the inhibition function, resulting

¹An exception is when the finishing times of the go and the stop process are perfectly negatively correlated, in which case the mean method provides valid SSRT estimates if context independence holds (Colonius & Diederich, 2018).

in overestimation of SSRTs (Verbruggen, Chambers, & Logan, 2013).² Since most choice RT (CRT) distributions are skewed to the right and the degree of skew can vary between groups and experimental conditions, the mean method can mask differences or result in spurious SSRT effects. The integration method is less sensitive to the skewness of the go RT distribution, but it assumes that the finishing time of the stop process is constant, which is clearly implausible as SSRTs —just like go RTs— are certainly variable.

Parametric approaches, which assume a specific form for the finishing-time distributions, require computationally expensive model fitting in order to estimate the parameters of the go RT and the SSRT distribution. They, however, provide more information about SSRT in the form of its full distribution (Matzke, Dolan, Logan, Brown, & Wagenmakers, 2013; Matzke, Love, et al., 2013). They are also more flexible: parametric models can explicitly account for choice errors (Matzke, Curley, Gong, & Heathcote, 2019) and can deal with trigger failures occurring when the stop signal does not activate the stop process (Matzke, Love, & Heathcote, 2017), which are known to invalidate non-parametric estimation methods (Logan, 1994). Although standard parametric approaches typically assume stochastic and context independence, parametric modeling does not necessarily require these assumptions. For example, context independence is not required if the finishing-time distribution of the go process on go trials does not share parameters with the finishing-time distribution of the go process on stop-signal trials.

Anticipated Response Inhibition Tasks

Leunissen, Zandbelt, Potocanac, Swinnen, and Coxon (2017) argued that *Anticipated Response Inhibition* (ARI) tasks are not only of interest in their own right, but that they also have better measurement properties than the traditional approach. The ARI task is a variant of the stop-signal paradigm where the go task requires a response that coincides with a predictably timed event. They provided three reasons why ARI tasks are better for

²Note that this problem can be remedied by estimating the mean of the inhibition function directly as opposed to approximating it using mean SSD (Logan & Cowan, 1984). However, this approach requires a relatively large number of stop-signal trials and hence it is not used frequently.

SSRT estimation than the traditional choice-based approach and tasks based on a simple detection of the go stimulus.³

First, anticipated response time (ART) distributions are more symmetric than CRT distributions, so are more suitable for the mean method. Second, ART varies less between trials than CRT, at least for the anticipation of short (e.g., < 1 second) intervals. Between-trial variability increases measurement noise and so decreases the reliability of measures derived from RTs (Rouder & Haaf, 2019). Lastly, in choice-based tasks, participants often slow their go responses in an effort to beat the stop signal and maximize the probability of stopping. This is not a problem for SSRT estimation if the slowing is constant, but if it is progressive and varies across trials, it can again add measurement noise and so reduce reliability, and can also cause systematic biases (Verbruggen et al., 2013). In the standard choice-based task, slowing is a fruitful strategy because it improves choice accuracy. In contrast, slowing reduces the accuracy of anticipated responses, which by definition require a constant ART, and so this problem is naturally less pronounced in ARI tasks than in the choice-based approach.

Even before these advantages were highlighted by Leunissen et al. (2017), ARI tasks were becoming increasingly popular because of their clear relevance to real-world action control. Although this interest is still relatively recent, in a search of the literature, we identified 57 papers using variants of the ARI procedure based on a variety of go stimuli and stop signals. In almost two-third of the reviewed ARI studies, the go stimulus was a rising bar, where participants were required to press or release a button when a rising bar intersected a fixed target line. On stop-signal trials, the bar stopped rising before reaching the fixed target line, indicating that participants should stop their response (e.g., Coxon, Stinear, & Byblow, 2007; He et al., 2019; Yamanaka, Kadota, & Nozaki, 2013). In one of the studies, stopping was indicated by a change in the color of the rising bar (Verbruggen, Adams, & Chambers, 2012). In the remaining one-third of the studies, go stimuli included a

³Detection go tasks are rarely used because they are often associated with occasional anticipatory responses made by guessing when the go stimulus will appear. Without knowledge of which trials are based on detection and which on anticipation, and a model that can take this knowledge into account, SSRT estimation is problematic (e.g., Matzke et al., 2019).

moving pointer on a circular clock-type display (Carlsen, Chua, Timothy Inglis, Sanderson, & Franks, 2008; Coxon, Stinear, & Byblow, 2006), a runway of illuminating LED lights (Lobjois, Benguigui, & Bertsch, 2005), and the movement of “gamified” stimuli (e.g., fruit dropping from a tree, Smittenaar et al., 2015; a race car moving along a track, Gilbert et al., 2019). Stop signals included a halt in the motion of the go stimulus (He et al., 2018; MacDonald, Coxon, Stinear, & Byblow, 2014), an auditory stimulus (Ilmane & LaRue, 2011; LaRue, 2007), or a visual stimulus (e.g., a STOP sign, Marinovic, Reid, Plooy, Riek, & Tresilian, 2010, 2011). Some studies have extended this basic paradigm to investigate selective stopping (e.g., stopping a bimanual go response with one hand while continuing to respond with the other hand, Cirillo, Cowie, MacDonald, & Byblow, 2018; Zandbelt & Vink, 2010) or used manipulations such as indicating, via a cue, the probability of stop-signal occurrence (Kleerekooper et al., 2016; Vink, Kaldewaij, Zandbelt, Pas, & du Plessis, 2015).

Two-third of the 57 papers estimated SSRT in at least one of their experiments. The use of ARI-based SSRT estimates has been gaining popularity in recent years: prior to 2010 only about 20% of the studies (3 out of 14), whereas after 2010 nearly 80% of the studies (35 out of 43) reported SSRT estimates based on the independent race model. ARI-based SSRT estimates have been used to study inhibitory control across the lifespan, including children (van Hulst et al., 2018), adolescents (Vink et al., 2014) and young and older adults (Hermans et al., 2018), in healthy as well as clinical populations (e.g., schizophrenia, Zandbelt, van Buuren, Kahn, & Vink, 2011; post-traumatic stress disorder, van Rooij, Geuze, Kennis, Rademaker, & Vink, 2015). Inference using SSRT estimates has often been critical to conclusions made in ARI studies, let these be between-group/condition comparisons or correlations between SSRT and neural measures. Most studies used the integration method or the mean method in combination with interpolation to estimate SSRTs.

To our knowledge, parametric approaches have never been used to estimate SSRTs in ARI tasks, although the methods of Matzke, Dolan, et al. (2013) and Matzke, Love, and Heathcote (2017) are applicable, as the ex-Gaussian parametric distribution form they

assume is able to accommodate symmetric as well as right skewed distributions. The genesis of this paper was an attempt to use Matzke and colleagues’ parametric approach with ARI data. However, parametric assessment of the assumptions of the race model made it immediately obvious that context independence was often clearly violated. In the next section, we discuss a potential cause: a long-known phenomena in time estimation called the *filled-interval illusion*.

Context Independence and the Filled-Interval Illusion

It is a robust and long-established finding that filled intervals are perceived as longer than empty intervals (Ornstein, 1969). For example, participants produce a longer interval when matching a previously presented interval filled with events compared to an empty interval (Thomas & Brown, 1974), and also provide longer verbal estimates of time (Wearden, Norton, Martin, & Montford-Bebb, 2007). The filled-interval illusion is observed from sub-second to much longer time scales, and is induced by a range of filling materials, from a series of short tones (Thomas & Brown, 1974) and continuous tones (Wearden et al., 2007), to attending to complex stimuli such as line drawings (Ornstein, 1969). The illusion can be substantial (e.g., more than a 50% increase in Wearden et al., 2007). Timing ability (i.e., duration perception and production) is usually attributed to a pacemaker processes that accumulates timing pulses, or an analogous continuous quantity (Simen, Vlasov, & Papadakis, 2016), at a constant rate. Wearden et al. (2007) found that the filled-interval illusion likely results from an increase in the rate of that process.

If the production of anticipated responses relies on timing ability, the filled-interval illusion —due to its pervasive nature— could affect performance in ARI tasks, producing violations of context independence. How such an effect might manifest is unclear, because the timing literature typically relies on different type of tasks than ARI paradigms. ARTs may be slowed in the presence of a stop signal, analogous to the longer times produced to match filled intervals in the design of Thomas and Brown (1974). Conversely, if an anticipated response is produced after a fixed number of pacemaker pulses, and the occurrence

of the stop signal quickens the pacemaker, ARTs on stop trials may speed up relative to responses on go trials. In both cases, there should be an accompanying increase or decrease in ART variability, as timing has a scale-invariance property, such that the coefficient of variation (i.e., the standard deviation divided by the mean) is constant (Simen et al., 2016).

Non-Parametric Tests of Context Independence

The standard non-parametric test of context independence compares mean go RT to the mean of RTs on stop-signal trials that escaped inhibition – also known as signal-respond RTs. If context independence holds, signal-respond RTs should be faster than go RTs. This test is based on the independent race model’s prediction that the signal-respond RT distribution and the go RT distribution share a common lower bound, but their cumulative distribution functions (CDFs) diverge at higher quantiles. The common lower bound implies that mean signal-respond RT is shorter than mean go RT (Colonius, Özyurt, & Arndt, 2001; Osman, Kornblum, & Meyer, 1986). In other words, the independent race model assumes that signal-respond RTs are samples from the go RT distribution, where the stop signal causes longer samples to be increasingly censored because they lose the race with the stop process, so no response occurs and no RT is recorded. Hence, the average of the remaining (generally fast) signal-respond RTs will be lower than the mean of the go RTs. Although the standard test based on distribution means may pick up violations caused by the filled-interval illusion lengthening ARTs, it will be ineffective if they are shortened.

Here we developed a stronger test of context independence based on the *fan effect* (Colonius & Diederich, 2018): the strict ordering and fanning of the CDF of go RTs and signal-respond RTs as a function of SSD. The upper left panel of Figure 1a shows the CDF of go RTs and the CDF of signal-response RTs on five SSDs in artificial choice-based stop-signal data. The CDFs are non-decreasing functions rising from zero at a time (t) corresponding to the minimum possible RT to a value of one as t increases, with the height of the curve corresponding to the probability of observing $RT < t$. When context independence holds, the go CDF and the signal-respond CDFs at the various SSDs start at the same value

of t , but the signal-respond CDFs rise more quickly, causing a fanning pattern before the functions come back together near one. The shorter the SSD, the steeper the rise of the CDF of the signal-respond RTs so there is a strict ordering of the CDFs: signal-respond CDFs at lower SSDs always occur on the left of those at higher SSDs. When SSDs are determined by the staircase tracking algorithm (e.g., Logan, Schachar, & Tannock, 1997), as is most commonly the case, early and late SSDs will have only a few RTs, so the corresponding CDFs are poorly estimated. However, as shown in the upper middle panel of Figure 1a, the strict ordering of the go and the signal-respond CDF and the characteristic fan shape also hold when signal-respond RTs are collapsed across SSD, enabling a more reliable test of context independence in data sets with a limited number of stop-signal trials.

In order to investigate context independence, we use the CDF of signal-respond RTs collapsed over SSD to create *delta functions* (de Jong, Liang, & Lauber, 1994; Speckman, Rouder, Morey, & Pratte, 2008). Delta functions show the horizontal difference between CDFs (i.e., the difference in RTs corresponding to the same probability, which are called the distribution quantiles) as a function of probability.⁴ As we subtract the signal-respond quantiles from the go quantiles, context independence predicts that the delta function is never negative and usually clearly positive: the signal-respond CDF “dominates” the go CDF. As shown in the upper right panel of Figure 1a, delta functions provide a powerful and sensitive visualization tool that makes the fan effect evident as an increasing trend, reflecting the increasing separation between the go and the signal-respond CDFs. This pattern of delta functions —positive and increasing— does not assume any particular parametric form: distributions with a substantial right skew (as in choice-based tasks, Figure 1a) and more symmetric distributions (as in ARI tasks, Figure 1b) both produce increasing delta functions, but the increase is typically smaller in the symmetric case.

One might attempt to construct a non-parametric test of context independence based on patterns that *cannot* be produced by a race model that assumes context independence,

⁴The points of delta functions are often plotted against the average of the two quantiles on the x-axis, but that difference is immaterial here. Note also that, because the CDFs typically come back together only at very high probabilities, the corresponding decrease in the delta function to zero can rarely be observed in practice.

such as decreasing delta functions or failures of dominance. Although this is clearly better than a simple test on mean go and signal-respond RT, it would still produce a relatively weak and partial assessment, detecting only severe violations. In the next section, we describe a parametric test that is not only sensitive, but is also able to account for trigger failures that distort the fan effect even when context independence holds.

Figure 1. The fan effect and delta functions. Upper panels: Cumulative distribution function (CDF) of go RTs and signal-respond RTs for a go process finishing-time distribution with relatively large variance and appreciable right skew (i.e., choice-based tasks). Lower panels: CDFs of go RTs and signal-respond RTs for a go process finishing-time distribution with relatively small variance and negligible right skew (i.e., ARI tasks). Left column: CDF of signal-respond RTs as a function of stop-signal delay (SSD). Middle column: CDF of signal-respond RTs collapsed across SSD. Right column: Delta functions based on signal-respond RTs collapsed across SSDs. The figures are based on synthetic data (measured in seconds) generated from ex-Gaussian go and stop finishing-time distributions.

Trigger Failures and the Fan Effect. Parametric modeling has found trigger failures to be pervasive in choice-based stop-signal tasks. Trigger failures refer to failures to initiate the inhibition process in response to the stop signal. They occur at higher rates

in clinical populations (e.g., schizophrenia, Matzke, Hughes, et al., 2017; ADHD, Weigard, Heathcote, Matzke, & Huang-Pollock, 2019) and when healthy participants are fatigued (Skippen et al., 2019). Even under standard conditions, trigger failures have been estimated to occur on 5 – 10% of trials on average, both when the go choice is easy (Matzke, Hughes, et al., 2017) and when it is difficult (Matzke et al., 2019). Trigger failures are thought to reflect mechanisms related to attention rather than inhibition, such as failures in encoding the stop signal, and are known to inflate non-parametric SSRT estimates (see also Band, van der Molen, & Logan, 2003).

As shown in Figure 2 — which illustrates the effects of the highest average level of trigger failures that have been reported in healthy participants, 18% — trigger failures can cause the strict ordering of CDFs to break down.⁵ At this level, trigger failures substantially decrease the slope of the delta function in the choice-based task and almost entirely flatten it in the ARI task (compare Figures 1 and 2, which apart from trigger failures were generated using the exact same go and stop finishing-time distributions). The Supplementary Materials (<https://osf.io/tw46u/>) show that with 50% trigger failures, the delta function is still increasing for the choice, but is decreasing for the ARI setting. Note, however, that trigger failures can never cause a failure of dominance (i.e., negative delta function values), and because of the high rate required, we believe that in practice they would rarely, if ever, cause decreasing delta functions.

A Parametric Test of Context Independence

The variation in delta functions shown in Figures 1 and 2 makes it clear that in only extreme cases can violations of context independence be unambiguously detected by a simple characterization of the delta functions (e.g., as strongly decreasing or negative). We instead propose a more sensitive procedure that compares the observed delta function to the delta function produced by fitting a flexible parametric race model that assumes context independence. In particular, our proposed method is based on the BEESTS model (Matzke,

⁵Failures to trigger the go process (i.e., go failures), which are also often observed in choice-based stop-signal tasks, do not substantially modulate the shape of delta functions.

(a) Choice-based task with 18% trigger failures.

(b) ARI task with 18% trigger failures.

Figure 2. The fan effect and delta functions in the presence of trigger failures. Upper panels: Cumulative distribution function (CDF) of go RTs and signal-respond RTs for a go process finishing-time distribution with relatively large variance and appreciable right skew (i.e., choice-based tasks). Lower panels: CDFs of go RTs and signal-respond RTs for a go process finishing-time distribution with relatively small variance and negligible right skew (i.e., ARI tasks). Left column: CDF of signal-respond RTs as a function of stop-signal delay (SSD). Middle column: CDF of signal-respond RTs collapsed across SSD. Right column: Delta functions based on signal-respond RTs collapsed across SSDs. The figures are based on synthetic data (measured in seconds) generated from ex-Gaussian go and stop finishing-time distributions with 18% trigger failures.

Dolan, et al., 2013; Matzke, Love, & Heathcote, 2017; Matzke et al., 2019), a Bayesian parametric race model that treats the finishing times of both the go and the stop process as random variables and therefore enables the estimation of the entire distribution of go RTs and SSRTs. BEESTS is flexible because it characterizes the go and stop finishing times with ex-Gaussian distributions, which are known to provide a very accurate description of observed RT distributions (Heathcote, Popiel, & Mewhort, 1991; Hohle, 1965; Matzke & Wagenmakers, 2009). The ex-Gaussian distribution is the sum of a normal random variable (which is symmetric) and an exponential random variable (which is strongly right skewed). When the exponential component has higher variance, the ex-Gaussian is sufficiently skewed to accurately describe CRT distributions, whereas when it has low variance, the more symmetric case characteristic of ART distributions is accommodated. Importantly, BEESTS allows the effects of trigger failures as well as go failures to be accounted for.

BEESTS infers the posterior distribution of the model parameters by updating the prior distributions with incoming data (Heathcote et al., 2019). The posterior distributions reflect knowledge about the parameters after the data have been observed and capture uncertainty due to the full range of relevant factors including those related to the experimental design, such as the number of trials per participant. In the Supplementary Materials, we present a brief overview of the basic concepts of Bayesian parameter estimation. For comprehensive introductions to Bayesian methods in general and Bayesian cognitive modeling in particular, the reader is referred to Edwards, Lindman, and Savage (1963), Farrell and Lewandowsky (2018), Gelman and Hill (2007), Kruschke (2010), Lee and Wagenmakers (2013), and Wagenmakers et al. (2018).

Once the posterior distributions are obtained, we use posterior predictive simulations to evaluate whether the model, which assumes context independence, can adequately account for the observed data (Gelman, Meng, & Stern, 1996). In particular, for each participant, we use 1,000 samples from the joint posterior distribution to simulate 1,000 data sets from the experimental design, using the observed SSDs and the observed number of go and stop-signal trials. The simulated data sets are then used to calculate *posterior pre-*

dictive delta functions. As we use the entire joint posterior to generate the predicted delta functions, we not only account for sampling error, but also take into account the uncertainty of the parameter estimates. To provide a compact summary of the results, we average the observed and posterior predictive delta functions over participants; as both observed and predicted data are treated in the same way, there is no risk of the averaging distorting the results. Systematic discrepancies between the observed and the predicted delta functions are taken as evidence that the data cannot be explained by a race model that assumes context independence. Note that this procedure is not only sensitive to qualitative violations, such as failures of dominance, but can also detect more subtle quantitative violations, such as discrepancies between the slopes of the observed and predicted delta functions.

Context Independence Violations in Empirical Data

In this section, we report the results of applying our parametric method of testing context independence to a range of empirical data sets, both from our own laboratories and from Leunissen et al. (2017).

Data Sets

We first present two stop-signal conditions from Matzke et al. (2019) in an easy and difficult choice-based go task with a visual stop signal, demonstrating that the race model assuming context independence and allowing for trigger-and go failures provides an accurate account of observed delta functions.⁶ We then examined a range of different ARI data sets with different types of stop signals. The use of different stop signals was motivated by the possibility that some stop signals may affect the processing of the go stimulus even when they do not replace it. This could lead to context independence violations additional to, or instead of, the filled-interval illusion. As shown in Table 1, we analyzed four ARI data

⁶Note that we did not examine the choice stop-signal data reported by Leunissen et al. (2017) because their task design replaced the go stimulus with the stop signal. Clearly, this could lead to a violation of context independence at short SSDs. In particular, SSDs less than 0.1 second must lead to an apparent decrease in stimulus intensity due to Bloch’s law, and even at longer SSDs there may still be some weakening of the go stimulus that could influence the finishing times of the go process, leading to further violations of context independence.

sets from Leunissen et al. (2017), and data from three new ARI tasks (Experiments 1-3; see Supplementary Materials for the full details of these experiments). All of the ARI tasks required participants to anticipate when a bar rising at a constant rate intersected a marker, defining a short time interval of less than 1 second. In the four Leunissen et al. (2017) data sets and our Experiment 1, the stop signal involved a halt in the rise of the bar. If the timing of the anticipated response is, at least in part, based on the motion of the bar, it is possible that this type of stop signal causes a violation of context independence, instead of or in addition to the filled-interval illusion.

To further investigate this possibility, in Experiment 2 we instead changed the color of the bar on stop trials, but the bar kept moving at the same rate. Experiment 3 had the same type of stop signal used in the choice-based task of Matzke et al. (2019), a red square that appeared around the rising bar at a sufficient distance that violations of context independence due to perceptual interactions were unlikely. Experiments 2 and 3 also aimed to improve data quality (Smith & Little, 2018) relative to the other data sets by collecting more than double the number of trials per participant (see Table 1). Experiment 2 used a “gamified” format in an attempt to maintain motivation throughout the longer task (Hamari et al., 2016). Experiment 3 compared the same gamified format in one condition with a standard format more typical of most psychology experiments; the standard format used the same software platform but with the game-related details stripped away.

Table 1

Overview of the data sets, estimated probability of trigger failures (P_{TF}), and sample size (N = number of participants; n = number of stop-signal trials).

Data set	Task type	Condition	Stop signal	Mean P_{TF}	N/n
Matzke et al. (2019)	choice	easy	square	6	53/84
Matzke et al. (2019)	choice	difficult	square	5	53/84
Leunissen et al. (2017) Exp. 1	ARI	release	bar stops	2	17/90
Leunissen et al. (2017) Exp. 1	ARI	press	bar stops	1	17/90
Leunissen et al. (2017) Exp. 2	ARI	release, young adults	bar stops	3	13/90
Leunissen et al. (2017) Exp. 2	ARI	release, older adults	bar stops	2	12/90
Experiment 1	ARI	release	bar stops	2	7/100
Experiment 2	ARI	gamified	color change	12	64/230
Experiment 3	ARI	gamified	square	18	9/240
Experiment 3	ARI	standard	square	29	9/240

Fitting the Independent Race Model

Data Pre-Processing. For the choice-based stop-signal data, we followed the data-cleaning procedures described in Matzke et al. (2019); the data are available at <https://osf.io/m26u/>. For the ARI data sets, we removed trials with ARTs shorter than 0.4s; the data are available at <https://osf.io/tw46u/>. All data sets were fit on the second scale.

Models and Prior Distributions. For the choice-based stop-signal data, we used the BEESTS model developed by Matzke et al. (2019), which assumes context independence and uses one racer to represent each choice on the go task and one to represent the stop process. All racers have ex-Gaussian finishing-time distributions, and the model allows for both trigger failures and go failures. The ARI models differed only in that there was one rather than two ex-Gaussian racers for the go process (see Matzke, Dolan, et al., 2013; Matzke, Love, et al., 2013; Matzke, Love, & Heathcote, 2017). The implementation of the models is available in the Dynamic Models of Choice software (Heathcote et al., 2019) at <https://osf.io/pbwx8/>.

We fit the data of each participant separately using Differential Evolution Markov Chain Monte Carlo sampling as detailed in Heathcote et al. (2019). We did not use hierarchical methods because we wanted to obtain the best possible fit for each participant without influence from the group. As shown in the Supplemental Materials, we used weakly informative (truncated) normal prior distributions for both models. The trigger failure and go failure parameters were estimated on the probit scale and were transformed back to the probability scale for inference.

As described earlier, we used posterior predictive simulations to assess the descriptive accuracy of the models. Importantly, we did not only do this for the delta functions, but also for other aspects of the data, such as inhibition functions, median signal-respnd RTs as a function of SSD, and the distribution of go RTs and signal-respnd RTs.

Results

Table 1 provides the average trigger failure estimates for each data set. These ranged from small to the substantial levels used in our simulations reported earlier (see Figure 2). As trigger failures give the model a limited ability to mimic context-independence violations, it is possible that estimates associated with the ARI data sets are inflated. We discuss this further in the Discussion. In what follows, we present the observed and posterior predictive average delta functions for each data set. The full set of posterior predictive results are presented in the Supplementary Materials.

Figure 3 shows the observed and predicted delta functions averaged across participants for the Matzke et al. (2019) data. The observed delta functions follow the expected steeply increasing pattern in both conditions. The model provides an excellent description of these functions, suggesting that the race model with context independence describes these data very well.

As shown in Figure 4, the first three Leunissen et al. (2017) data sets (Experiment 1 press and release conditions and Experiment 2 young adults) resulted in increasing observed average delta functions. However, the independent race model does not provide an adequate

(a) *Easy condition.*(b) *Difficult condition.*

Figure 3. Average delta functions in the Matzke et al. (2019) choice-based stop-signal data sets. The red line shows the observed delta functions averaged across participants. The black lines show 1,000 posterior predictive delta functions averaged across participants obtained from fitting the data with the standard BEESTS model assuming context independence; the darker the area, the denser the predicted delta functions.

approximation of the observed average delta functions; even after accounting for trigger failures, the slope of the delta functions predicted by the model is systematically steeper than the slope of the observed delta functions. That is, segments of the observed delta functions are on the edge or outside the black/gray area representing the model's predictions.

The bottom right panel of Figure 4 shows that the fourth Leunissen et al. (2017) data set (older adults in Experiment 2) displays severe violations, resulting in a decreasing average delta function, at least for the faster responses. Correspondingly, the failure of the model is striking. Note that even in this case, trigger failure rate was estimated as quite low, likely because, although decreasing, the observed delta function remains strongly positive. With hindsight, the failure of context independence is evident in the CDFs plotted in Leunissen et al.'s Figure 5, but the much clearer violation in Figure 4 demonstrates why we prefer to use delta functions and inference based on a parametric model.

Figure 5 shows results for the new experiments. Experiment 1 displays a slightly

(a) *Experiment 1: Release condition.* (b) *Experiment 1: Press condition.* (c) *Experiment 2: Young adults.* (d) *Experiment 2: Older adults.*

Figure 4. Average delta functions for Experiments 1 and 2 in Leunissen et al. (2017). The red line shows the observed delta functions averaged across participants. The black lines show 1,000 posterior predictive delta functions averaged across participants obtained from fitting the data with the standard BEESTS model assuming context independence; the darker the area, the denser the predicted delta functions.

(a) *Experiment 1.* (b) *Experiment 2.* (c) *Experiment 3: Standard.* (d) *Experiment 3: Gamified.*

Figure 5. Average delta functions in our Experiments 1 - 3. The red line shows the observed delta functions averaged across participants. The black lines show 1,000 posterior predictive delta functions averaged across participants obtained from fitting the data with the standard BEESTS model assuming context independence; the darker the area, the denser the predicted delta functions.

exaggerated version of the pattern for the first three Leunissen et al. (2017) data sets. The remaining data sets display more extreme violations. For Experiment 2, the observed average delta function is clearly decreasing. The large sample size and excellent data quality of this experiment resulted in very tight model predictions that remain slightly increasing, despite the relatively high — 12% — trigger failure rate, and so they clearly diverge from the observed data. The violations in Experiment 3 are so extreme, with the slower portions of the observed delta functions being negative, that a model is not needed to diagnose the failure of context independence. Recall that the independent race model cannot produce negative delta functions because this pattern implies a failure of the strict ordering of the signal-respond RT and go RT distributions. The estimated trigger failure rates in the two Experiment 3 data sets were higher than in any other, but no matter how high they become, they cannot explain the clear failure of dominance displayed in the observed data.

As shown in the Supplementary Materials, the results of the posterior predictive simulations for other aspects of the data indicate that delta plots often provide more fine-grained information about model misfit than inhibition functions, median signal-respond RTs, and the distribution of go RTs and signal-respond RTs. For instance, the clear violation of the independent race model exhibited by the older participants in Leunissen et al.’s (2017) Experiment 2 is not apparent in the go RTs and signal-respond RTs when their distribution is considered in isolation and not relative to each other, as is done by delta functions. Similarly, model predictions for the inhibition function and the median signal-respond RTs do not seem to indicate misfit, perhaps with the exception of RTs on late SSDs that seem to increase slightly more steeply than the observed data.

Modeling Context Independence Violations

Here we propose a new parametric race model that accounts for the pervasive context independence violations in ARI tasks described in the previous section. The proposed model is based on Matzke, Dolan, et al.’s (2013) ex-Gaussian BEESTS approach. The ex-Gaussian μ and σ parameters quantify the mean and standard deviation of the Gaussian component,

and τ reflects the mean (and equivalently standard deviation) of the exponential component.

The probability density function of the ex-Gaussian distribution is:

$$f(t; \mu, \sigma, \tau) = \frac{1}{\tau} \exp\left(\frac{\mu - t}{\tau} + \frac{\sigma^2}{2\tau^2}\right) \Phi\left(\frac{t - \mu}{\sigma} - \frac{\sigma}{\tau}\right), \text{ for } \sigma > 0, \tau > 0, \quad (3)$$

where Φ is the standard normal distribution function, defined as

$$\Phi\left(\frac{t - \mu}{\sigma} - \frac{\sigma}{\tau}\right) = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\frac{t - \mu}{\sigma} - \frac{\sigma}{\tau}} \exp\left(-\frac{y^2}{2}\right) dy. \quad (4)$$

The distribution function of the ex-Gaussian distribution is:

$$F(t; \mu, \sigma, \tau) = \Phi\left(\frac{t - \mu}{\sigma}\right) - \exp\left(\frac{\sigma^2}{2\tau^2} - \frac{t - \mu}{\tau}\right) \Phi\left(\frac{t - \mu}{\sigma} - \frac{\sigma}{\tau}\right), \quad (5)$$

and its mean and variance equal

$$E = \mu + \tau \quad (6)$$

and

$$\text{Var} = \sigma^2 + \tau^2, \quad (7)$$

respectively. BEESTS relies on a separate set of ex-Gaussian parameters to describe the finishing-time distribution of the go process and the finishing-time distribution of the stop process: μ_{go} , σ_{go} , and τ_{go} for the go, and μ_{stop} , σ_{stop} , and τ_{stop} for the stop process.

As previously discussed, the standard BEESTS model does not require context independence, although it does benefit from the parsimony it offers and hence assumes that the finishing-time distribution of the go process is the same on go trials and stop-signal trials. That is, in ARI models with one go and one stop racer, the same ex-Gaussian *go* parameters that account for the observed distribution of ARTs on go trials are assumed to account for the underlying distribution of RTs on stop-signal trials.

Here we develop a new “Context-Independence Violation” BEESTS model (BEESTS-CV) that relaxes this assumption, enabling us to directly investigate the nature of the violations in cases where they were most evident: Leunissen et al.’s (2017) Experiment 2 and our Experiment 3. Note that BEESTS-CV continues to assume stochastic independence. Rather than estimating entirely separate ex-Gaussian parameters for the finishing-time distribution of the go process on stop-signal trials, we reasoned that a slightly more parsimonious formulation might suffice whereby this distribution is linearly related to the observed finishing-time distribution of the go process on go trials (i.e., go RT distribution). In particular, we estimate a parameter that uniformly shifts the go RT distribution, M , and one that scales it, S , so that the finishing-time distribution of the go process on stop-signal trials is described by the following parameters:

$$\mu_{go_{stop}} = \mu_{go} + M, \quad (8)$$

$$\sigma_{go_{stop}} = \sigma_{go} \times S, \quad (9)$$

$$\tau_{go_{stop}} = \tau_{go} \times S. \quad (10)$$

In terms of the filled-interval illusion, positive values of M suggest that the anticipated interval is lengthened by the occurrence of the stop signal, and negative values of M that it is shortened. Similarly, values of $S > 1$ suggest that the anticipated interval is made longer and more variable by the occurrence of the stop signal, and $S < 1$ that it is made shorter and less variable (see Equation 6 and Equation 7). The implementation of BEESTS-CV is available at: <https://osf.io/tw46u/>. The prior distributions are available in the Supplementary Materials.

Results

Figure 6 shows that BEESTS-CV provides an excellent account of the observed delta functions. The new model does not only accommodate the relatively subtle misfit to Leunissen et al.’s (2017) younger participants in Experiment 2, but also the negative

slopes found for their older participants and the failures of dominance in our Experiment 3. As for the earlier results, the full set of posterior predictive results are presented in the Supplementary Materials. BEESTS-CV provides a clearly more accurate account than the standard BEESTS model also to other aspects of the data, such as the inhibition functions, median signal-respond RTs, and the distribution of go RTs and signal-respond RTs. Critically, after accounting for context independence violations, BEESTS-CV provides the same excellent account of ART distributions on go trials as it did for CRT distributions. This indicates that any misfit of the standard BEESTS with respect to delta functions is unlikely to be due to the parametric assumption of ex-Gaussian finishing-time distributions.

Given BEESTS-CV accurately describes the data sets, we now examine its parameters in order to understand the nature of the context independence violations, and to determine how these distort estimates of inhibitory ability. Table 2 provides parameter estimates for both the standard BEESTS model that assumes context independence and the new BEESTS-CV model that relaxes this assumption. It is evident that in all cases the violation mainly takes the form of an increase in S by between 22% – 41%, which manifests in an increase both in the mean and the variability of the finishing-time distribution of the go process on stop trials relative to go trials. The uniform shift M is positive in all cases, but it is relatively small, with the 95% credible intervals indicating that it is negligible in Leunissen et al.’s (2017) experiment but not in our Experiment 3. Note that an increase in S has a larger proportional effect on the variability of the finishing-time distribution than on its mean because the mean is also affected by $\mu_{go} + M$ (see Equation 6) and because $\mu_{go} \gg \tau_{go}$. This underscores again that assessing violations of context independence by comparing mean go RT and mean signal-respond RT is insufficient. Although comparing the variability in go RTs and signal-respond RTs is more informative in this case, we believe that the only generally effective approach is parametric modeling.

Table 2 shows that the mis-specified standard BEESTS model consistently produced an upward bias in trigger failures by a factor of 1.3 – 1.5. The absolute margin was substantial, almost 10%, in the *Standard* condition of our Experiment 3 where trigger failures

(a) *Experiment 2: Young adults.* (b) *Experiment 2: Older adults.* (c) *Experiment 3: Standard.* (d) *Experiment 3: Gamified.*

Figure 6. Average delta functions for Experiment 2 in Leunissen et al. (2017) and our Experiment 3. The red line shows the observed delta functions averaged across participants. The black lines show 1,000 posterior predictive delta functions averaged across participants obtained from fitting the data with the Context-Independence Violation BEESTS model; the darker the area, the denser the predicted delta functions.

were most common. Table 2 also shows the mean SSRT estimates produced by the models, which are obtained as $\mu_{stop} + \tau_{stop}$. In all cases, the mis-specified standard BEESTS model created a downward bias in SSRT estimates relative to BEESTS-CV. The bias was mostly relatively small, ranging between 4 – 11 ms, but for the *Standard* condition of our Experiment 3 it was a more substantial 25 ms.

Table 2

Means of the posterior distribution of the standard BEESTS and BEESTS-CV estimates averaged across participants and the corresponding 95% credible intervals (in parentheses) for Leunissen et al.'s (2017) Experiment 2 and our Experiment 3.

Parameter	CI Young	CV Young	CI Older	CV Older	CI Game	CV Game	CI Standard	CV Standard
P_{TF}	3.0 (0.9-5.6)	2.3 (0.7-4.5)	2.1 (0.4-5.1)	1.6 (0.3-5.2)	17.7 (23.8-28.6)	13.4 (9.6-17.9)	29.1 (25.3-33.7)	19.5 (15.4-23.5)
P_{GF}	1.3 (1-1.8)	1.3 (1-1.7)	11.1 (10.1-12.2)	11.0 (9.9-12.1)	2.6 (2.3-3)	2.6 (2.3-3)	1.2 (1-1.5)	1.2 (1-1.5)
μ_{stop}	146 (135-154)	153 (144-161)	163 (147-176)	176 (162-188)	184 (156-211)	228 (210-244)	205 (190-223)	225 (208-239)
σ_{stop}	20 (12-32)	16 (9-25)	31 (19-50)	24 (15-39)	56 (26-119)	23 (11-44)	25 (9-66)	22 (10-45)
τ_{stop}	23 (13-37)	20 (12-30)	35 (22-54)	31 (20-45)	76 (37-123)	44 (23-71)	31 (11-56)	36 (19-61)
SSRT	169 (163-175)	173 (168-179)	198 (189-208)	208 (198-216)	261 (238-286)	272 (262-285)	236 (223-251)	261 (252-273)
μ_{go}	824 (820-828)	823 (820-827)	856 (851-861)	856 (851-861)	691 (689-692)	690 (689-692)	691 (689-692)	692 (690-693)
σ_{go}	40 (39-42)	38 (37-40)	67 (65-69)	63 (62-66)	29 (28-29)	27 (27-28)	27 (27-28)	27 (26-28)
τ_{go}	15 (12-19)	16 (13-20)	11 (7-16)	11 (8-16)	23 (21-24)	22 (20-23)	21 (19-22)	18 (17-20)
M	- -	2 (-2-7)	- -	9 (-1-18)	- -	8 (5-11)	- -	12 (9-15)
S	- -	1.22 (1.13-1.32)	- -	1.31 (1.21-1.43)	- -	1.41 (1.32-1.5)	- -	1.41 (1.33-1.48)

Note. CI = Standard BEESTS model with context independence. CV = Context-Independence Violation BEESTS

model. P_{TF} = probability of trigger failures. P_{GF} = probability of go failures. SSRT = $\mu_{stop} + \tau_{stop}$. NB:

Estimation was carried out on the second scale, but the estimates are reported on the millisecond scale.

We also calculated non-parametric SSRT estimates for these data sets using the mean method and the integration method that accounts for omissions on go trials (Verbruggen et al., 2019). With the exception of the *Standard* condition of Experiment 3, where the mean

method produced an SSRT estimate 5 ms slower than BEESTS-CV, the mean method typically produced under-estimates by 9 – 25 ms, with average estimates of 163 ms and 183 ms for Leunissen et al. (2017)’s younger and older adults, respectively, and 263 ms and 266 ms for our *Gamified* and *Standard* conditions. The percentage of variance shared between the BEESTS-CV and the mean method estimates ranged from 7% and 31% for older and young participants, respectively, to 39% and 73% for the *Standard* and *Gamified* conditions of Experiment 3. The integration method also produced under-estimates, but performed better on average, with values of 170 ms and 202 ms for the younger⁷ and older participants, respectively, and 260 ms for both the *Gamified* and *Standard* conditions of Experiment 3. Shared variance ranged from 62% and 38% for older and young participants, respectively, to 32% and 79% for the *Standard* and *Gamified* conditions.

Finally, we performed a model selection analysis, whereby we compared the standard BEESTS model assuming context independence with BEESTS-CV assuming either the full context-independence violation parameterization ($M > 0$ and $S > 0$) or only allowing scaling (i.e., fixing M to zero). We based our comparison on the Deviance Information Criterion (DIC, Spiegelhalter, Best, Carlin, & van der Linde, 2002), which weighs the improvement in goodness-of-fit obtained by estimating extra parameters against the associated increase in the model’s flexibility, and hence the possibility of over-fitting. Lower DIC values indicate that the model strikes a better balance between goodness-of-fit and parsimony, where a difference of greater than ten is often considered to provide strong evidence for the chosen model. In the Leunissen et al. (2017) data sets, the support for the full BEESTS-CV over the standard BEESTS was strong (41 units for both young and older adults), whereas the support for $M > 0$ was more equivocal, with the full BEESTS-CV winning by 17 units from the restricted BEESTS-CV ($M = 0$) for the young group, but losing by 5 units for the older group. For our Experiment 3, the support for the full BEESTS-CV was stronger, winning by 158 and 236 units from the standard BEESTS model for the *Gamified* and

⁷Note that the SSRT estimate reported by Leunissen et al. (2017) for younger participants was 7 ms lower than our estimate of 180 ms. This difference results from slight differences in their implementation of the integration method: Leunissen et al. did not account for omissions on go trials, discarded error RTs, and removed RTs more than three standard deviation from mean go RT.

Standard conditions, respectively, and by 45 and 61 units from the model assuming $M = 0$. Hence, the model selection results largely confirmed our conclusions based on the parameter estimates.

Discussion

The anticipated response inhibition (ARI) task measures inhibitory ability in arguably one of the most common and important contexts for action control: the stopping of an ongoing response that is timed to coincide with a predicted event. It is not only an increasingly popular method for understanding cognitive control (e.g., Coxon et al., 2006; He et al., 2018; MacDonald et al., 2014; Verbruggen et al., 2012), it has also been claimed to provide a better measure of response inhibition than the traditional choice-based stop-signal paradigm. In particular, Leunissen et al. (2017) have argued that ARI tasks provide a more reliable measurement of the speed of the inhibitory process (i.e., stop-signal reaction time or SSRT), and to avoid methodological problems associated with the strategic slowing of choice responses. In all past applications, ARI-based SSRT estimates have been obtained through non-parametric methods based on the independent race model (Logan & Cowan, 1984). In order for those estimates to be valid, the occurrence of the stop signal must not change the speed of the anticipated response or “go” process.

Here we first developed a new and sensitive test to check this “context independence” assumption based on delta functions. Context independence predicts that delta functions, which plot the difference between RTs on go and signal-respond (i.e., failed inhibition) trials as a function of overall speed, must be positive and—even in the presence of realistic levels of trigger failures—increasing. Delta functions estimated for each SSD separately enable the assessment of the ordering of delta functions implied by the fan effect. However, this approach would require data quality that is almost never obtained in real stop-signal studies. We have therefore opted for an alternative approach, where we collapsed signal-respond RTs across SSDs in a single distribution before computing the delta function. This approach necessarily loses information about the ordering of signal-respond RT distributions as a

function of SSD, but it allows us to detect subtle violations of context independence relative to the go RTs even in scarce data. Our method provides more fine-grained information about model misfit than the standard non-parametric test of independence using mean RTs or the assessment of descriptive accuracy using inhibition functions and signal-respond RTs as a function of SSD, or the distribution of go RTs and signal-respond RTs considered in isolation. Importantly, our parametric model-based approach explicitly acknowledges that goodness-of-fit is not an all-or-none phenomenon; violations of context independence are often a matter of degree which necessitates the use of methods that are not only sensitive to qualitative violations, such as failures of dominance, but can also detect more subtle quantitative violations, such as discrepancies between the slopes of the observed and predicted delta functions.

We applied our parametric test of context independence to two choice-based and eight ARI data sets. The ARI data came from paradigms that all used the most common method of eliciting an anticipation of the time at which a steadily increasing bar intersected a marker, but differed in the stop signal. We raised the concern that context independence might be violated in ARI tasks due to a pervasive perceptual effect, the filled-interval illusion (Ornstein, 1969), whereby the occurrence of a stop signal changes the perception of time relative to trials without a stop signal. Consistent with this concern, the ARI data sets displayed varying degrees of context independence violations. Clear violations occurred with all of the different stop signals, including the most common method in the literature, a halt in the rising bar, but also with a change in the color of the bar, and even with the same type of stop signal we used in the choice-based task (i.e., square). The variants of the ARI task we examined are not exhaustive. Hence, there may be other variants in which context independence holds, and the methods that we developed can be used to investigate this possibility.⁸ However, given that the cases we examined cover much of the past literature,

⁸For instance, consider an experimental setup where both go and stop-signal trials involve a change in the color of the rising bar: on go trials, the bar turns blue and participants should continue responding, whereas on stop-signal trials, the bar turns red indicating that participants should cancel their response. This setup would essentially turn the ARI task into a selective-stopping task requiring a choice about the color change. Although this task may avoid context independence violations as a result of the filled-interval illusion, it would invite other problems associated with the use of the independent race model in the context

and the likelihood that the filled-interval illusion has a pervasive influence, our results raise serious concerns about using ARI tasks in combination with the standard independent race model to measure SSRT.

Similar to many other studies (e.g., Camalier et al., 2007; de Jong et al., 1990; Hanes & Carpenter, 1999; Hanes & Schall, 1995; Osman et al., 1986), our analysis did not find violations of independence in the choice-based stop-signal data. In particular, we analyzed two conditions reported in Matzke et al. (2019) and found that the race model assuming context independence provided an excellent account of the observed delta functions, inhibition functions, median signal-respond RTs as a function of SSD, and the distribution of go and signal-respond RTs. Note, however, that a few other studies have reported violations of independence (e.g., Colonius et al., 2001; Logan, 1981; Logan, Cowan, & Davis, 1984), especially at short SSDs that typically feature only a small number of signal-respond RTs. Although our findings suggest that context independence violations are less common, and if present, less severe in choice-based tasks, the origins of the mixed results remain uncertain. We believe that our parametric approach coupled with high quality stop-signal data has the potential to bring new insights about the presence and extent of potential independence violations in choice-based paradigms.

In order to address the pervasive context independence violations in ARI tasks, and to better understand why they occur, we developed a new parametric model. The model is based on Matzke, Dolan, et al.'s (2013) ex-Gaussian BEESTS approach and can accommodate violations of context independence. We applied this to data sets where context independence violations were the most extreme: Leunissen et al.'s (2017) Experiment 2 comparing the performance of younger and older participants, and our Experiment 3, comparing performance between standard and gamified ARI tasks. BEESTS-CV provided an accurate account of the data of both the younger participants, where the violation was small, and the data of older participants, where the violation was large enough to result in a decreasing delta function. The violations were almost entirely accounted for by a pro-

of selective stopping (Bissett & Logan, 2011).

portional increase in the scale of anticipated response times (ARTs) on stop-signal trials relative to go trials, by 22% for the younger participants and 31% for the older participants. An even larger increase in scale of 41% provided an accurate account of the severe violations in our Experiment 3, which produced not only decreasing delta functions but also negative delta function values for the slowest responses.

We compared parameter estimates from the standard BEESTS and BEESTS-CV to investigate the distortions in parametric estimates of inhibitory ability that erroneously assume context independence. The standard BEESTS consistently underestimated SSRT, with an average downward bias of 13 ms. There was, however, a much greater effect on estimates of the percentage of trigger failures, cases in which the stop process fails to enter the race when a stop signal occurs, with an inflation by up to 50% in the worst case. This finding is consistent with the tendency of trigger failures to flatten the slope of the delta function, providing a limited ability to mimic context independence violations. However, even when this tendency was removed in the fits of BEESTS-CV, the estimated trigger failure rate remained substantial in our Experiment 3. Failing to account for trigger failures has been shown to cause SSRTs to be overestimated, with serious consequences for understanding the nature of inhibitory processing and group differences (Matzke, Hughes, et al., 2017; Skippen et al., 2019). It appears that the same concerns apply to SSRT estimates derived from ARI tasks, at least in some cases, making non-parametric methods of estimating SSRT unattractive as they cannot account for trigger failures.

Despite this shortcoming, we also examined non-parametric SSRT estimates in order to assess how their use might have affected past conclusions in the ARI literature. The mean method typically resulted in underestimation of SSRT, with an average of 15 ms, with the exception of the *Standard* condition of Experiment 3, where it produced an SSRT estimate 5 ms slower than BEESTS-CV. The integration method fared much better, generally only slightly underestimating SSRT, with an average difference of 6 ms between integration and BEESTS-CV estimates. Note that these results must be interpreted in the context of the finding that failures to account for trigger failures cause overestimation of SSRT, and

suggest a trade-off occurring between the effects of trigger failures and context independence violations. Whether such fortuitous trade-offs led to reasonably accurate SSRT estimates in the past literature remains a matter of conjecture that only retrospective application of the BEESTS-CV model could resolve. In any case, we believe that it would be unwise to rely on non-parametric SSRT estimates in future research, especially given BEESTS-CV enables both SSRTs and trigger failures to be estimated without making the potentially invalid and misleading assumption that context independence holds in ARI tasks.

If BEESTS-CV is to be adopted as the default approach to ARI analysis it would be desirable in future work to investigate how it can be further refined. A promising avenue for simplification that is consistent with the scale-invariance property of timing models (Simen et al., 2016) is to model context independence violations only in terms of a scaling parameter. Although scaling effects were clearly the dominant factor in our results, model selection did favor a small uniform shift in the ART distributions as well, at least for some of the data sets. Whether this holds more broadly, or whether a pure scaling model is both sufficient and advantageous because of its simplicity, remains to be seen. It would also be desirable to more closely examine its measurement properties using parameter-recovery and model-recovery studies (Heathcote et al., 2019). Such studies could also determine whether, and if so in what designs, ARI offers a more reliable framework for estimating SSRTs than the traditional choice-based paradigm.

Conclusion

In summary, the parametric BEESTS-CV analysis developed here provides a method that researchers can confidently use to investigate action control in the context of timed responding. Not only does it address what appears to be the pervasive violation of context independence associated with the filled-interval illusion, it also enables inhibition-based effects to be disentangled from attention-based effects associated with failures to act on the stop signal. Further, it need not be taken on trust that our approach is equally applicable to the many potential variants of the ARI task that are likely to be of interest in future

research because of the availability of powerful and sensitive new ways of testing for context independence and BEESTS-CV's ability to account for them. This not only enables a wider investigation of inhibitory control beyond the narrow focus of traditional choice-based approaches, but also enables researchers to access the advantages of ARI tasks in cases where strategic slowing effects are of concern.

Acknowledgements

DM is supported by a Veni grant (451-15-010) from the Netherlands Organization of Scientific Research (NWO). CSS was supported by U01DA041106 from NIDA. AW was supported by T32AA007477 from NIAAA. RP was supported through an Australian Government Research Training Program stipend. AH is supported by an ARC Discovery project (DP120102907).

Data Availability Statement

The ARI data from Experiments 1-3 are available at <https://osf.io/tw46u/>.

References

- 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.
- Aron, A. R., Robbins, T. W., & Poldrack, R. A. (2014). Inhibition and the right inferior frontal cortex: One decade on. *Trends in Cognitive Sciences*, *18*, 177–185.
- Badcock, J. C., Michie, P., Johnson, L., & Combrinck, J. (2002). Acts of control in schizophrenia: Dissociating the components of inhibition. *Psychological Medicine*, *32*, 287–297.
- Band, G. P. H., van der Molen, M. W., & Logan, G. D. (2003). Horse-race model simulations of the stop-signal procedure. *Acta Psychologica*, *112*, 105–142.
- 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.
- Camalier, C., Gotler, A., Murthy, A., Thompson, K., Logan, G., Palmeri, T., & Schall, J. (2007). Dynamics of saccade target selection: Race model analysis of double step and search step saccade production in human and macaque. *Vision Research*, *47*, 2187–2211.
- Carlsen, A. N., Chua, R., Timothy Inglis, J., Sanderson, D. J., & Franks, I. M. (2008). Motor preparation in an anticipation-timing task. *Experimental Brain Research*, *190*, 453–461.
- Cirillo, J., Cowie, M. J., MacDonald, H. J., & Byblow, W. D. (2018). Response inhibition activates distinct motor cortical inhibitory processes. *Journal of Neurophysiology*, *119*, 877–886.
- Colonius, H. (1990). A note on the stop-signal paradigm, or how to observe the unobservable. *Psychological Review*, *97*, 309–312.
- Colonius, H., & Diederich, A. (2018). Paradox resolved: Stop signal race model with negative dependence. *Psychological Review*, *125*, 1051–1058.
- Colonius, H., Özyurt, J., & Arndt, P. A. (2001). Countermanding saccades with auditory stop signals: Testing the race model. *Vision Research*, *41*, 1951–1968.
- Coxon, J. P., Stinear, C. M., & Byblow, W. D. (2006). Intracortical inhibition during volitional inhibition of prepared action. *Journal of Neurophysiology*, *95*, 3371–3383.
- Coxon, J. P., Stinear, C. M., & Byblow, W. D. (2007). Selective inhibition of movement. *Journal of Neurophysiology*, *97*, 2480–2489.
- de Jong, R., Coles, M. G., Logan, G. D., & Gratton, G. (1990). In search of the point of no return: The control of response processes. *Journal of Experimental Psychology: Human Perception*

- and Performance*, 16, 164–182.
- de Jong, R., Liang, C.-C., & Lauber, E. (1994). Conditional and unconditional automaticity: a dual-process model of effects of spatial stimulus-response correspondence. *Journal of Experimental Psychology: Human Perception and Performance*, 20, 731–750.
- Edwards, W., Lindman, H., & Savage, L. J. (1963). Bayesian statistical inference for psychological research. *Psychological Review*, 70, 193–242.
- Farrell, S., & Lewandowsky, S. (2018). *Computational modeling of cognition and behavior*. Cambridge University Press.
- Fillmore, M. T., Rush, C. R., & Hays, L. (2002). Acute effects of oral cocaine on inhibitory control of behavior in humans. *Drug and Alcohol Dependence*, 67, 157–167.
- Forstmann, B. U., Keuken, M. C., Jahfari, S., Bazin, P.-L., Neumann, J., Schäfer, A., ... Turner, R. (2012). Cortico-subthalamic white matter tract strength predicts interindividual efficacy in stopping a motor response. *Neuroimage*, 60, 370–375.
- Gelman, A., & Hill, J. (2007). *Data analysis using regression and multilevel/hierarchical models*. Cambridge, UK: Cambridge University Press.
- Gelman, A., Meng, X., & Stern, H. (1996). Posterior predictive assessment of model fitness via realized discrepancies. *Statistica Sinica*, 6, 733–807.
- Gilbert, D. L., Huddleston, D. A., Wu, S. W., Pedapati, E. V., Horn, P. S., Hirabayashi, K., ... Mostofsky, S. H. (2019). Motor cortex inhibition and modulation in children with ADHD. *Neurology*, 93, e599–e610.
- Hamari, J., Shernoff, D. J., Rowe, E., Coller, B., Asbell-Clarke, J., & Edwards, T. (2016). Challenging games help students learn: An empirical study on engagement, flow and immersion in game-based learning. *Computers in Human Behavior*, 54, 170–179.
- Hanes, D. P., & Carpenter, R. H. S. (1999). Countermanding saccades in humans. *Vision Research*, 39, 2777–2791.
- Hanes, D. P., & Schall, J. D. (1995). Countermanding saccades in macaque. *Visual Neuroscience*, 12, 929–937.
- He, J. L., Fuelscher, I., Coxon, J., Barhoun, P., Parmar, D., Enticott, P. G., & Hyde, C. (2018). Impaired motor inhibition in developmental coordination disorder. *Brain and Cognition*, 127, 23–33.
- He, J. L., Fuelscher, I., Coxon, J., Chowdhury, N., Teo, W. P., Barhoun, P., ... Hyde, C. (2019). Individual differences in intracortical inhibition predict motor-inhibitory performance. *Exper-*

- imental Brain Research*, 237, 2715–2727.
- Heathcote, A., Lin, Y.-S., Reynolds, A., Strickland, L., Gretton, M., & Matzke, D. (2019). Dynamic models of choice. *Behavior Research Methods*, 5, 961–985.
- Heathcote, A., Popiel, S. J., & Mewhort, D. J. (1991). Analysis of response time distributions: An example using the Stroop task. *Psychological Bulletin*, 109, 340–347.
- Hermans, L., Leunissen, I., Pauwels, L., Cuypers, K., Peeters, R., Puts, N. A., ... Swinnen, S. P. (2018). Brain GABA levels are associated with inhibitory control deficits in older adults. *Journal of Neuroscience*, 38, 7844–7851.
- Hohle, R. H. (1965). Inferred components of reaction times as functions of foreperiod duration. *Journal of Experimental Psychology*, 69, 382–386.
- Hughes, M. E., Fulham, W. R., Johnston, P. J., & Michie, P. T. (2012). Stop-signal response inhibition in schizophrenia: Behavioural, event-related potential and functional neuroimaging data. *Biological Psychology*, 89, 220–231.
- Ilmane, N., & LaRue, J. (2011). Postural and focal inhibition of voluntary movements prepared under various temporal constraints. *Acta Psychologica*, 136, 1–10.
- Kleerekooper, I., van Rooij, S. J., van den Wildenberg, W. P., de Leeuw, M., Kahn, R. S., & Vink, M. (2016). The effect of aging on fronto-striatal reactive and proactive inhibitory control. *NeuroImage*, 132, 51–58.
- Kruschke, J. K. (2010). Bayesian data analysis. *Wiley Interdisciplinary Reviews: Cognitive Science*, 1, 658–676.
- LaRue, J. (2007). A new methodology providing evidence of two distinct processes in the production of hand/foot simultaneous responses. *Human Movement Science*, 26, 757–768.
- Lee, M. D., & Wagenmakers, E.-J. (2013). *Bayesian modeling for cognitive science: A practical course*. Cambridge, UK: Cambridge University Press.
- Leunissen, I., Zandbelt, B. B., Potocanac, Z., Swinnen, S. P., & Coxon, J. P. (2017). Reliable estimation of inhibitory efficiency: To anticipate, choose or simply react? *European Journal of Neuroscience*, 45, 1512–1523.
- Lobjois, R., Benguigui, N., & Bertsch, J. (2005). Aging and tennis playing in a coincidence-timing task with an accelerating object: The role of visuomotor delay. *Research Quarterly for Exercise and Sport*, 76, 398–406.
- Logan, G. D. (1981). Attention, automaticity, and the ability to stop a speeded choice respons. In J. B. Long & A. Baddeley (Eds.), *Attention and performance IX* (pp. 205–222). Hillsdale,

NJ: Erlbaum.

- 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. 189–240). San Diego, CA: Academic Press.
- Logan, G. D., & Cowan, W. B. (1984). On the ability to inhibit thought and action: A theory of an act of control. *Psychological Review*, *91*, 295–327.
- Logan, G. D., Cowan, W. B., & Davis, K. A. (1984). On the ability to inhibit simple and choice reaction time responses: A model and a method. *Journal of Experimental Psychology: Human Perception and Performance*, *10*, 276–291.
- Logan, G. D., Schachar, R. J., & Tannock, R. (1997). Impulsivity and inhibitory control. *Psychological Science*, *8*, 60–64.
- MacDonald, H. J., Coxon, J. P., Stinear, C. M., & Byblow, W. D. (2014). The fall and rise of corticomotor excitability with cancellation and reinitiation of prepared action. *Journal of Neurophysiology*, *112*, 2707–2717.
- Marinovic, W., Reid, C. S., Plooy, A. M., Riek, S., & Tresilian, J. R. (2010). Delayed inhibition of an anticipatory action during motion extrapolation. *Behavioral and Brain Functions*, *6*, 22.
- Marinovic, W., Reid, C. S., Plooy, A. M., Riek, S., & Tresilian, J. R. (2011). Corticospinal excitability during preparation for an anticipatory action is modulated by the availability of visual information. *Journal of Neurophysiology*, *105*, 1122–1129.
- Matzke, D., Curley, S., Gong, C. Q., & Heathcote, A. (2019). Inhibiting responses to difficult choices. *Journal of Experimental Psychology: General*, *148*, 124–142.
- 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.
- Matzke, D., Hughes, M., Badcock, J. C., Michie, P., & Heathcote, A. (2017). Failures of cognitive control or attention? The case of stop-signal deficits in schizophrenia. *Attention, Perception, & Psychophysics*, *79*, 1078–1086.
- Matzke, D., Love, J., & Heathcote, A. (2017). A Bayesian approach for estimating the probability of trigger failures in the stop-signal paradigm. *Behavior Research Methods*, *49*, 267–281.
- Matzke, D., Love, J., Wiecki, T. V., Brown, S. D., Logan, G. D., & Wagenmakers, E.-J. (2013). Release the BEESTS: Bayesian estimation of ex-Gaussian stop-signal reaction time distributions. *Frontiers in Quantitative Psychology and Measurement*, *4*:918, doi: 10.3389/fpsy-

- syg.2013.00918.
- Matzke, D., Verbruggen, F., & Logan, G. (2018). The stop-signal paradigm. In E.-J. Wagenmakers & J. T. Wixted (Eds.), *Stevens' handbook of experimental psychology and cognitive neuroscience, Volume 5: Methodology (4th ed.)* (p. 383-427). John Wiley & Sons, Inc.
- 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.
- Miyake, A., Friedman, N. P., Emerson, M. J., Witzki, A. H., Howerter, A., & Wager, T. D. (2000). The unity and diversity of executive functions and their contributions to complex “frontal lobe” tasks: A latent variable analysis. *Cognitive Psychology*, 41, 49–100.
- Ornstein, E. (1969). *On the experience of time*. London: Penguin Books.
- 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.
- Ridderinkhof, K. R., van den Wildenberg, W. P., Segalowitz, S. J., & Carter, C. S. (2004). Neurocognitive mechanisms of cognitive control: The role of prefrontal cortex in action selection, response inhibition, performance monitoring, and reward-based learning. *Brain and Cognition*, 56, 129–140.
- Rouder, J. N., & Haaf, J. M. (2019). A psychometrics of individual differences in experimental tasks. *Psychonomic Bulletin & Review*, 26, 452–467.
- Schachar, R., & Logan, G. D. (1990). Impulsivity and inhibitory control in normal development and childhood psychopathology. *Developmental Psychology*, 26, 710–720.
- Schachar, R., Mota, V. L., Logan, G. D., Tannock, R., & Klim, P. (2000). Confirmation of an inhibitory control deficit in attention-deficit/hyperactivity disorder. *Journal of Abnormal Child Psychology*, 28, 227–235.
- Simen, P., Vlasov, K., & Papadakis, S. (2016). Scale (in)variance in a unified diffusion model of decision making and timing. *Psychological Review*, 123, 151–181.
- Skippen, P., Matzke, D., Heathcote, A., Fulham, W., Michie, P., & Karayanidis, F. (2019). Reliability of triggering inhibitory process is a better predictor of impulsivity than SSRT. *Acta Psychologica*, 192, 104-117.
- Smith, P. L., & Little, D. R. (2018). Small is beautiful: In defense of the small-N design. *Psychonomic Bulletin & Review*, 25, 2083–2101.

- Smittenaar, P., Rutledge, R. B., Zeidman, P., Adams, R. A., Brown, H., Lewis, G., & Dolan, R. J. (2015). Proactive and reactive response inhibition across the lifespan. *PLOS ONE*, *10*, e0140383.
- Speckman, P. L., Rouder, J. N., Morey, R. D., & Pratte, M. S. (2008). Delta plots and coherent distribution ordering. *The American Statistician*, *62*, 262–266.
- Spiegelhalter, D. J., Best, N. G., Carlin, B. P., & van der Linde, A. (2002). Bayesian measures of model complexity and fit (with discussions). *Journal of the Royal Statistical Society Series B*, *64*, 583–616.
- Thomas, E. C., & Brown, I. (1974). Time perception and the filled-duration illusion. *Perception & Psychophysics*, *16*, 449–458.
- van Hulst, B. M., de Zeeuw, P., Vlaskamp, C., Rijks, Y., Zandbelt, B. B., & Durston, S. (2018). Children with ADHD symptoms show deficits in reactive but not proactive inhibition, irrespective of their formal diagnosis. *Psychological Medicine*, *48*, 2508–2514.
- van Rooij, S. J., Geuze, E., Kennis, M., Rademaker, A. R., & Vink, M. (2015). Neural correlates of inhibition and contextual cue processing related to treatment response in PTSD. *Neuropsychopharmacology*, *40*, 667–675.
- Verbruggen, F., Adams, R., & Chambers, C. D. (2012). Proactive motor control reduces monetary risk taking in gambling. *Psychological Science*, *23*, 805–815.
- Verbruggen, F., Aron, A. R., Band, G. P., Beste, C., Bissett, P. G., Brockett, A. T., ... Boehler, C. N. (2019). A consensus guide to capturing the ability to inhibit actions and impulsive behaviors in the stop-signal task. *eLife*, *8*, e46323.
- Verbruggen, F., Chambers, C. D., & Logan, G. D. (2013). Fictitious inhibitory differences: How skewness and slowing distort the estimation of stopping latencies. *Psychological Science*, *24*, 352–362.
- Verbruggen, F., Stevens, T., & Chambers, C. D. (2014). Proactive and reactive stopping when distracted: An attentional account. *Journal of Experimental Psychology: Human Perception and Performance*, *40*, 1295.
- Vink, M., Kaldewaij, R., Zandbelt, B. B., Pas, P., & du Plessis, S. (2015). The role of stop-signal probability and expectation in proactive inhibition. *European Journal of Neuroscience*, *41*, 1086–1094.
- Vink, M., Zandbelt, B. B., Gladwin, T., Hillegers, M., Hoogendam, J. M., van den Wildenberg, W. P., ... Kahn, R. S. (2014). Frontostriatal activity and connectivity increase during

- proactive inhibition across adolescence and early adulthood. *Human Brain Mapping*, 35, 4415–4427.
- Wagenmakers, E.-J., Marsman, M., Jamil, T., Ly, A., Verhagen, J., Love, J., . . . Morey, R. (2018). Bayesian inference for psychology. part I: Theoretical advantages and practical ramifications. *Psychonomic Bulletin & Review*, 25, 35–57.
- Wearden, J. H., Norton, R., Martin, S., & Montford-Bebb, O. (2007). Internal clock processes and the filled-duration illusion. *Journal of Experimental Psychology: Human Perception and Performance*, 33, 716–729.
- Weigard, A., Heathcote, A., Matzke, D., & Huang-Pollock, C. (2019). Cognitive modeling suggests that attentional failures drive longer stop-signal reaction time estimates in ADHD. *Clinical Psychological Science*, 7, 856–872.
- 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.
- Yamanaka, K., Kadota, H., & Nozaki, D. (2013). Long-latency TMS-evoked potentials during motor execution and inhibition. *Frontiers in Human Neuroscience*, 7, 751.
- Zandbelt, B. B., van Buuren, M., Kahn, R. S., & Vink, M. (2011). Reduced proactive inhibition in schizophrenia is related to corticostriatal dysfunction and poor working memory. *Biological Psychiatry*, 70, 1151–1158.
- Zandbelt, B. B., & Vink, M. (2010). On the role of the striatum in response inhibition. *PLoS ONE*, 5, e13848.