6

Control in Response Inhibition

Frederick Verbruggen (University of Exeter) and Gordon D. Logan (Vanderbilt University)

Stopping prepared but no longer relevant responses is a simple act of cognitive control that supports flexible and goal-directed behaviour. In the past two decades, response inhibition has received much attention across research domains. Cognitive psychologists and neuroscientists have explored the cognitive and neural mechanisms of response inhibition, developmental scientists have studied the 'rise and fall' of inhibitory control capacities across the life span, and clinical researchers have examined correlations between individual differences in response inhibition and behaviours such as substance abuse, overeating, and risk taking. In this chapter, we provide a selective review of recent behavioural and computational work on response inhibition (for a review of the neuroscience literature, see Chapter 16 by Forstmann & Alkemade in this volume).

In the response-inhibition literature, individual or group differences are often attributed to variation in the effectiveness of a single inhibitory control function. But many processes contribute to stopping an action. Response inhibition requires an interplay between three basic and computationally well-defined 'reactive' processes (signal detection, action selection, and suppression of motor output), which are regulated and influenced by sets of processes that take place on different timescales: outcome monitoring, advance preparation, rule acquisition and maintenance, associative learning, and development (Verbruggen, McLaren, & Chambers, 2014). Here we will discuss mainly 'reactive', 'proactive', and 'associatively mediated' forms of inhibitory control (for a review of key developmental transitions, see Munakata, Snyder, & Chatham, 2012).

How to Stop a Response

Reactive inhibitory control in response to changes in the environment or internal state is often studied in tasks such as the go/no-go task (Donders, 1868/1969) and the stop-signal task (Lappin & Eriksen, 1966; Logan & Cowan, 1984; Vince, 1948). In the go/no-go task, subjects are instructed to respond when a go stimulus appears (e.g., an 'O'), but to withhold their response when a no-go stimulus appears (e.g., an 'X'). In the stop-signal task, subjects perform a primary go task, such as responding to the identity of a stimulus (e.g., press left

when an 'O' appears, and right when an 'X' appears). On a minority of the trials, an extra visual or auditory signal appears after a variable stop-signal delay (SSD), instructing subjects to withhold the planned go response.

We will focus on the stop-signal paradigm because it allows researchers to estimate the covert latency of response inhibition: the stop-signal reaction time (SSRT). SSRT has become an established marker for reactive inhibitory control (Verbruggen, Chambers, & Logan, 2013), and many studies have used the stop-signal paradigm to develop detailed accounts of the cognitive processes underlying reactive inhibitory control and to examine what factors influence them.

The Race Between Going and Stopping

Since the 1980s, virtually all stop-signal research relied on a mathematical model because SSRT is unobservable without it. The model is the Logan and Cowan (1984) independent race model, which describes response inhibition as a race between a go process, triggered by the presentation of a go stimulus, and a stop process, triggered by the presentation of the stop signal. When the stop process finishes before the go process, response inhibition is successful, and no response is emitted (signal-inhibit trial); when the go process finishes before the stop process, response inhibition is unsuccessful, and the response is incorrectly emitted (signal-respond trial). The independent race model mathematically relates signal-respond reaction time (RT), no-stop-signal RT, and inhibition functions and provides several methods for estimating SSRT. Logan and Cowan wrote it as a race between generic finishing time distributions without specifying the processes that gave rise to the finishing times. This generality is a virtue because it allows the model to be applied to any subject population (including people, monkeys, and rats), effector (hands, eyes, feet), or experimental condition (for a review of models, see Verbruggen & Logan, 2009c). This generality is also a limitation: the model does not specify the underlying processes, so it can only describe and not explain reactive control. Fortunately, several models of the processes underlying the stop-signal task have appeared in the literature to address these issues.

General and Special Race Models Logan, Van Zandt, Verbruggen, and Wagenmakers (2014) extended the independent race model to include choice in the go process and the potential for choice in the stop process. They proposed a general race model that assumed choice was implemented as a race between the alternative responses, assuming one runner for each alternative response in the go task and one runner for the stop process. They proposed special race models in which the runners were characterised as specific stochastic accumulators. The best-fitting model represented each runner as a 'diffusion' to a single bound, characterised by three main parameters: non-decision time, the rate of stochastic accumulation, and the threshold on the accumulator. The first runner to reach its threshold determined the outcome of the trial. If it was a go response, then the response was executed. On no-stop-signal trials, it could be correct or erroneous. If the trial was a stop-signal trial (so there was a stop runner in the race), then the winning go response produced a signal-respond trial. If the trial was a stop-signal trial and the stop runner reached threshold first, the winning stop response produced a signal-inhibit trial. The general and special race models make stronger commitments to underlying processes than the original independent race model. Their assumptions may be controversial and may be falsified (Ratcliff & Smith, 2004; Teodorescu & Usher, 2013), but all specific assumptions are controversial, and falsifiability is a virtue. Importantly, the stronger commitments to underlying processes allow stronger inferences. For example, they allowed Logan et al. (2014) to test hypotheses about capacity sharing between stopping and going.

Research on dual tasks indicates that central processing capacity is limited, resulting in a performance decrement when stimuli for the two tasks appear close together in time (i.e., at the equivalent of short SSDs; Pashler, 1994). Stop and go processes do not seem to share capacity in this way (e.g., Logan, 1981; Logan & Burkell, 1986; Verbruggen, Schneider, & Logan, 2008). Yamaguchi, Logan, and Bissett (2012) measured SSRT in a dual-task experiment in which subjects had to stop one of two go tasks. They found that SSRT was unaffected by the temporal overlap of the two tasks and no different in single- and dual-task conditions. Logan et al. (2014) tested capacity sharing more rigorously by using the diffusion race model. They identified capacity limitations with processing rates in the stochastic accumulators, distinguishing between fixed capacity, in which the rate parameters for all the runners in the race sum to a constant, and unlimited capacity, in which the rate parameters for the different runners are unconstrained and independent (Bundesen, 1990; Townsend & Ashby, 1983). They tested these hypotheses about processing rates in a stop-signal experiment that manipulated the number of choice alternatives (2, 4, or 6) in the go task. They found evidence of strong capacity limitations in the go process: adding more go runners to the race slowed the rate at which each runner accumulated evidence. But they found no evidence of capacity sharing between the go process and the stop process: the rate of accumulation in the stop process was the same no matter how many choice alternatives there were in the go task. This converges on more qualitative analyses in previous research on simple stopping (Logan, 1981; Logan & Burkell, 1986; Yamaguchi et al., 2012) and contrasts with recent findings in selective stopping (reviewed below).

It is important to note that this conclusion could not be reached by applying the Logan and Cowan (1984) race model. The conclusion is based on changes in the parameters of the stochastic accumulators in the special race models. Below we present another application of the special race model. Thus, the additional complexity and stronger commitment to mechanism produces stronger conclusions.

Interactive Race Models and Their Alternatives   The turn of the century marked a substantial increase in interest in the neural mechanisms that underlie response inhibition. The Logan and Cowan (1984) race model has been useful in this pursuit, but its failure to specify the underlying computational mechanisms limits its ability to specify the underlying neural mechanisms. Boucher, Palmeri, Logan, and Schall (2007) developed the interactive race model, which explained what happened at the end of the race when the stop process won. They modelled the go process as a single diffusion to a threshold and modelled the stop process as a unit that inhibited the growth of go activation after an afferent delay following a stop signal. They made strong linking propositions that identified the go process with movement-related neurons and the stop process with fixation neurons in frontal eye fields, which control eye movements. The model accounted for behaviour as well as a version of the independent race model but in addition, accounted for firing rates recorded from movement and fixation neurons in monkeys performing a countermanding task (i.e., the saccadic version of the stop-signal task). Movement neuron activity is at baseline at the beginning of a trial and rises to a threshold value 10–20 ms before an eye movement occurs. On trials in which stopping is successful, the activity rises towards a threshold following the same trajectory as on no-stop-signal trials, and then falls abruptly before reaching the threshold. The transition from rising to falling coincides with SSRT measured from the monkeys' behaviour, as if some kind of change happens just before SSRT to drive down go activation and inhibit the response. Boucher et al. (2007) identified the change as inhibition from fixation neurons projected onto the movement neurons.

Logan, Yamaguchi, Schall, and Palmeri (2015) challenged the linking proposition connecting the mechanism of stopping to fixation neurons in frontal eye fields by

investigating a blocked input model that inhibits responses by blocking the input that drives them. Like Boucher et al. (2007), they implemented the go process as a single diffusion to a threshold, but unlike Boucher et al., they assumed that the stop process sets the growth rate for the go process to zero, removing the 'drive' that causes the rise to a threshold. They showed that this model accounted for behaviour as well as the interactive race model but fit the neural data better. Specifically, the interactive race model required very strong inhibition from stop (fixation) units to go (movement) units, and that strong inhibition produced much steeper decay in movement unit activation after SSRT than what was observed in the neurons. Logan et al. extended both models to account for steady-state fixation activity, constraining the activity of fixation neurons in the fixation period at the beginning of the trial. This improved the neural predictions of the interactive race model, but worsened behavioural predictions. The blocked input model fit the neural data just as well but fit the behavioural data better. Logan et al. also considered a boosted fixation model, which multiplied the activity in the fixation unit to stop responses and found that it fit the behavioural data as well as the blocked input model, though it predicted steady-state fixation activity following stop trials that was 3 to15 times as high as steady-state fixation activity before stop trials. On balance, the blocked input model did better than interactive race or boosted fixation. Nevertheless, it is possible that responses can be inhibited either by blocking the input or by inhibiting the growth of activation in movement neurons.

Cortical-Basal Ganglia Model   Imaging, lesion, and animal studies indicate that response inhibition relies on an extensive cortico-basal-ganglia network (see Chapter 16 by Forstmann & Alkemade in this volume). Wiecki and Frank (2013) proposed a computational model that instantiates this network, extending Frank's (2006) basal ganglia model to include cortical structures involved in response inhibition (see also Chapter 21 by Bhandari, Badre, & Frank in this volume). On go trials, saccadic responses are generated via the frontal eye fields and superior colliculus. On stop-signal trials, responses are inhibited via the 'hyperdirect' pathway: via the right inferior frontal gyrus and the subthalamic nucleus, to the substantia nigra pars reticulata, which then inhibits the superior colliculus. Responses are successfully stopped if the hyperdirect pathway inhibits the superior colliculus before it reaches the threshold at which it releases a response. This model correctly predicts neural dynamics in key structures during stop-signal tasks, correctly predicts behaviour, and correctly predicts the effects of lesions and drugs.

The Wiecki and Frank (2013) model is an important advance because it provides a platform for evaluating hypotheses about interactions between brain structures in response inhibition and other tasks that is strongly constrained by physiology, behaviour, and computation. However, the model is complex (it has to capture the interactions between the many neural structures) and difficult to evaluate by standard model fitting techniques from mathematical psychology. Ratcliff and Frank (2012) fit the diffusion model to the output of the basal ganglia model, showing how diffusion model parameters relate to activation of structures in the basal ganglia model. It would be interesting to fit the Logan et al. (2014) special race model to the output of the Wiecki and Frank (2013) model to relate its parameters to activation of its structures in stop-signal tasks.

Stopping Revisited?   The first step in successfully stopping a response is nearly always detecting the stop signal (e.g., a traffic light turning red; Verbruggen, McLaren, et al., 2014). The contribution of non-inhibitory processes is still largely neglected in the 'reactive' inhibitory control literature, even though computational work indicates that most of SSRT is occupied by afferent processes (Boucher et al., 2007; Logan et al., 2014, 2015; Salinas & Stanford,

2013), as indicated in several behavioural studies (e.g., Cavina-Pratesi, Bricolo, Prior, & Marzi, 2001; van den Wildenberg & van der Molen, 2004; Verbruggen, Stevens, & Chambers, 2014). For example, perceptual distractors substantially prolong SSRTs in a visual stop-signal task, especially when the visual stop signals occurred in the periphery (Verbruggen, Stevens, et al., 2014). Other studies have shown that stopping deficits in certain clinical populations could be due to impairments in signal detection rather than in inhibition per se (e.g., Bekker et al., 2005). Thus, a failure to detect the signal quickly can have important consequences and lead to 'stopping' deficits.

   More generally, the studies discussed in this and previous sections highlight the importance of focusing on the underlying processes, such as signal detection, rather than general and unitary control functions. They also have important consequences for the interpretation of the SSRT. SSRT describes the chain of processes involved in an act of control that results in a response being withheld (Logan et al., 2014; Verbruggen, McLaren, et al., 2014; Verbruggen & Logan, 2015). It includes the duration of perceptual, decisional, and (inhibitory) motor-related processes. Thus, in simple stop-signal tasks and their many variants, SSRT reflects more than the duration of a single neural inhibitory process, and researchers should consider at which processing stages differences between groups or conditions arise (see Verbruggen, McLaren, et al., 2014).

Are Stopping and Going Really Independent?

Beginning with Logan and Cowan (1984), models of response inhibition have always assumed that stopping and going are independent, at least for most of their durations. The models are based on two assumptions about independence: stochastic independence and context independence. Stochastic independence requires that the finishing times of the stop and go process on a given trial are uncorrelated (i.e., $P(\text{stop} < t \text{ and } \text{go} < t) = P(\text{stop} < t) \times P(\text{go} < t)$). This assumption simplifies the mathematics of the race considerably. Context independence assumes that the stop process is unaffected by the presence of the go process, and vice versa. This assumption allows go RT on no-stop-signal trials to be used as an estimate of go RTs on stop-signal trials, to predict signal-respond RTs and calculate SSRT. If context independence is violated, then estimates of SSRT are invalid. Thus, the assumptions should not be taken lightly. Recently, two lines of evidence have challenged the context independence assumption. One addresses selective versions of the stop signal task (Bissett & Logan, 2014; Verbruggen & Logan, 2015), and one addresses short SSDs in all stop tasks (Bissett, 2014).

Selective Stopping   Most stop-signal tasks involve 'simple' stopping: there is only one stop signal and only one response to it, namely, stopping the ongoing response. In selective stopping tasks, subjects perform some kind of discrimination before deciding whether to stop. In stimulus selective stopping tasks, the go task is accompanied by a stop signal or an ignore signal; these signals are usually in the same modality but differ along one or more stimulus dimensions. Subjects are instructed to stop the go response when the stop signal occurs and to continue the go response when the ignore signal occurs (e.g., Bissett & Logan, 2014). In motor selective stopping tasks, there is one stop signal, and subjects are instructed to stop one response (e.g., the left hand) but not the other (e.g., the right hand) when the stop signal occurs (e.g., De Jong, Coles, & Logan, 1995). In conditional motor selective stopping tasks, subjects make two concurrent responses to the go tasks, for example, pressing keys with the index fingers in response to one go stimulus and pressing keys with the middle fingers in response to another go stimulus. The stop signal instructs subjects to stop one of the two

responses (e.g., the left hand response) but continue with the other (e.g., the right hand response) when a stop signal occurs (e.g., Aron & Verbruggen, 2008).

Subjects succeed at stopping 'selectively' in all of these tasks, but their success comes with a cost: signal-respond RT is often longer than no-stop-signal RT, and RT for the responses that are not supposed to be inhibited is longer than no-stop-signal RT in most experiments (for a review and a re-analysis of published data, see Bissett & Logan, 2014). The longer signal-respond RTs indicate violations of context independence, which invalidates conclusions based on the application of the independent race model to the data and challenges the common assumption that selective stopping increases the complexity of the stop process without affecting the go process (e.g., van de Laar, van den Wildenberg, van Boxtel, & van der Molen, 2010). The longer RTs on ignore trials suggest that going and stopping share processing capacity (Bissett & Logan, 2014; Verbruggen & Logan, 2015); in some cases, subjects may even stop all responses completely first and then restart them if the trial requires it. This challenges the use of the ignore signal in selective stopping tasks as a control for the requirements to attend to the stop signal (e.g., Sharp et al., 2010).

Verbruggen and Logan (2015) tested the hypothesis that the go and stop processes share capacity in selective stopping tasks by manipulating the consistency of mapping between signals and the requirement to stop or ignore in response to the signal. In consistent mapping conditions, each signal played the same role throughout the experiment; in varied mapping conditions, the role changed repeatedly over the course of the experiment. Following Shiffrin and Schneider (1977) and others, we assumed that varied mapping conditions would demand more capacity than the consistent mapping conditions, and so should produce larger violations of context independence. That is what we found.

These selective stopping results are interesting in contrast with simple stopping, where increasing the capacity demands of the go process has no effect on the stop process (see above). We propose that this is due to the low selection demands in standard stop-signal tasks. This does not imply that capacity sharing can never occur in these tasks. The stop rate parameters depend on the discriminability, intensity, and modality of the stop signal (e.g., van der Schoot, Licht, Horsley, & Sergeant, 2005), which could be interpreted as a capacity limitation (Logan et al., 2014). Furthermore, the influence of visual distractors on stopping (see above) is consistent with the idea that stimuli compete for limited processing capacity (Bundesen, 1990; Duncan, 2006). Finally, SSRT is prolonged on incongruent trials in interference control tasks (e.g., Verbruggen, Liefooghe, & Vandierendonck, 2004). This 'functional dependence' could also be interpreted as a capacity limitation. Thus, it seems that under certain circumstances, capacity sharing may occur in simple stop-signal tasks.

Short SSDs   Bissett (2014) looked for violations of context independence in simple stop-signal tasks, and found strong violations at SSDs shorter than 250 ms in almost every experiment he analysed. These violations were unexpected from the stop-signal literature, where the independent race model has reigned supreme since the 1980s, but they would be expected from the dual-task literature (see above). These violations may not have been apparent in published stop-signal studies. Many studies now use a tracking procedure that adjusts SSDs to produce successful inhibition on 50% of the trials. This procedure usually produces a bell-shaped distribution of SSDs with a mode well away from the short SSDs where context independence is violated. Moreover, researchers rarely report performance as a function of SSD. Instead, they often collapse across SSD, masking possible violations at short SSDs. Researchers should not try to avoid violations of context independence by

excluding short SSDs from their experiments. Subjects are sensitive to the range of SSDs and routinely slow their go task performance when SSDs are long (Logan, 1981). Instead, we suggest that investigators use a broad range of SSDs (e.g., by using the tracking procedure) and then remove the data from SSDs shorter than 200 ms before calculating SSRT. At present, these recommendations are only provisional. More research needs to be done before we can make stronger recommendations.

Balancing Cognitive Demands

Proactive Control

In response-inhibition tasks, people often adjust their behaviour in anticipation of future acts of inhibitory control (*proactive inhibitory control*). Proactive adjustments are made at the beginning of a block or on a trial-by-trial basis, suggesting a flexible cognitive system that adjusts itself quickly in response to novel contextual information (Verbruggen & Logan, 2009b). Proactive inhibitory control has received much attention in the cognitive neuroscience literature (for reviews, see Aron, 2011; Stuphorn & Emeric, 2012), and researchers have explored to what extent individual differences and developmental trends could reflect differences in the deployment of proactive control (e.g., Chevalier, Kelsey, Wiebe, & Espy, 2014). Here we will primarily focus on studies examining which control adjustments are made.

Proactive Adjustments of Attentional and Response Settings   People adjust response settings when they anticipate a stop signal. In a series of experiments, we presented cues indicating whether signals were relevant for the next trial(s). We hypothesised that subjects balance going quickly and stopping by adjusting response thresholds in the go task (see above). Increasing the response threshold increases the finishing time of the go process, but decreases both the probability of an incorrect go response on no-stop-signal trials and the probability of a response on stop-signal trials. Our behavioural results and subsequent diffusion model fits supported the response-threshold account (Logan et al., 2014; Verbruggen & Logan, 2009b).

Further studies suggest that proactive inhibitory control can also directly modulate motor activity (e.g., Jahfari, Stinear, Claffey, Verbruggen, & Aron, 2010; Lo, Boucher, Paré, Schall, & Wang, 2009). For example, Lo et al. (2009) observed high tonic activation of fixation neurons before the onset of a stop signal in a countermanding task, and argued that this was activation modulated by a top-down proactive control signal. Suppression of prestimulus motor activity is also observed in go/no-go and choice RT paradigms, which points to a general proactive control mechanism that prevents premature responding (Frank, 2006).

Most studies have focused on preparatory adjustments of response- or motor-related processes. However, in some tasks, perceptual or attentional settings are also adjusted when subjects anticipate stop signals. The distractor study mentioned above contrasted no-stop-signal blocks with stop-signal blocks in which visual stop signals were presented in the centre of the screen and blocks in which the signals could occur in the periphery (Verbruggen, Stevens, et al., 2014). On half of the no-stop-signal trials in each block type, perceptual distractors were presented throughout the trial. These distractors had the largest effect on go responses in the noncentral signal blocks compared with central-signal blocks and no-stop-signal blocks, suggesting that subjects widened their attentional focus for

detecting the signal in the periphery, making them more vulnerable to the distractors on no-stop-signal trials. Another study measured event-related potentials in the go task (Elchlepp, Lavric, Chambers, & Verbruggen, 2016). Both early and late go-ERP components were modulated when subjects expected a stop signal to occur. Thus, proactive inhibitory control involves adjusting the settings of both perceptual and response processes to balance the competing demands of the go and stop tasks (Elchlepp et al., 2016). Importantly, similar modulations of ERP components are observed in non-inhibitory tasks, suggesting similarities between various forms of proactive control (Braver, 2012; Chapter 9 by Chiew & Braver in this volume).

Influences of Proactive Control on Stopping   Proactive control does not necessarily influence SSRT (Logan & Burkell, 1986; Verbruggen & Logan, 2009b). Whether or not SSRT is influenced may depend on how proactive control adjustments are implemented (e.g., response slowing vs. enhanced monitoring; Wiecki & Frank, 2013). Proactive inhibitory control can also influence *how* people stop their responses. Brain stimulation studies suggest that both task-relevant and irrelevant muscles are suppressed on stop-signal trials, indicating that stopping has global effects on the motor system (Aron, 2011). However, stopping becomes more selective when it can be prepared in advance (e.g., Aron & Verbruggen, 2008; Smittenaar, Guitart-Masip, Lutti, & Dolan, 2013).

Proactive control could even lead to a 'prepared control reflex' (Verbruggen, Best, Bowditch, Stevens, & McLaren, 2014; Verbruggen, McLaren, et al., 2014). When attention is proactively allocated and responses are prepared, goal-directed behaviour may not require much control anymore (see also Chapter 7 by Hommel in this volume); instead, responses could be activated easily by stimuli in the environment (for a recent review, see Meiran, Cole, & Braver, 2012). Response inhibition can be triggered by task-irrelevant primes (e.g., van Gaal, Ridderinkhof, van den Wildenberg, & Lamme, 2009; Verbruggen & Logan, 2009a), but these priming effects are only observed in contexts in which subjects are instructed to stop occasionally (Chiu & Aron, 2014; Verbruggen & Logan, 2009a). These findings are consistent with the prepared reflex idea: once subjects have proactively adjusted attentional and response settings, the 'stop response' can be activated easily by both task-relevant and task-irrelevant information in the environment. Note that the 'prepared reflex idea' might also explain why going and stopping do not share (much) capacity in standard stop-signal tasks (see above).

Post-Signal Slowing

The proactive control literature indicates that people set detection and response criteria at the beginning of a trial or run of trials. These criteria are further adjusted after each trial. Response latencies are often slower after stop-signal trials than after no-stop-signal trials (Bissett & Logan, 2011, 2012; Emeric et al., 2007; Nelson, Boucher, Logan, Palmeri, & Schall, 2010; Rieger & Gauggel, 1999; Verbruggen & Logan, 2008b; Verbruggen, Logan, Liefooghe, & Vandierendonck, 2008). Bissett and Logan (2011) contrasted several accounts of post-stop-signal slowing, and found most support for a strategic adjustment account that proposes that stop-signal presentation encourages subjects to shift priority from the go task to the stop task. Such a shift produces longer response latencies after a signal trial and can reduce SSRT (Bissett & Logan, 2012). However, the SSRT reduction is only observed when the stop-signal modality remains the same (Bissett & Logan, 2012). This could indicate that the post-signal strategic adjustments influence attentional settings. Alternatively, the modality specificity could indicate stop-signal priming (see below).

Contributions From Memory and Long-Term After-Effects of Stopping

Stimulus-Specific Sequential Effects

The post-signal slowing is usually more pronounced when the stimulus, stimulus category, or stimulus features of the previous trial are repeated (Bissett & Logan, 2011; Rieger & Gauggel, 1999; Verbruggen, Logan, et al., 2008). We have attributed this stimulus-specific slowing to the retrieval of stimulus–stop associations: a go stimulus becomes associated with a 'stop' representation on a stop trial; when it is repeated on the next go trial, the stop representation is activated via associative retrieval, and this will suppress the go response (Verbruggen & Logan, 2008b; Verbruggen, Logan, et al., 2008). Activation of these stop representations has global effects on responding (Giesen & Rothermund, 2014), which suggests that they are general ('stop responding') rather than response-specific (e.g., 'stop the left response').

Associative Learning and Automatic Inhibition

The stimulus-specific stop effects are observed up to 20 trials after the presentation of the stop signal (Verbruggen & Logan, 2008b). Such long-term associations may support the development of automatic response inhibition (Logan, 1988, 1990). In a series of experiments, we trained subjects to stop their response to a specific stimulus, and then reversed the stimulus–stop mappings in a test phase. In this test phase, subjects were slower to respond to stimuli previously associated with stopping compared with stimuli that they had not seen before or stimuli that were inconsistently associated with going and stopping (Verbruggen & Logan, 2008a). On the basis of this finding, we proposed the 'automatic inhibition' hypothesis: inhibitory control in go/no-go and stop-signal tasks can be triggered automatically via the retrieval of stimulus–stop associations from memory.

Subsequent work indicated that different associations can be learned. For example, subjects can learn stimulus-signal associations rather than stimulus–stop associations (Verbruggen, Best, et al., 2014). Such associations between the stop items and the stop signal will prime the representation of the stop-signal detection rather than the stop goal or stop response. As discussed above, signal detection plays a critical role in successful stopping. Thus, by priming the representation of the stop signal, learning could lead to improvements in stopping performance without influencing responding on go trials. The signal learning idea can also account for the modality specificity of post-stop signal improvements in SSRT (see above).

Stopping responses to stimuli can also lead to devaluation of these stimuli and reduce their behavioural incentive (e.g., Kiss, Raymond, Westoby, Nobre, & Eimer, 2008; Wessel, O'Doherty, Berkebile, Linderman, & Aron, 2014). Research on Pavlovian conditioning suggests links between an aversive system and avoidance behaviour (withdrawal, suppression) and between an appetitive system and approach behaviour (Dickinson & Balleine, 2002). Thus, when subjects have to stop their response to a specific stimulus, a link between this stimulus and the aversive/avoidance system may be created, impacting their motivational value. To account for these findings, McLaren and Verbruggen (2016) and Verbruggen, Best, et al. (2014) proposed an associative stop system that combines elements of associative learning theories, the interactive race model, and a Pavlovian model of motivational systems.

In sum, associative learning plays an important role in response-inhibition tasks. It is important to note that the expression of stimulus-specific training effects strongly depends on attentional settings (Best, Lawrence, Logan, McLaren, & Verbruggen, 2016). After learning,

subjects also form expectations about the occurrence of a stop signal, which influences subsequent performance. These findings suggest a strong interplay between learning and 'top-down' control processes (Best et al., 2016).

Applications of Automatic Inhibition

Automatic inhibition can have general effects on behaviour. For example, consistent pairing of food-related pictures to stopping in a go/no-go or stop-signal paradigm reduces subsequent food consumption (e.g., Houben & Jansen, 2011; Lawrence, Verbruggen, Morrison, Adams, & Chambers, 2015), and may even lead to weight loss (Lawrence et al., 2015; Veling, van Koningsbruggen, Aarts, & Stroebe, 2014). Recent meta-analyses indicate that no-go/stop training has small but robust effects on food and alcohol consumption (Allom, Mullan, & Hagger, 2015; Jones et al., 2016). Go/no-go training may have larger effects on subsequent appetitive behaviour than stop-signal training. Go/no-go tasks often use a more consistent stimulus–stop pairing, and the probability of responding is usually lower than in stop-signal tasks, which could influence the strength of stimulus–stop associations (Verbruggen & Logan, 2008a, 2008b).

The stimulus-specific stop training opens promising new avenues for treatment. However, more work is required to determine how the training protocols can be optimised to produce large and long-lasting effects. This will also require a careful study of the underlying mechanisms and a better understanding of how task characteristics influence what is learned.

General Learning and Carry-Over Effects

Some studies have also explored whether people can generally learn to stop their responses and whether practice effects can transfer to other domains. The results are mixed. Some studies found that SSRT decreased over practice (e.g., Berkman, Kahn, & Merchant, 2014; Chevalier, Chatham, & Munakata, 2014; Logan & Burkell, 1986), whereas others failed to find such an SSRT decrease (e.g., Enge et al., 2014) General stop-signal task training may induce a shift from reactive to proactive control (Berkman et al., 2014; Chevalier, Chatham, et al., 2014). This idea could help to reconcile the inconsistent training results, as not all forms of proactive control lead to SSRT improvements (see above).

General stop training and encouraging people to be cautious in stop-signal blocks can reduce risk taking in gambling tasks (Stevens et al., 2015; Verbruggen, Adams, & Chambers, 2012) and alcohol-seeking behaviour (Jones, Cole, Goudie, & Field, 2011). However, these effects are small and short-lived (Stevens et al., 2015; Verbruggen et al., 2013), which could potentially explain why some studies failed to observe far-transfer effects (e.g., Enge et al., 2014).

Conclusions

In the past two decades, we have seen an exponential increase in the number of studies examining the neurocognitive processes underlying response inhibition and correlations with behaviours outside the lab (Verbruggen, Chambers, et al., 2013). This work has led to detailed accounts of reactive inhibitory control (Logan et al., 2014). More recently, several studies have also shown how people adjust settings of lower-level perceptual and response systems to find a balance between going quickly and stopping in response to changes in the environment or internal state. Evidence that learning plays a critical role when stopping a

response has challenged theory and opened up an exciting new literature on the training and transfer of response inhibition. Above all, the work reviewed in this chapter highlights again that stopping or withholding a response relies on an interplay between many processes that take place on different time scales.

Acknowledgement

References

Allom, V., Mullan, B., & Hagger, M. (2015). Does inhibitory control training improve health behaviour? A meta-analysis. *Health Psychol Rev*, 1–38.

Aron, A. R. (2011). From reactive to proactive and selective control: Developing a richer model for stopping inappropriate responses. *Biological Psychiatry*, *69*, e55–68.

Aron, A. R., & Verbruggen, F. (2008). Stop the presses: Dissociating a selective from a global mechanism for stopping. *Psychological Science*, *19*, 1146–1153.

Bekker, E. M., Overtoom, C. C., Kooij, J. J., Buitelaar, J. K., Verbaten, M. N., & Kenemans, J. L. (2005). Disentangling deficits in adults with attention-deficit/hyperactivity disorder. *Archives of General Psychiatry*, *62*, 1129–1136.

Berkman, E. T., Kahn, L. E., & Merchant, J. S. (2014). Training-induced changes in inhibitory control network activity. *The Journal of Neuroscience*, *34*, 149–157.

Best, M., Lawrence, N. S., Logan, G. D., McLaren, I. P. L., & Verbruggen, F. (2016). Should I Stop or Should I Go? The Role of Associations and Expectancies. *Journal of Experimental Psychology: Human Perception and Performance*, *42*(1), 115–137.

Bissett, P. G. (2014). *Evaluating the independent race model for the stop signal paradigm: Context independence is violated at short stop signal delays*. Unpublished dissertation, Vanderbilt University.

Bissett, P. G., & Logan, G. D. (2011). Balancing cognitive demands: Control adjustments in the stop-signal paradigm. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *37*, 392–404.

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

Bissett, P. G., & Logan, G. D. (2014). Selective stopping? Maybe not. *Journal of Experimental Psychology: General*, *143*, 455–472.

Boucher, L., Palmeri, T. J., Logan, G. D., & Schall, J. D. (2007). Inhibitory control in mind and brain: An interactive race model of countermanding saccades. *Psychological Review*, *114*, 376–397.

Braver, T. S. (2012). The variable nature of cognitive control: A dual mechanisms framework. *Trends in Cognitive Sciences*, *16*, 106–113.

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

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

Chevalier, N., Chatham, C. H., & Munakata, Y. (2014). The practice of going helps children to stop: The importance of context monitoring in inhibitory control. *Journal of Experimental Psychology: General*, *143*(3), 959–965.

Chevalier, N., Kelsey, K. M., Wiebe, S. A., & Espy, K. A. (2014). The temporal dynamic of response inhibition in early childhood: An ERP study of partial and successful inhibition. *Developmental Neuropsychology*, *39*, 585–599.

Chiu, Y. C., & Aron, A. R. (2014). Unconsciously triggered response inhibition requires an executive setting. *Journal of Experimental Psychology: General*, *143*, 56–61.

De Jong, R., Coles, M. G. H., & Logan, G. D. (1995). Strategies and mechanisms in nonselective and selective inhibitory motor control. *Journal of Experimental Psychology: Human Perception and Performance*, *21*, 498–511.

Dickinson, A., & Balleine, B. (2002). The role of learning in the operation of motivational systems. In C. R. Gallistel (Ed.), *Stevens' handbook of experimental psychology* (Vol. *3*, 3rd ed., pp. 497–534). New York, NY: Wiley.

Donders, F. C. (1868/1969). On the speed of mental processes. *Acta Psychologica*, *30*, 412–431.

Duncan, J. (2006). EPS mid-career award 2004: Brain mechanisms of attention. *Quarterly Journal of Experimental Psychology*, *59*, 2–27.

Elchlepp, H., Lavric, A., Chambers, C. D., & Verbruggen, F. (2016). Proactive inhibitory control: A general biasing account. *Cognitive Psychology*, *86*, 27–61.

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

Enge, S., Behnke, A., Fleischhauer, M., Küttler, L., Kliegel, M., & Strobel, A. (2014). No evidence for true training and transfer effects after inhibitory control training in young healthy adults. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *40*, 987–1001.

Frank, M. J. (2006). Hold your horses: A dynamic computational role for the subthalamic nucleus in decision making. *Neural Networks*, *19*, 1120–1136.

Giesen, C., & Rothermund, K. (2014). You better stop! Binding 'stop' tags to irrelevant stimulus features. *Quarterly Journal of Experimental Psychology*, *67*, 809–832.

Houben, K., & Jansen, A. (2011). Training inhibitory control. A recipe for resisting sweet temptations. *Appetite*, *56*, 345–349.

Jahfari, S., Stinear, C. M., Claffey, M., Verbruggen, F., & Aron, A. R. (2010). Responding with restraint: What are the neurocognitive mechanisms? *Journal of Cognitive Neuroscience*, *22*(7), 1479–1492.

Jones, A., Cole, J., Goudie, A., & Field, M. (2011). Priming a restrained mental set reduces alcohol-seeking independently of mood. *Psychopharmacology*, *218*(3), 557–565.

Jones, A., Di Lemma, L. C. G., Robinson, E., Christiansen, P., Nolan, S., Tudur-Smith, C., & Field, M. (2016). Inhibitory control training for appetitive behaviour change: A meta-analytic investigation of mechanisms of action and moderators of effectiveness. *Appetite*, *97*, 16–28.

Kiss, M., Raymond, J. E., Westoby, N., Nobre, A. C., & Eimer, M. (2008). Response inhibition is linked to emotional devaluation: Behavioural and electrophysiological evidence. *Frontiers in Human Neuroscience*, *2*, 13.

Lappin, J. S., & Eriksen, C. W. (1966). Use of a delayed signal to stop a visual reaction-time response. *Journal of Experimental Psychology*, *72*, 805–811.

Lawrence, N. S., O'Sullivan, J., Parslow, D., Javaid, M., Adams, R. C., Chambers, C. D., … Verbruggen, F. (2015). Training response inhibition to food is associated with weight loss and reduced energy intake. *Appetite*, *95*, 17–28.

Lawrence, N. S., Verbruggen, F., Morrison, S., Adams, R. C., & Chambers, C. D. (2015). Stopping to food can reduce intake. Effects of stimulus-specificity and individual differences in dietary restraint. *Appetite*, *85*, 91–103.

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

Logan, G. D. (1981). Attention, automaticity, and the ability to stop a speeded choice response. In J. Long & A. D. Baddeley (Eds.), *Attention and performance IX* (pp. 205–222). Hillsdale, NJ: Erlbaum.

Logan, G. D. (1988). Toward an instance theory of automatization. *Psychological Review*, *95*, 492–527.

Logan, G. D. (1990). Repetition priming and automaticity: Common underlying mechanisms? *Cognitive Psychology*, *22*, 1–35.

Logan, G. D., & Burkell, J. (1986). Dependence and independence in responding to double stimulation—A comparison of stop, change, and dual-task paradigms. *Journal of Experimental Psychology: Human Perception and Performance*, *12*, 549–563.

Logan, G. D., & Cowan, W. B. (1984). On the ability to inhibit thought and action: A theory of an act of control. *Psychological Review*, *91*(3), 295–327.

Logan, G. D., Van Zandt, T., Verbruggen, F., & Wagenmakers, E. J. (2014). On the ability to inhibit thought and action: General and special theories of an act of control. *Psychological Review*, *121*, 66–95.

Logan, G. D., Yamaguchi, M., Schall, J. D., & Palmeri, T. J. (2015). Inhibitory control in mind and brain 2.0: Blocked input models of saccadic countermanding. *Psychological Review*, *122*, 115–147.

McLaren, I. P. L., & Verbruggen, F. (2016). Association, inhibition, and action. In R. Murphy & R. Honey (Eds.), *The Wiley Handbook on the Cognitive Neuroscience of Learning* (pp. 489–514). Chichester, England: Wiley-Blackwell.

Meiran, N., Cole, M. W., & Braver, T. S. (2012). When planning results in loss of control: Intention-based reflexivity and working-memory. *Frontiers in Human Neuroscience*, *6*, 104.

Munakata, Y., Snyder, H. R., & Chatham, C. H. (2012). Developing cognitive control: Three key transitions. *Current Directions in Psychological Science*, *21*, 71–77.

Nelson, M. J., Boucher, L., Logan, G. D., Palmeri, T. J., & Schall, J. D. (2010). Nonindependent and nonstationary response times in stopping and stepping saccade tasks. *Attention, Perception & Psychophysics*, *72*, 1913–1929.

Ratcliff, R., & Frank, M. J. (2012). Reinforcement-based decision making in corticostriatal circuits: Mutual constraints by neurocomputational and diffusion models. *Neural Computation*, *24*, 1186–1229.

Ratcliff, R., & Smith, P. L. (2004). A comparison of sequential sampling models for two-choice reaction time. *Psychological Review*, *111*, 333–367.

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

Salinas, E., & Stanford, T. R. (2013). The countermanding task revisited: Fast stimulus detection is a key determinant of psychophysical performance. *Journal of Neuroscience*, *33*(13), 5668–5685.

Schevernels, H., Bombeke, K., Van der Borght, L., Hopf, J.-M., Krebs, R. M., & Boehler, C. N. (2015). Electrophysiological evidence for the involvement of proactive and reactive control in a rewarded stop-signal task. *NeuroImage*, *121*, 115–125.

Sharp, D. J., Bonnelle, V., De Boissezon, X., Beckmann, C. F., James, S. G., Patel, M. C., & Mehta, M. A. (2010). Distinct frontal systems for response inhibition, attentional capture, and error processing. *Proceedings of the National Academy of Sciences of the United States of America*, *107*, 6106–6111.

Shiffrin, R. M., & Schneider, W. (1977). Controlled and automatic human information processing: II. Perceptual learning, automatic attending, and a general theory. *Psychological Review*, *84*, 127–190.

Smittenaar, P., Guitart-Masip, M., Lutti, A., & Dolan, R. J. (2013). Preparing for selective inhibition within frontostriatal loops. *The Journal of Neuroscience*, *33*, 18087–18097.

Stevens, T., Brevers, D., Chambers, C. D., Lavric, A., McLaren, I. P., Mertens, M.,…Verbruggen, F. (2015). How does response inhibition influence decision making when gambling? *Journal of Experimental Psychology: Applied*, *21*, 15–36.

Stuphorn, V., & Emeric, E. E. (2012). Proactive and reactive control by the medial frontal cortex. *Frontiers in Neuroengineering*, *5*, 9.

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

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

van de Laar, M. C., van den Wildenberg, W. P. M., van Boxtel, G., & van der Molen, M. W. (2010). Processing of global and selective stop signals: Application of Donders' subtraction method to stop-signal task performance. *Experimental Psychology*, *57*, 149–159.

van den Wildenberg, W. P., & van der Molen, M. W. (2004). Additive factors analysis of inhibitory processing in the stop-signal paradigm. *Brain and Cognition*, *56*, 253–266.

van der Schoot, M., Licht, R., Horsley, T. M., & Sergeant, J. A. (2005). Effects of stop signal modality, stop signal intensity and tracking method on inhibitory performance as determined by use of the stop signal paradigm. *Scandinavian Journal of Psychology*, *46*, 331–341.

van Gaal, S., Ridderinkhof, K. R., van den Wildenberg, W. P., & Lamme, V. A. (2009). Dissociating consciousness from inhibitory control: Evidence for unconsciously triggered response inhibition in the stop-signal task. *Journal of Experimental Psychology: Human Perception and Performance*, *35*, 1129–1139.

Veling, H., van Koningsbruggen, G. M., Aarts, H., & Stroebe, W. (2014). Targeting impulsive processes of eating behavior via the internet. Effects on body weight. *Appetite*, *78*, 102–109.

Verbruggen, F., & Logan, G. D. (2008a). Automatic and controlled response inhibition: Associative learning in the go/no-go and stop-signal paradigms. *Journal of Experimental Psychology: General*, *137*, 649–672.

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

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

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

Verbruggen, F., & Logan, G. D. (2009c). Models of response inhibition in the stop-signal and stop-change paradigms. *Neuroscience and Biobehavioral Reviews*, *33*, 647–661.

Verbruggen, F., & Logan, G. D. (2015). Evidence for capacity sharing when stopping. *Cognition*, *142*, 81–95.

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

Verbruggen, F., Adams, R. C., van't Wout, F., Stevens, T., McLaren, I. P., & Chambers, C. D. (2013). Are the effects of response inhibition on gambling long-lasting? *PLoS ONE*, *8*, e70155.

Verbruggen, F., Best, M., Bowditch, W. A., Stevens, T., & McLaren, I. P. L. (2014). The inhibitory control reflex. *Neuropsychologia*, *65*, 263–278.

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*(3), 352–362.

Verbruggen, F., Liefooghe, B., & Vandierendonck, A. (2004). The interaction between stop signal inhibition and distractor interference in the flanker and stroop task. *Acta Psychologica*, *116*, 21–37.

Verbruggen, F., Logan, G. D., Liefooghe, B., & Vandierendonck, A. (2008). Short-term aftereffects of response inhibition: Repetition priming or between-trial control adjustments? *Journal of Experimental Psychology: Human Perception and Performance*, *34*, 413–426.

Verbruggen, F., McLaren, I. P. L., & Chambers, C. D. (2014). Banishing the control homunculi in studies of action control and behavior change. *Perspectives on Psychological Science*, *9*, 497–524.

Verbruggen, F., Schneider, D. W., & Logan, G. D. (2008). How to stop and change a response: The role of goal activation in multitasking. *Journal of Experimental Psychology. Human Perception and Performance*, *34*(5), 1212–28.

Verbruggen, F., Stevens, T., & Chambers, C. D. (2014). Proactive and reactive stopping when distracted: An attentional account. *Journal of Experimental Psychology: Human Perception and Performance*, *40*, 1295–1300.

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

Wessel, J. R., O'Doherty, J. P., Berkebile, M. M., Linderman, D., & Aron, A. R. (2014). Stimulus devaluation induced by stopping action. *Journal of Experimental Psychology: General*, *143*, 2316–2329.

Wiecki, T. V., & Frank, M. J. (2013). A computational model of inhibitory control in frontal cortex and basal ganglia. *Psychological Review*, *120*, 329–355.

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