Associatively-Mediated Stopping: Training Stimulus-Specific Inhibitory Control

William A. Bowditch Frederick Verbruggen & I.P.L. McLaren
School of Psychology, University of Exeter

Response inhibition is often considered to be a deliberate act of cognitive control. However, our and other research suggests that by repeatedly pairing stimuli with an inhibitory response, inhibition can become automatized. Currently, relatively little research has focused on the nature of the associative structure that underpins stimulus-specific inhibitory training. This paper investigated what associations can be learnt in stop-signal training tasks, distinguishing between indirect priming of the stop signal and direct activation of a stop response. We employed a novel paradigm, where colored cues are stochastically paired with a number of stop-signals, and demonstrated that cues consistently paired with stopping reduced commission errors and slowed reaction times. Furthermore, we showed that manipulating the pairings between stimuli and stop signals, in a manner that favored the formation of stimulus-stop associations, produced enhanced stop learning effects on reaction time, but not probability of responding. Our results suggest that perceptual processes supporting signal detection (priming) as well as inhibitory processes are involved in inhibitory control training, and that inhibition training may benefit from reducing the contingency between stimuli and stop-signals.

Keywords: Inhibitory control; Stop-signal training; Response inhibition; Associative learning; Stimulus specific training.

The ability to exert executive control over our behavior is key; without this fundamental ability we would haphazardly engage in whatever behaviors are prompted by our current environment (Brazzelli & Spinnler, 1998; Lhermitte, 1983; O'Reilly, 2006). In this paper we take control to be a multifaceted concept (Miyake et al., 2000), of which inhibition is a core component that facilitates goal-directed behavior through the suppression of otherwise prepotent responses (Verbruggen & Logan, 2008c)[1].

Theories of executive control typically ascribe inhibitory control to a deliberate top-down process that selectively modulates bottom-up environment-driven processes (Miller & Cohen, 2001; O'Reilly, 2006; Ridderinkhof, van den Wildenberg, Segalowitz, & Carter, 2004; Verbruggen & Logan, 2008c). However, a growing body of research suggests that response inhibition, in certain situations, can itself operate automatically in a bottom-up stimulus-driven fashion akin to the automaticity observed in learnt response execution (Logan 1988; Schneider and Shiffrin 1977; Shiffrin and Schneider 1977; for a review see: Verbruggen, Best, Bowditch, Stevens, and McLaren 2014). Verbruggen and Logan (2008a, Experiment 5) demonstrated that, by pairing stimuli with an inhibitory response, inhibitory processes could become somewhat automatized. Participants were trained on a stop-signal task, where the stimulus category (i.e. living or non-living) determined the correct response (i.e. left or right). Subsets of stimuli were repeatedly paired with the requirement to respond or withhold the response throughout training, allowing for the formation of stimulus-specific associations. Upon test, when the stimulus mappings were reversed, participants were slower to respond to stimuli previously paired with stopping in comparison to stimuli associated with responding. In a similar paradigm, Lenartowicz, Verbruggen, Logan, and Poldrack (2011) demonstrated that stimulus-specific slowing on

Acknowledgements: This work was supported by a studentship from the Economic and Social Research Council [ES/J50015X/1] to WAB, an Economic and Social Research Council Grant (ES/J00815X/1) to FV and IPLM, and a starting grant from the European Research Council (ERC) to FV under the European Union's Seventh Framework Program (FP7/2007-2013)/ ERC Grant Agreement No. 312445.

Corresponding author: William A. Bowditch, wb221@exeter.ac.uk, Washington Singer Laboratories, University of Exeter, Perry Road, Exeter, EX4 4QG

[1] Whilst the concept of inhibition is often invoked in explaining behaviour (or the lack of it), its direct involvement in many tasks is debatable (MacLeod, Dodd, Sheard, Wilson, & Bibi, 2003). There is, however, little doubt that inhibitory control is directly involved in the cancellation of an already initiated motor response, and thus our research employs tasks that require the active suppression of a prepotent motor response; such as the go/no-go (Donders, 1969) and stop-signal tasks (Verbruggen & Logan, 2008a).

no-signal trials was accompanied by increased activation of the right inferior frontal gyrus (rIFG), a region typically implicated in response inhibition (Aron, Robbins, & Poldrack, 2004, 2014; Chambers, Garavan, & Bellgrove, 2009). This last result suggests that the stimuli in question were associated with some "stop center". In this paper we investigate the mechanisms by which stimulus-specific stop effects are learnt within the stop-signal paradigm, differentiating between perceptual and response processes, with a view to enhance inhibitory control training paradigms.

Verbruggen, McLaren, and Chambers (2014) have proposed a theoretical framework that ascribes action control to three fundamental cognitive processes: signal detection, action selection and action execution. The present paper explores the role played by signal detection and its interaction with associative learning in stop-signal paradigms. Signal-detection is undoubtedly essential for successful response inhibition; computational models suggest that a significant portion of stop-signal reaction time (SSRT) reflects non-inhibitory detection processes (Boucher, Palmeri, Logan, & Schall, 2007; Logan, Van Zandt, Verbruggen, & Wagenmakers, 2014) and increasing the difficulty of signal detection by introducing irrelevant perceptual distractors impairs SSRT (particularly when the stop signal could occur in the periphery; Verbruggen, Stevens, and Chambers 2014). Furthermore, there is some evidence that stimulus detection may indeed be enhanced, in an implicit, associative manner, through repeated pairing; both detection and recognition are augmented in visual search when distractors (which act as cues) consistently co-occur with the same target stimulus, even when location varies randomly (Chun & Jiang, 1999). Our point of departure in this paper is to note that by definition, the signal to stop and the act of stopping are entirely confounded within the stop-signal paradigm. Thus, when a cue that consistently precedes a signal trial is presented, there are at least two events that can be predicted: firstly, the imminent presentation of the stop signal and secondly, the impending requirement to withhold ones response. Crucially, these consequences have rather different cognitive requirements. The former does not require the involvement of motor inhibition and operates at a perceptual signal-detection level, whilst the latter does require inhibition or preparation for its initiation.

We can distinguish between at least two possible types of associatively-mediated pathway to action inhibition (see Figure 1): One makes use of a *direct* associative link between the cue associated with stopping and some representation of stopping itself, variously termed a "stop center" or "stop goal" (we shall use the former designation). The other, *indirect* associative pathway, operates by means of a link between the cue to the representation of the stop signal used in the experiment, and exploits the ability of that (active) representation to inhibit ongoing actions; for simplicity, we have assumed that the latter is achieved via a link from the signal to the same representation of stopping utilised by the direct pathway, but we acknowledge that this does not have to be the case. Note that these direct and indirect associative pathways do not necessarily map onto the direct and indirect cortical-subcortical pathways (Nambu, Tokuno, & Takada, 2002).

Both associative pathways are capable of producing associatively-mediated stopping effects, by which we mean producing slowing of RTs when that cue is presented on a no-signal trial and/or reduced errors of commission (lower p(respond|signal)) on stop trials, and both will typically be involved in stimulus-specific stop effects. The mechanism is straightforward for the direct pathway: It enables the cue to activate the representation that leads to stopping, which slows a go response on a go trial and helps avoid an erroneous action on a stop trial. The case for the indirect pathway can be equally straightforward if we simply assume that the cue activates the signal representation sufficiently to allow it to in turn then activate the stop centre. There is, however, another possibility inherent in this arrangement of links and representations, which is that the activation passed to the signal representation is not sufficient to result in any appreciable activation that can then be passed on to the stop centre. Instead, this input primes signal detection, allowing easier and more rapid detection of the stop signal when it occurs, as it already has some sub-threshold input applied to it. Whilst detection of the stop signal is essential to successfully stop, and thus its enhancement may be advantageous on signal trials, this scenario would have little behavioral consequence on trials where the stop signal does not occur. As reaction time measures are gathered on no-signal trials, we would not expect to observe much slowing if enhanced signal detection is what drives an associatively-mediated stop effect. This arrangement naturally leads to the prediction that the indirect pathway can lead to effects on p(respond|signal) in the absence of any effect on RT (Verbruggen, Best, et al., 2014). By contrast, the direct pathway is constrained to affect both p(respond|signal) and RT.

The indirect associative pathway is reliant on stable contingencies between cues and stop signals. Therefore, manipulating the contingencies between cues and stop signals can bias the relative strength of the direct and indirect associative pathways. This can be straightforwardly implemented by systematically varying the number of stop signals such that cues are either (A) presented with a single stop signal, or (B) presented with multiple stop signals that are equally distributed across all cues (see Table 1). Table 2 gives further insight into this manipulation. It gives the contingency (defined as (P(event | target cue)-P(event | no target cue)) $\times$ 100) relating the cue to either the signal(s) used or stopping. Inspection of the table reveals that in the single-signal case the contingencies for the signal and for stopping are obviously the same. The implication is that both associations will, other things being equal, be learned to a similar extent. We can use performance in

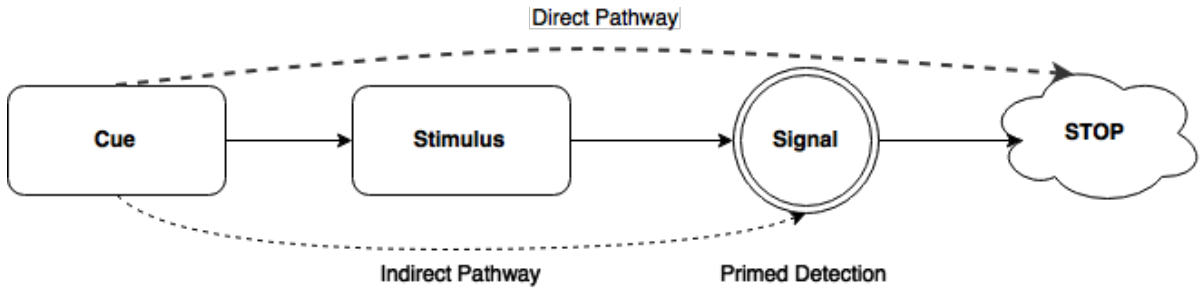

Figure 1. Diagrammatic representation of possible pathways to stopping.The direct pathway (bold dashed line) depicts associations from the cue to the stop centre that are not mediated via the stop signal representation. The indirect pathway (dashed line) depicts associations to the stop signal representation, which can then trigger activation in the stop centre via the link (solid arrow) that already exists.

this condition as a baseline for predicting what will happen in our other condition. In the multiple-signal case the pattern of contingencies changes - now the contingencies for stopping are substantially higher than those for the signal, favoring the formation of direct cue to stop associations, particularly as the contingent relationship to the signal is now so weak. Thus, we can argue that the shift to multiple signals should bring about a quite substantial shift in the relative strengths of the pathways involved in any associatively-mediated stopping, and this in turn should lead to stronger effects on RT in such a multiple-signal condition.

Our analysis thus far can only be part of the story, as close inspection of the table reveals that the contingencies for single- and multiple-signal conditions are the same for stopping, but different for the signal. Thus, one reading of the contingency table is that the strength of the associative link to the stop centre should be equal in both groups in our experiments, but there should be more priming of the stop signal in the single group. There are two possible mechanisms, however, which suggest that reducing the stimulus-signal contingency will result in stronger stimulus-stop learning. One relies quite straightforwardly on the fact that ultimately both pathways attempt to activate the stop center. If an error correcting algorithm is in force for associative learning, as we believe is the case (McLaren et al., 2014; Verbruggen, Best, et al., 2014), then the more effective one pathway is the less effective the other pathway will be - they will compete for the ability to activate the stop center. One way of viewing this is as an example of the overshadowing phenomenon often found in associative learning (Mackintosh 1976; and see McLaren et al. 2014). As has recently been noted, however, there is another mechanism that can bring about overshadowing that may be particularly applicable to our single vs. multiple manipulation (Civile, Chamizo, Mackintosh, & McLaren, 2014). This appeals to generalization decrement (Pearce, 1987) and simply points out that if two stimuli (in this case the serial compound of the cue and the stop signal) both predict an outcome (stopping), then when one (say the cue) is presented, the activation of that outcome representation suffers from generalization decrement (i.e. a reduction in that activation) due to the other stimulus not being presented. It is easy to see how this might apply to the single-signal case. But the multiple-signal case explicitly trains reliance on the cue rather than the stop signal. Here the network (see McLaren, Forrest, and McLaren 2012 for an example of such a network) will form multiple representations that capture each cue and signal configuration's link to stopping. When the cue is presented on its own, it will only partially activate all these representations, but the summed effect on stopping will be strong. Therefore, there will be less generalization decrement than in the case where a single cue + signal representation is involved. As a consequence, both mechanisms predict that less overshadowing will be observed, and consequently the associatively-mediated effect on stopping will be greater, in the multiple-signal case than in the single-signal case. Our experiments test this prediction.

Experiment 1

Method

Subjects. Forty-two students from the University of Exeter participated in return for £5 cash or 1 course credit. The majority of which were right handed (97%), females (71%), with an average age of 22 years and 7 months.

Apparatus and Stimuli. The experiment was run on an iMac computer (20" display; Apple, California) using Matlab 2012b in conjunction with the Psychophysics Toolbox 3 (Brainard, 1997). The stimuli consisted of three circles (19mm diameter) arranged in a horizontal line and presented centrally on a 50% grey background, and separated by 22mm edge-to-edge. At fixation, the middle circle appeared as a white outline, which on each trial filled with one of four colors (see Table 1 and Figure 2). Subsequently, one of the peripheral circles (left or right) filled with white and participants responded with a spatially congruent key ('X' or '.>', with their left or right index finger). However, on signal trials the

				Stop-signal Colour					
				Single		Multiple			
Cue	Colour	Signal Trials	No-signal Trials	E	F...G...H	E	F	G	H
A	75% Stop	24	8	24	0	6	6	6	6
B	25% Stop	8	24	8	0	2	2	2	2
C	50% Stop	16	16	16	0	4	4	4	4
D	50% Stop	16	16	16	0	4	4	4	4

Table 1

Depicts the design and cue/stop signal pairings employed in Experiments 1 and 2. ABCD represent central cue colors; either blue (RGB: 000 000 255), yellow (255 255 000), violet (128 000 128) or brown (128 051 000). EFGH represent stop signal colors; these were orange (255 128 000), pink (255 170 204), red-brown (168 046 037), or turquoise (000 172 165).

peripheral circle filled with one of four colors after a variable delay, prompting participants to withhold their response. Incorrect responses (or failures to respond) were signaled with a 400Hz 150-millisecond tone delivered through loudspeakers.

Procedure

Each trial began with the presentation of a cue, when the central circle filled with one of four colors (Table 1) for 250ms. Following the colored cue, which remained on screen for the duration of the trial, one of the peripheral circles filled white, instructing the participant to execute a left or right response. On no-signal trials, the go stimulus remained on screen for 1000ms during which period the participant could respond. However, on some trials, following a variable stop-signal delay (SSD), the circle temporarily changed to one of four colors (stop signal) for 250ms, instructing the participant to withhold their response. The next trial commenced after a variable inter trial interval (between 250-500ms; average 375ms), during which the fixation screen was displayed.

The onset of the stop signal was varied systematically based on each participants' performance: initially the SSD was set at 250ms from stimulus onset, but after two consecutive successful stop trials it was increased by 50ms and each failure to stop resulted in a 50ms decrease. The SSD could therefore vary between 50–950ms. The tracking procedure applied only to control trials (50% stop), but experimental trials were yoked to the same SSD. The 2-up/one-down procedure typically results in a 30% probability of successfully responding to a stop trial (p(respond|signal)) and compensates for both within and between-participant differences (Verbruggen & Logan, 2009). We used this tracking procedure to ensure that stopping was successful on most signal trials, as previous research suggests that the outcome can influence learning in stop-signal tasks (Verbruggen & Logan, 2008b, 2008a).

	Contingency(signal)x100		Contingency(stop)x100	
	Single	Multiple	Single	Multiple
75% Stop	33.3	8.33	33.3	33.3
50% Stop	0	0	0	0
25% Stop	-33.3	-8.33	-33.3	-33.3

Table 2

Contingencies between cues and stop-signals or stopping. Defined as P(event | target cue)-P(event | not target cue) × 100 and therefore can vary between +100 and -100. A zero contingency means there is no predictive relationship.

Each cue color was stochastically predictive of whether the trial would involve the execution or inhibition of a response: one cue color was mostly paired with stopping (75% stop), one responding (25% stop), and two with both (50% stop) (see Table 1). Thus, the overall number of signal and no-signal trials was equal. Consistent with our previous work (Yeates, Jones, Wills, McLaren, & McLaren, 2012; Yeates, Jones, Wills, Aitken, & McLaren, 2013), the predictive value of the cue was not explicitly revealed to participants and they were simply told: "The central colored circle acts as a warning that the trial is about to begin".

The color cues were either paired with a single stop signal or distributed evenly across multiple stop signals (a between groups manipulation). In both cases participants were given the same instructions: to "stop if the filled circle changes from white to any color". However, the single stop-signal group only saw one color, randomly selected from a pool of four, as a stop signal. In the multiple-signal group, each cue was paired with four different colored stop signals equally often (see Table 1).

Participants completed 10 training blocks of 128 trials, followed by two test blocks of the same length, where all cues were non-predictive (all contingencies 50:50). If subjects had acquired stimulus-stop associations during training, we would expect this to influence performance at test. Between each block participants were given a 30 second break (minimum) and given feedback if performance substantially differed from the previous block. Specifically, if participants' reaction times slowed by 5% and reaction times were >300ms they were instructed to respond more rapidly. Similarly, if errors increased by 5% and were in excess of 5%, they were instructed to re-

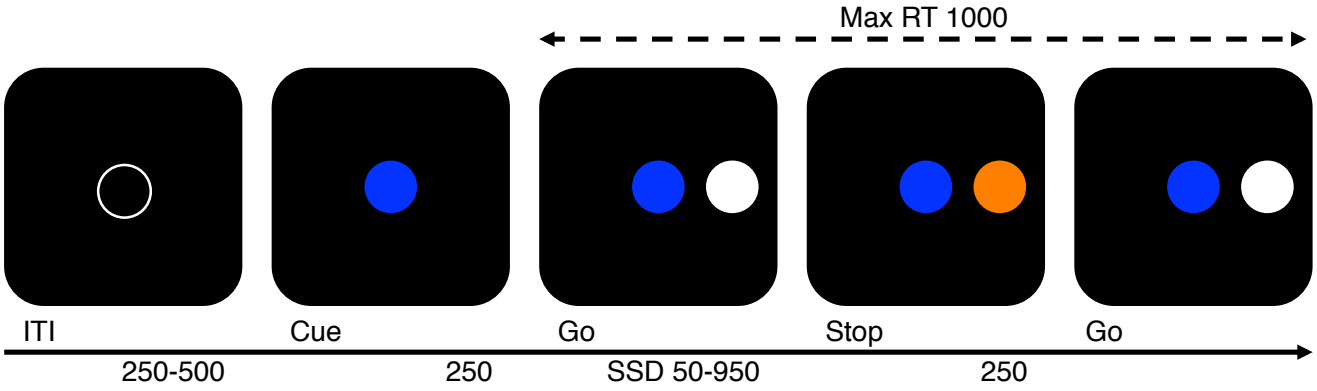

Figure 2. Depicts an example stop trial. All durations are in milliseconds. The central colored circle acts as the cue, the white circle to the left or right (right in this case) as the go stimulus, and if it changes color (orange in this case) this is the stop signal. A go trial progresses with the same time course but in the absence of the stop signal.

spond more accurately. Following the stop-signal task, participants were shown each central cue and asked to rate how much they would expect to respond or withhold responding on a scale of 1 (Not at all) to 9 (Definitely).

Analysis

All data was analysed using R (R Development Core Team, 2014), raw data and analysis scripts are available online (http://hdl.handle.net/10871/18105).

Two participants were excluded from the initial analysis: one for not stopping throughout the experiment and one due to technical difficulties that prevented them from completing the experiment.

Boxplot analysis identified five outliers: three had unusually low no-signal choice accuracy (< 75%) and two had unusually high p(respond | signal) (> 0.39), leaving seventeen participants in the single-signal group and eighteen participants in the multiple-signal group.

Results and discussion

Results are summarized in Table 3. We analyzed performance with a mixed ANOVA with trial type (75%, 50%, or 25% Stop) and block as within subject factors, and Group (multi vs. single) as a between subjects factor. Then we considered any of the interactions with Group that required further analysis.

Evidence of learning in measures of reaction time was observed across both training and test (see Figure 3, top panel). During training a main effect of trial type was observed (p<.01, $\hat{n}^2_G$ = .001); planned comparisons revealed that participants were slower to respond to trials cued by a 75% stop cue (M = 624, SD = 155) in comparison to those cued by a 50% (M = 617, SD = 154) (p<.01, $\hat{n}^2_G$ = .001) or 25% (M = 613, SD = 153) (p<.01, $\hat{n}^2_G$ = .001) stop cue. 50% stop cues and 25% stop cues did not significantly differ (p=.20, $\hat{n}^2_G$ = .000).

At test the effect of trial type was marginally significant (p<.06, $\hat{n}^2_G$ = .001); follow up comparisons revealed that 75% stop cues (M = 614, SD = 165) prompted significantly slower responses than 25% stop cues (M = 602, 163) (p<.04, $\hat{n}^2_G$ = .001), all other comparisons failed to reach significance (all p>.12, $\hat{n}^2_G \leq$ .001). The analysis revealed a three-way interaction between trial type, block, and group (multiple/single stop signals) during test (p<.04, $\hat{n}^2_G$ = .001), which was limited to the 75%/25% stop comparison (p<.03, $\hat{n}^2_G$ = .001). Whilst for the single-signal group, participants were initially slower to respond to 75% stop cued trials in comparison to 25% stop cued trials, the effect was markedly reduced by the second test block. Conversely, the multiple-signal group were slower to respond to 75% stop cued trials in comparison to 25% stop cued trials across both blocks, with the effect being somewhat larger in the second block of test. We investigated this interaction further by running separate contrasts for the 75% vs 25% comparison for each group in each test block. This revealed that there was a marginally significant effect for the first block of test in the single-signal group (p<.06, $\hat{n}^2_G$ = .007; but none in the second block of test, where numerically the effect reversed. The multiple-signal group exhibited the converse pattern, with no significant effect of 75% stop vs. 25% stop in the first block of test (the numerical effect was in the expected direction), but a marginally significant effect in the second (p<.08, $\hat{n}^2_G$ = .004). This pattern could suggest that there are roughly equivalent weak effects in both groups, and it was just chance that led to it manifesting in the first block for the single-signal group and the second block for the multiple-signal group. Alternatively, this result could suggest that the distributed signal training resulted in more robust learning, in the sense that the single-signal effects either diminished rapidly or were simply weaker and hence more

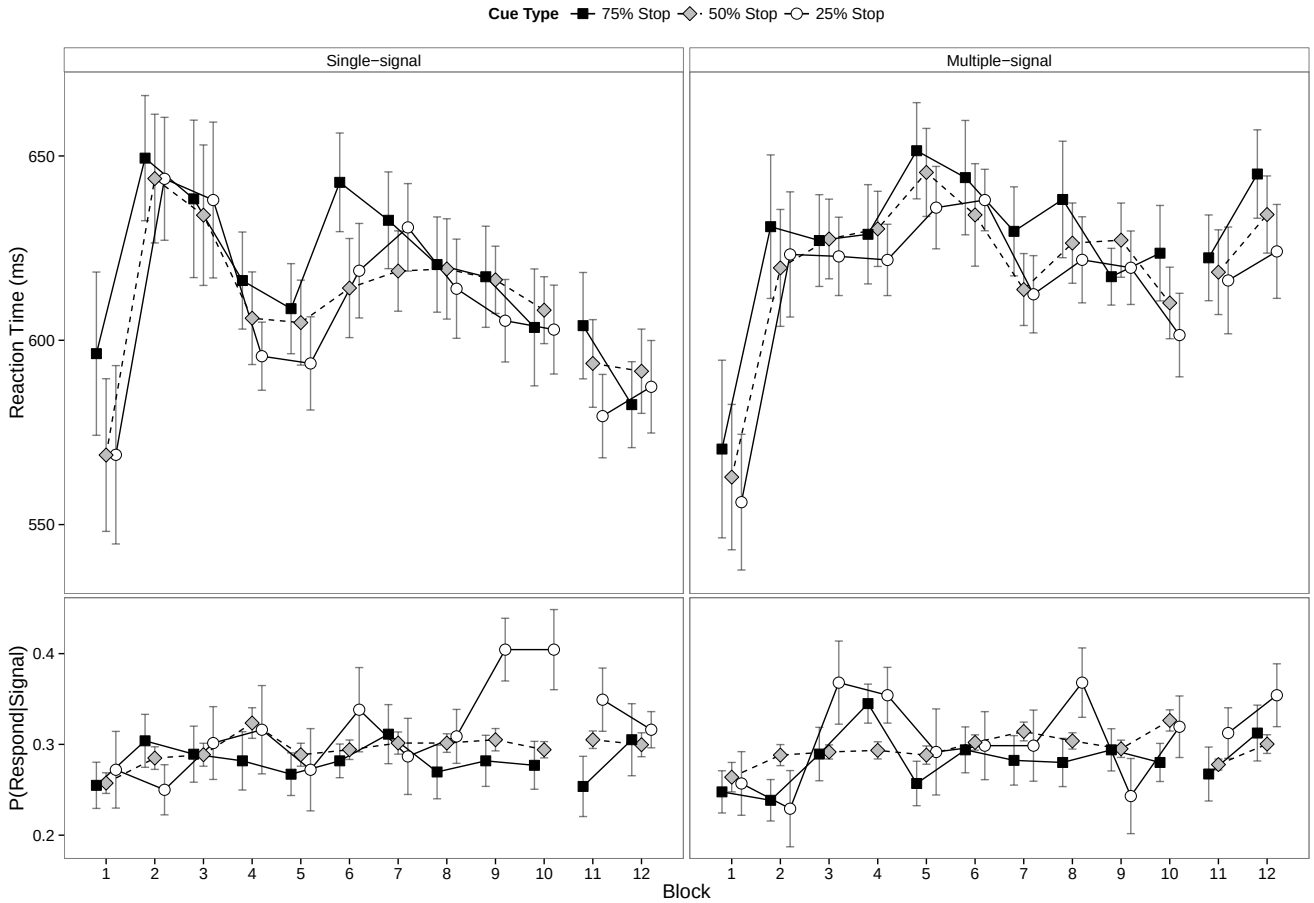

Figure 3. Mean reaction times of no-signal trial (top) and p(respond|signal) (bottom) for the single-signal (left) and multiple-signal groups (right) from Experiment 1. Error bars are normalised 95% confidence intervals (see Morey, 2008).

variable. We shall return to this point shortly.

In measures of *p*(respond|signal) we observed main effects of trial type in both training ($p<.01$, $\hat{n}^2_G = .009$) and test ($p<.02$, $\hat{n}^2_G = .035$) (see Figure 3 bottom panel) and there were no significant interactions with the group factor. Planned comparisons revealed that, during training, participants were less likely to make a commission error to trials cued by a 75% stop cue (M = 0.28, SD = 0.11) in comparison to a 25% stop cue (M = 0.31, SD = 0.17) ($p<.01$, $\hat{n}^2_G = .010$). 75% stop cues also significantly differed from controls (50% stop, M = 0.30, SD = 0.05) ($p<.03$, $\hat{n}^2_G = .007$). Similarly, during test, participants were less likely to make a commission error to trials preceded by a 75% stop cue (M = 0.28, SD = 0.14) in comparison to 25% stop cues (M = 0.33, SD = 0.13) ($p<.02$, $\hat{n}^2_G = .033$). However, during test only the 25% stop cues differed from 50% controls (M = 0.30, SD =0.04) ($p<.01$, $\hat{n}^2_G = .040$; 75% vs. 50% $p = .52$, $\hat{n}^2_G = .003$).

Overall, these results confirm that the contingencies were learned, as the 75% vs 25% difference is reliable, and involves slower responding on no-signal trials and fewer errors of commission on signal trials to the 75% stop cue relative to the 25% stop cue. The fact that there was no interaction in *p*(respond|signal) during test with the group factor was also expected and suggests that both single-signal and multiple-signal groups were equally able to benefit from the presence of a 75% stop cue in aiding them to withhold their response on a stop-signal trial. As we have indicated, the interaction with the group factor for RTs on no-signal trials during test could indicate that the effect on RTs was more robust in the multiple-signal group, but the involvement of block complicates its interpretation. Given the importance of this issue for our theoretical understanding of the basis of the associatively-mediated stopping effect we decided to replicate and extend Experiment 1 to clarify this result.

Experiment 2

The interaction observed in Experiment 1 is consistent with the idea that distributing multiple stop signals equally across cues influences the associatively mediated stopping

Table 3
Summary of Experiment 1

	DFn	DFd	SSn	SSd	F	p	p<.05	$\hat{n}^2_G$
Training								
Go Reaction Time								
cue type	2	66	22570.71	97132.60	7.67	.001	*	.001
75% stop vs. 25% stop	1	33	21446.61	59628.24	11.87	.002	*	.001
75% stop vs. 50% stop	1	33	10456.92	48818.39	7.07	.012	*	.001
25% stop vs. 50% stop	1	33	1952.54	37252.27	1.73	.198		.000
p(respond)								
cue type	2	66	0.13	0.86	5.10	.014	*	.009
75% stop vs. 25% stop	1	33	0.13	0.61	7.21	.011	*	.010
75% stop vs. 50% stop	1	33	0.03	0.22	5.15	.030	*	.007
25% stop vs. 50% stop	1	33	0.03	0.46	2.29	.139		.003
Test								
Go Reaction Time								
cue type	2	66	4979.56	54805.72	3.00	.058		.001
75% stop vs. 25% stop	1	33	4834.75	34151.77	4.67	.038	*	.001
75% stop vs. 50% stop	1	33	592.67	20961.43	0.93	.341		.000
25% stop vs. 50% stop	1	33	2041.92	27095.37	2.49	.124		.001
cue type:block:multiple/single	2	66	4333.73	42132.54	3.39	.043	*	.001
75% stop vs. 25% stop	1	33	4280.02	26233.07	5.38	.027	*	.001
Single Signal block 11	1	16	5117.246	19171.19	4.27	.055	^	.007
Single Signal block 12	1	16	200.51	11685.08	0.27	.607		.000
Multiple Signal Block 11	1	17	338.32	10236.30	0.56	.464		.000
Multiple Signal Block 12	1	17	3964.97	19292.27	3.49	.079	^	.004
75% stop vs. 50% stop	1	33	1525.48	23135.00	2.18	.150		.000
25% stop vs. 50% stop	1	33	695.08	13830.75	1.66	.207		.000
p(respond)								
cue type	2	66	0.09	0.69	4.29	.022	*	.035
75% stop vs. 25% stop	1	33	0.08	0.48	5.57	.024	*	.033
75% stop vs. 50% stop	1	33	0.00	0.32	0.42	.521		.003
25% stop vs. 50% stop	1	33	0.05	0.24	6.82	.013	*	.040

effect, presumably as the distribution of signals reduces the formation of cue-signal associations and therefore increases the relative strength of cue-stop associations. Experiment 2 sought to replicate this effect, using the same procedures as Experiment 1 but this time run in a group testing facility allowing us to test more participants.

Method

Subjects. Sixty-six students from the University of Exeter participated in return for 1 course credit or £5. The majority of which were right handed (89.6%) females (62.7%), with an average age of 20 years.

Apparatus & Stimuli. The stimuli were identical to those used in Experiment 1. However, Experiment 2 was run on PCs, with 19" monitors, in a multiple testing environment. Consequentially, error tones were presented through closed headphones, rather than loud speakers.

Procedure

The procedure was identical to that of Experiment 2. Participants were assigned to each stop-signal group serially, unless they were replacing an identified outlier.

Analysis

Two participants were excluded for using the incorrect response keys. A further four participants were removed for having unusually low no-signal choice accuracy (<75%) and two for having unusually high p(respond|signal) (>41%) as identified by box-and-whisker analysis. This left 60 participants in total, with 30 in each stop-signal group.

Results and discussion

The results of this analysis are summarized in Table 4 and Figure 4. Replicating Experiment 1, we observed a main effect of cue type in reaction times during training (p<.04,

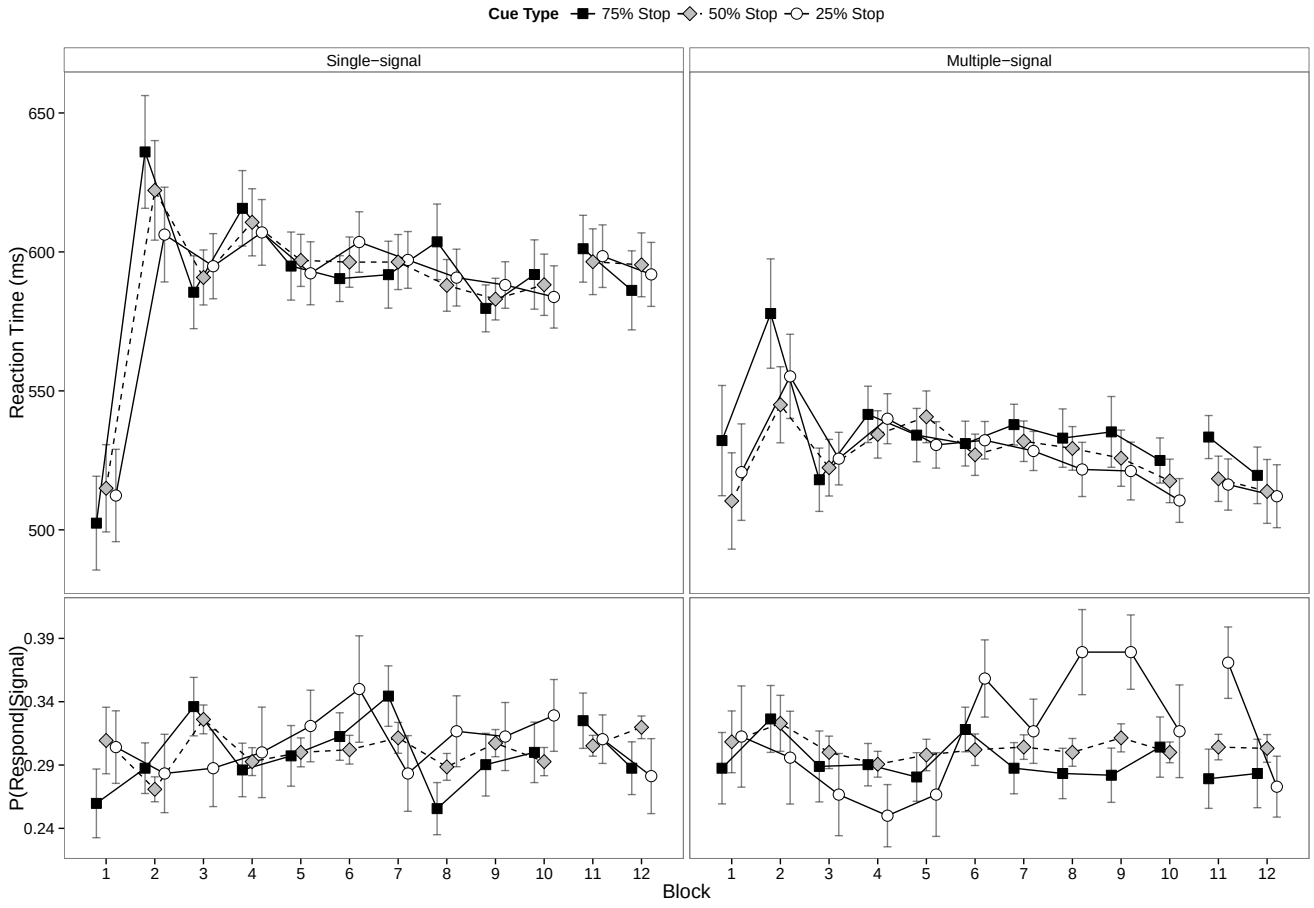

Figure 4. Mean reaction times of no-signal trials (top) and p(respond|signal) (bottom) for the single-signal (left) and multiple-signal groups (right) from Experiment 2. Error bars are normalised 95% confidence intervals (see Morey, 2008).

$\hat{n}^2_G$ = .000). Planned comparisons confirmed that 75% stop cues (M = 563, SD = 111) were significantly different to both 25% (M = 558, SD = 109) (p<.05, $\hat{n}^2_G$ = .000) and 50% (M = 559, SD = 113) (p<.02, $\hat{n}^2_G$ = .000) stop cues. However, 25% and 50% stop cues did not significantly differ (p=.81, $\hat{n}^2_G$ = .000). At test, whilst the main effect of cue type was not significant (p = .19, $\hat{n}^2_G$ = .000), a two way interaction between cue type and group was observed in measures of reaction time (p<.05, $\hat{n}^2_G$ = .001). Follow up comparisons found the interaction to be limited to the 75% vs. 25% stop cue (p<.05, $\hat{n}^2_G$ = .001) and 75% vs. 50% stop cue (p<.02, $\hat{n}^2_G$ = .001) comparisons. This suggests that the manipulation selectively influenced the 75% stop cues and not the 25% stop cues. The interaction reflects greater learning in the multiple-signal group, where the overall difference between 25% (M = 514, SD = 28) and 75% (M = 526, SD = 30) stop cues was 12ms (p<.02, $\hat{n}^2_G$ = .003), in comparison to the single-signal group where the difference was just 1ms (75%: M = 594, SD = 33; 25%: M = 595, SD = 32, p = .73, $\hat{n}^2_G$ = .000). Similarly, a significant difference was observed between 75% and 50% stop cues (M = 516, SD = 26) in the multiple-signal group (p<.01, $\hat{n}^2_G$ = .002), but not in the single-signal group (M = 596, SD = 28, p = .53, $\hat{n}^2_G$ = .000). No significant differences were observed between 25% and 50% stop cues (all p>.65). Thus, we now have clear evidence that the multiple-signal group shows a stronger stopping effect on the RT measure than the single-signal group. We should point out that the single-signal group were markedly slower in both training (p<.04, $\hat{n}^2_G$ = .051) and test (p<.02, $\hat{n}^2_G$ = .085) (M = 589, SD = 123; M = 595, SD = 140, respectively) in comparison to the multiple-signal group (M = 531, SD = 90; M = 519, SD = 104, respectively). We have no ready explanation for this effect, given that it did not occur in Experiment 1. Although this slowing could be interpreted as problematic for between-group comparisons, as slowing for the single-signal group may have obscured cue-specific slowing, we note that reaction times of this group are comparable to Experiment 1 (where a main effect of cue type was observed).

In the p(respond|signal) measure of performance, a significant cue type by block interaction was observed during

Table 4
Summary of Experiment 2

	DFn	DFd	SSn	SSd	F	p	p<.05	$\hat{n}^2_G$
Training								
Go Reaction Time								
cue type	2	116	8262.61	145565.40	3.29	.044	*	.000
75% stop vs. 25% stop	1	58	6799.12	96757.20	4.08	.048	*	.000
75% stop vs. 50% stop	1	58	5529.22	59973.96	5.35	.024	*	.000
25% stop vs. 50% stop	1	58	65.58	61616.98	0.06	.805		.000
multiple/single	1	58	1477576.00	20187171.00	4.25	.044	*	.051
p(respond)								
cue type	2	116	0.07	2.23	1.92	.159		.002
block	9	522	0.19	14.40	0.78	.623		.006
cue type:block	18	1044	0.47	12.78	2.13	.007	*	.014
first half	2	116	0.03	1.58	0.96	.363		.001
second half	2	116	0.24	1.96	7.02	.003	*	.016
75% stop vs. 25% stop	1	58	0.20	1.25	9.25	.004	*	.014
75% stop vs. 50% stop	1	58	0.00	0.52	0.29	.589		.001
25% stop vs. 50% stop	1	58	0.16	1.18	7.64	.008	*	.016
Test								
Go Reaction Time								
cue type	2	116	1890.22	65004.26	1.69	.191		.000
multiple/single	1	58	519677.60	5316395.87	5.67	.021	*	.085
cue type:multiple/single	2	116	3532.42	65004.26	3.15	.048	*	.001
75% stop vs. 25% stop	1	58	2872.02	39775.74	4.19	.045	*	.001
single	1	29	70.90	17576.00	0.12	.734		.000
multiple	1	29	4539.00	22200.00	5.93	.021	*	.003
75% stop vs. 50% stop	1	58	2405.99	25064.57	5.57	.022	*	.001
single	1	29	153.00	11198.00	0.40	.533		.000
multiple	1	29	3247.00	13866.00	6.79	.014	*	.002
25% stop vs. 50% stop	1	58	20.61	32666.09	0.04	.849		.000
single	1	29	15.70	17030.00	0.03	.871		.000
multiple	1	29	107.93	15637.00	0.20	.658		.000
p(respond)								
cue type	2	116	0.02	1.44	0.70	.496		.004
block	1	58	0.05	0.84	3.72	.059	^	.011
cue type:block	2	116	0.08	1.25	3.58	.031	*	.016
first block	2	116	0.06	1.30	2.49	.088	^	.025
75% stop vs. 25% stop	1	58	0.04	0.78	3.33	.073	^	.021
75% stop vs. 50% stop	1	58	0.00	0.56	0.02	.885		.001
25% stop vs. 50% stop	1	58	0.04	0.61	3.67	.060	^	.032
second block	2	116	0.04	1.40	1.60	.206		.015

both training (p<.01, $\hat{n}^2_G$ = .014) and test (p<.03, $\hat{n}^2_G$ = .016). Follow up comparisons revealed that differences in cue type were contingent on the amount of training; whilst the first half of training displayed no significant effect of cue type (p=.36, $\hat{n}^2_G$ = .001) the second half of training did (p<.01, $\hat{n}^2_G$ = .016). 75% stop cues (M = 0.30, SD = 0.06) resulted in significantly fewer errors than 25% stop cues (M = 0.33, SD = 0.09) (p<.01, $\hat{n}^2_G$ = .014), but did not differ from 50% cues (M = 0.30, SD = 0.01) (p=.59, $\hat{n}^2_G$ = .001). Additionally, 25% and 50% stop cues differed significantly (p<.01, $\hat{n}^2_G$ = .016). Conversely, at test, the effect of cue type was marginally significant during the first block (p = .09, $\hat{n}^2_G$ = .025), but had extinguished by the second block (p = .20, $\hat{n}^2_G$ = .015). Follow up comparisons, performed on the first block of test, revealed that 25% stop cues (M = 0.34, SD = 0.14) marginally differed from both 75% (M = 0.30, SD = 0.13) (p<.07, $\hat{n}^2_G$ = .021)

and 50% cues (M = 0.30, SD = 0.04) ($p<.06$, $\hat{n}^2_G$ = .032). However, 75% and 50% stop cues did not differ (p = .89, $\hat{n}^2_G$ = .001).

These results help us interpret the findings of Experiment 1: We can now be sure that the multiple-signal training regime results in more robust slowing to 75% stop cues, in measures of reaction time, than the single-signal variant. Similarly, in measures of p(respond|signal) a main effect of trial type was observed, albeit limited to the second half of training, and a marginal effect on test.

General Discussion

Experiments 1 and 2 revealed that manipulating the pairings between cues and stop signals, in a manner that reduces their contingent relationship, results in more robust cue-specific stop effects on reaction times, but does not affect p(respond|signal) (i.e. measures of p(respond|signal) do not interact with group). Across both experiments the multiple-signal groups show effects on both measures during test, but the single-signal groups do not always do so. The single-signal groups produce reliable effects on p(respond|signal) in both experiments, in the sense that there is evidence of an effect of cue type and no significant interaction with Group, but the evidence for any effect on RT is mixed. There is a reliable effect in Experiment 1, but none in Experiment 2, and the effect in the multiple-signal group in this last experiment is significantly different to that in the single-signal group. Given the results of Experiment 1, this is perhaps the sort of pattern that should be expected for this group on the RT measure, and suggests a relatively weak effect of the indirect associative pathway on RTs.

Associative learning of stop signals

We have proposed that the condition employing a single stop signal emphasizes an indirect link from cue to stopping via the signal representation, whereas the condition employing multiple stop signals shifts the emphasis to a direct association from cue to stop centre (see Best, Lawrence, Logan, McLaren, & Verbruggen, 2015).

The careful reader might wonder why the stop signal itself, which is a 100% valid cue for stopping, does not always (eventually) overshadow the cue (which at most is 75% valid in our experiments). As a corollary, surely the signal will become the stimulus most strongly associated with stopping and we should be using this as a cue in some test phase where we change the signals used to denote stopping? However, there is a theoretical reason to doubt this logic; the signal's timing in relation to stopping is not ideal for associative learning (the interval between signal and response is too short, See Mackintosh 1974, p. 57), whereas that of the cue is (and quite deliberately so). It may be that this allows the cue equal status with the signal in forming a serial compound that becomes associated with stopping, and that this then leads to our current pattern of results. On the other hand, it may be that the signal is entirely ineffective in associating to the stop outcome, and that the only associations in play are those involving the cue. Supporting this view, research assessing how the relative speed of the stop process (as indexed by SSRT) changes with practice is mixed and does not always yield any significant improvement (J. R. Cohen & Poldrack, 2008; Logan & Burkell, 1986). Thus, suggesting that response inhibition may not benefit from acquired associations between stop signals and stopping. If the latter is the case, then we must appeal to the competitive version of overshadowing alluded to earlier, but if the former is true, then the generalization decrement version of the overshadowing account is also viable. Future research that explicitly compares the effectiveness of the cue and the stop signal in producing associatively-mediated slowing of go responses after stop training would help us to decide between these alternatives.

We do not, at present, have the direct evidence for priming of the stop signal representation that would substantiate our analysis of the single signal group's performance as being due to the indirect pathway that we have identified. Our evidence is indirect, and inferred from the fact that the effect on RTs in the single signal groups is rather weak compared to that on p(respond|signal) (see also Verbruggen, Best, et al., 2014, for a similar pattern of results). One way in which we could attempt to rectify this in the future would be to train cues using a single signal, and then instruct our participants that the stop signal was no longer effective. If the indirect pathway was in play, and was sufficiently strong to generate an effect on RTs, then this instruction should immediately abolish any such effect. Our data suggest that these provisos will be quite difficult to meet, however, and we would first need to find a method of strengthening this indirect influence to the point where it would affect RTs reliably before making any such attempt. As matters stand, in our experiments there would have been little effect on RTs to influence using this manipulation. The corollary, however, that the training based on use of multiple stop signals should be unaffected by this type of instructional manipulation is easily done and something that we intend to look at in the future. Additionally, we could utilize neuroimaging techniques with fine temporal resolution (such as EEG) to establish whether training influences perceptual or response related processes.

The role for signal detection that we have identified provides us with an alternative to the view that the rIFG is exclusively responsible for automatic 'inhibition'. Lenartowicz and colleagues (2011) found increased activation on no-signal trials for stimuli that had previously been associated with stopping. They argued that this reflected automatic activation of the stop response. However, as we have demonstrated that signal detection can also become learned and subregions of rIFG have been implicated in stimulus detection processes (Dodds, Morein-Zamir, & Robbins, 2011; Hampshire, Chamberlain,

Monti, Duncan, & Owen, 2010), it is possible that increased rIFG activation on no-signal trials represents priming of the stop signal (even though it is not actually presented).

A final theoretical issue is whether we are justified in classifying our results for the 75% cues as denoting associatively mediated inhibition. Another way of interpreting our results would be to say that the 25% cue becomes an excitatory or "go" stimulus, promoting more rapid responding and leading to more errors on stop trials. To explore this idea, we can use the 50% cues as a baseline because these cues are neither associated with going nor with stopping. Thus, the differences between 25% and 50% cue types in Experiments 1 and 2 indicates that the 25% cue becomes an excitatory or "go" stimulus. Importantly, the differences between the 50% baseline and the 75% cue also makes the case for the 75% cue having an inhibitory influence on responding. Thus, we propose that both excitatory and inhibitory effects occur (for a more elaborate discussion of this issue, see Best et al., 2015).

Implications for stop training programmes

The implications of this research for inhibition training are clear; if transfer effects are due to associative learning, then the introduction of multiple stop signals should result in greater stimulus-stop learning and potentially enhance the effectiveness of this type of training. Certainly our evidence suggests that the cue-specific effect is more robust in the multiple-signal group, and we can see no reason why this would not be expected to apply in inhibition training with stimuli such as foods (e.g. Lawrence, Verbruggen, Morrison, Adams, & Chambers, 2015; Veling, Aarts, & Papies, 2011) or alcohol (e.g. Jones & Field, 2013) that use a stop-signal paradigm similar to ours. It remains to be seen whether this approach is to be preferred to the use of a Go/NoGo paradigm for training purposes. The latter has the advantage that the cues used to signal a NoGo trial are 100% reliable as they do not suffer from the failure rate inherent in the tracking procedure used in stop-signal tasks. Whilst the feature used to signal a NoGo trial will suffer from a poor temporal relationship to stopping in terms of generating any associative learning, this can be solved by presenting the target associative cue before the signal in this paradigm as in Best et al. (2015) or Veling et al. (2011). Indeed, a recent meta-analysis suggests that the Go/NoGo paradigm results in a greater reduction of alcohol or food consumption than the stop-signal paradigm (Jones et al., 2015), yet none of these tasks used multiple stop signals. There is, however, at least one reason to think that the stop-signal paradigm will produce more potent associative effects. In inhibition training with animals, one of two standard procedures is often used. A conditioned and unconditioned stimulus can be explicitly unpaired, thus A+ B- will give B some inhibitory properties, or a conditioned inhibition procedure can be used of the form A+ AB-. This latter has been shown to result in stronger inhibitory responding to B (see McLaren & Verbruggen, 2015). One theoretical analysis of this result is simply to say that having A predict the outcome unless it is paired with B generates a larger prediction error, and hence stronger learning, than a design that effectively relies on the context to do this (as in A+ B-). Clearly the A+ AB- version is more like the stop-signal procedure and thus, by analogy, may be expected to produce stronger associatively-mediated stopping. It may be that a feature negative design (Jenkins & Sainsbury, 1969), combining the best aspects of both Go/NoGo and stop-signal methodologies will eventually prove the most effective (and we have some preliminary data that suggests this might be the case). This hybrid approach would effectively use a fixed SSD of zero, such that the signal/feature would appear at the same time as what would otherwise be the Go stimulus, once again producing a large prediction error to drive learning.

Conclusion

The present study sought to demonstrate that both signal detection and automatic response suppression are implicated in cue specific stop-signal tasks. The results suggest that arranging cues and stop-signals in a manner that reduces the contingency between them results in more robust slowing of reaction times on go trials preceded by a stop cue. We suggest that this enhancement arises because this configuration encourages the formation of direct stimulus-stop associations, rather than the formation of stimulus-signal associations that would primarily prime the detection of the stop signal. This finding has particularly interesting implications for applied settings. Firstly, it suggests that, if transfer effects (such as reduced food consumption) are the result of an acquired association with a stop response, multiple stop-signals should be employed to maximize stimulus-stop learning. Secondly, it suggests that stop-signal detection could also be enhanced through training if a single stop signal is employed. This may be of particular use in training subjects to more readily notice cues that prepare them to inhibit a response; e.g. a driver looking out for a red traffic light. We hope to continue our investigation of different learning/training designs so as to shed further light on these possibilities and help develop optimal techniques for inhibition training.

References

Aron, A. R., Robbins, T. W., & Poldrack, R. A. (2004). Inhibition and the right inferior frontal cortex. *Trends in cognitive sciences*, *8*(4), 170–7. doi:10.1016/j.tics.2004.02.010

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*(4), 177–185. doi:10.1016/j.tics.2013.12.003

Best, M., Lawrence, N. S., Logan, G. D., McLaren, I. P. L., & Verbruggen, F. (2015). Should I stop or should I go? The role of associations and expectancies. *Journal of Experimental Psychology: Human Perception and Performance, in press*.

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*(2), 376–97. doi:10.1037/0033-295X.114.2.376

Brainard, D. H. (1997). The Psychophysics Toolbox. *Spatial vision*, *10*(4), 433–6.

Brazzelli, M. & Spinnler, H. (1998). An example of lack of frontal inhibition: the 'utilization behaviour'. *European Journal of Neurology*, *5*(4), 347–353. doi:10.1046/j.1468-1331.1998.540347.x

Chambers, C. D., Garavan, H., & Bellgrove, M. A. (2009). Insights into the neural basis of response inhibition from cognitive and clinical neuroscience. *Neuroscience and biobehavioral reviews*, *33*(5), 631–46. doi:10.1016/j.neubiorev.2008.08.016

Chun, M. M. & Jiang, Y. (1999). Top-Down Attentional Guidance Based on Implicit Learning of Visual Covariation. *Psychological Science*, *10*(4), 360–365. doi:10.1111/1467-9280.00168

Civile, C., Chamizo, V., Mackintosh, N. J., & McLaren, I. P. L. (2014). The effect of disrupting configural information on rats performance in the Morris water maze. *Learning and Motivation*, *48*, 55–66. doi:10.1016/j.lmot.2014.07.001

Cohen, J. R. & Poldrack, R. A. (2008). Automaticity in motor sequence learning does not impair response inhibition. *Psychonomic Bulletin & Review*, *15*(1), 108–115. doi:10.3758/PBR.15.1.108

Dodds, C. M., Morein-Zamir, S., & Robbins, T. W. (2011). Dissociating inhibition, attention, and response control in the frontoparietal network using functional magnetic resonance imaging. *Cerebral Cortex*, *21*(5), 1155–1165. doi:10.1093/cercor/bhq187

Donders, F. (1969). On the speed of mental processes. *Acta Psychologica*, *30*, 412–431. doi:10.1016/0001-6918(69)90065-1

Hampshire, A., Chamberlain, S. R., Monti, M. M., Duncan, J., & Owen, A. M. (2010). The role of the right inferior frontal gyrus: inhibition and attentional control. *NeuroImage*, *50*(3), 1313–1319. doi:10.1016/j.neuroimage.2009.12.109

Jenkins, H. M. & Sainsbury, R. S. (1969). The development of stimulus control through differential reinforcement. *Fundamental issues in associative learning*, 123–161.

Jones, A., Di Lemma, L. C. G., Robinson, E., Christiansen, P., Nolan, S., Tudur-Smith, C., & Field, M. (2015). Inhibitory control training for appetitive behaviour change: A meta- analysis. Manuscript Submitted for Publication.

Jones, A. & Field, M. (2013). The effects of cue-specific inhibition training on alcohol consumption in heavy social drinkers. *Experimental and clinical psychopharmacology*, *21*(1), 8–16. doi:10.1037/a0030683

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. doi:10.1016/j.appet.2014.11.006

Lenartowicz, A., Verbruggen, F., Logan, G. D., & Poldrack, R. A. (2011). Inhibition-related activation in the right inferior frontal gyrus in the absence of inhibitory cues. *Journal of cognitive neuroscience*, *23*(11), 3388–99. doi:10.1162/jocn_a_00031

Lhermitte, F. (1983). 'Utilization behaviour' and its relation to lesions of the frontal lobes. *Brain : a journal of neurology*, *106*(2), 237–255. doi:10.1093/brain/106.2.237

Logan, G. D. (1988). Toward an instance theory of automatization. *Psychological Review*, *95*(4), 492–527. doi:10.1037/0033-295X.95.4.492

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*(4), 549–563. doi:10.1037//0096-1523.12.4.549

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*(1), 66–95. doi:10.1037/a0035230

Mackintosh, N. J. (1974). *The Psychology of Animal Learning*. London: Academic Press.

Mackintosh, N. J. (1976). Overshadowing and stimulus intensity. *Animal learning & behavior*, *4*, 186–192. doi:10.3758/BF03214033

MacLeod, C. M., Dodd, M. D., Sheard, E. D., Wilson, D. E., & Bibi, U. (2003). In Opposition to Inhibition. In B. H. Ross (Ed.), *Psychology of learning and motivation - advances in research and theory* (Vol. 43, pp. 163–214). London: Academic Press. doi:10.1016/S0079-7421(03)01014-4

McLaren, I. P. L., Forrest, C. L. D., McLaren, R. P., Jones, F. W., Aitken, M., & Mackintosh, N. J. (2014). Associations and propositions: The case for a dual-process account of learning in humans. *Neurobiology of Learning and Memory*, *108*, 185–195. doi:10.1016/j.nlm.2013.09.014

McLaren, I. P. L., Forrest, C. L., & McLaren, R. P. (2012). Elemental representation and configural mappings: Combining elemental and configural theories of associative

learning. *Learning & Behavior*, *40*, 320–333. doi:10.3758/s13420-012-0079-1

McLaren, I. P. L. & Verbruggen, F. (2015). Association and inhibition. In R. A. Murphy & R. C. Honey (Eds.), *The wiley blackwell handbook on the cognitive neuroscience of learning.* Chichester, England: John Wiley and sons.

Miller, E. K. & Cohen, J. D. (2001). An integrative theory of prefrontal cortex function. *Annual review of neuroscience*, *24*, 167–202. doi:10.1146/annurev.neuro.24.1.167

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*(1), 49–100. doi:10.1006/cogp.1999.0734

Morey, R. D. (2008). Confidence Intervals from Normalized Data: A correction to Cousineau (2005). *Tutorials in Quantitative Methods for Psychology*, *4*(2), 61–64. doi:10.3758/s13414-012-0291-2

Nambu, A., Tokuno, H., & Takada, M. (2002). Functional significance of the cortico-subthalamo-pallidal 'hyperdirect' pathway. *Neuroscience Research*, *43*(2), 111–117. doi:10.1016/S0168-0102(02)00027-5

O'Reilly, R. C. (2006). Biologically based computational models of high-level cognition. *Science (New York, N.Y.) 314*(2006), 91–94. doi:10.1126/science.1127242

Pearce, J. M. (1987). A model for stimulus generalization in Pavlovian conditioning. *Psychological review*, *94*(1), 61–73. doi:10.1037/0033-295X.94.1.61

R Development Core Team. (2014). R: A Language and Environment for Statistical Computing. Vienna.

Ridderinkhof, K. R., van den Wildenberg, W. P. M., Segalowitz, S. J., & Carter, C. S. (2004). Neurocognitive mechanisms of cognitive control: the role of prefrontal cortex in action selection, response inhibition, performance monitoring, and reward-based learning. *Brain and cognition*, *56*(2), 129–40. doi:10.1016/j.bandc.2004.09.016

Schneider, W. & Shiffrin, R. M. (1977). Controlled and automatic human information processing: I. Detection, search, and attention. *Psychological Review*, *84*(1), 1–66. doi:10.1037/0033-295X.84.1.1

Shiffrin, R. M. & Schneider, W. (1977). Controlled and automatic human information processing: II. Perceptual learning, automatic attending and a general theory. *Psychological Review*, *84*(2), 127–190. doi:10.1037/0033-295X.84.2.127

Veling, H., Aarts, H., & Papies, E. K. (2011). Using stop signals to inhibit chronic dieters' responses toward palatable foods. *Behaviour research and therapy*, *49*(11), 771–80. doi:10.1016/j.brat.2011.08.005

Verbruggen, F., Best, M., Bowditch, W. A., Stevens, T., & McLaren, I. P. L. (2014). The inhibitory control reflex. *Neuropsychologia*, *65*, 263–78. doi:10.1016/j.neuropsychologia.2014.08.014

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*(4), 649–72. doi:10.1037/a0013170

Verbruggen, F. & Logan, G. D. (2008b). Long-term aftereffects of response inhibition: memory retrieval, task goals, and cognitive control. *Journal of experimental psychology. Human perception and performance*, *34*(5), 1229–1235. doi:10.1037/0096-1523.34.5.1229

Verbruggen, F. & Logan, G. D. (2008c). Response inhibition in the stop-signal paradigm. *Trends in cognitive sciences*, *12*(11), 418–24. doi:10.1016/j.tics.2008.07.005

Verbruggen, F. & Logan, G. D. (2009). Models of response inhibition in the stop-signal and stop-change paradigms. *Neuroscience and biobehavioral reviews*, *33*(5), 647–61. doi:10.1016/j.neubiorev.2008.08.014

Verbruggen, F., McLaren, I. P. L., & Chambers, C. D. (2014). Banishing the control homunculi in studies of action control and behaviour change. *Perspectives on Psychological Science.* doi:10.1177/1745691614526414

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*(4), 1295–1300. doi:10.1037/a0036542

Yeates, F., Jones, F. W., Wills, A., Aitken, M., & McLaren, I. P. L. (2013). Implicit Learning : A Demonstration and a Revision of a Novel SRT Paradigm. *Proceedings of the 35th Annual Conference of the Cognitive Science Society*.

Yeates, F., Jones, F. W., Wills, A., McLaren, R. P., & McLaren, I. P. L. (2012). Implicit learning: A demonstration and a novel SRT paradigm. In N. Miyake, D. Peebles, & R. P. Cooper (Eds.), *Proceedings of the 34th annual meeting of the cognitive science society* (pp. 1185–90). Austin.