Behavioral and Neural Evidence for Item-specific Performance Monitoring

Chris Blais and Silvia Bunge

Abstract

■ How cognitive control is recruited and implemented has become a major focus of researchers in cognitive psychology and neuroscience. Current theories posit that cognitive control operates at the level of general rules—for example, in a Stroop task, "attend to the color of the stimulus." Here we report behavioral evidence suggesting that cognitive control is implemented much more locally, operating at the level of specific stimuli appearing in a task block. In addition, we report neural evidence that many of the regions implicated in cognitive control on the Stroop task, including anterior cingulate cortex and dorsolateral prefrontal cortex, operate at a local level. ■

INTRODUCTION

Cognitive control refers to the set of processes which allow us to flexibly guide our behavior. Although research on the components of cognitive control, such as working memory, response inhibition, and response selection, has long been prevalent, the past decade has seen a substantial increase into the investigation of how cognitive control is recruited. According to the widely accepted conflict monitoring hypothesis (Botvinick, Braver, et al., 2004; Botvinick, Cohen, & Carter, 2004; Botvinick, Braver, Barch, Carter, & Cohen, 2001), anterior cingulate cortex (ACC) detects conflict between active mental representations and recruits dorsolateral prefrontal cortex (DLPFC) and other brain regions to reduce that conflict by activating goal-relevant representations. This hypothesis is built on a solid base of behavioral, neural, and computational data; nonetheless, a number of open questions remain. One of the more recent issues to emerge regards the grain at which this ACC–DLPFC control loop operates (see Verguts & Notebaert, 2008; Blais, Robidoux, Risko, & Besner, 2007).

Central to this discussion has been the *proportion-congruent effect* in the Stroop task. In common variants of this task, participants must indicate the color that a word is printed in, while ignoring the meaning of the word. On incongruent trials, there is a mismatch between the color of the stimulus and the color word, such as the word "blue" printed in green ink. Such stimuli require participants to select between competing responses, unlike congruent stimuli like the word "blue" printed in blue ink. The Stroop effect measures the performance decrement (i.e., slower response times and/or lower accuracy) observed on incongruent trials relative to congruent trials.

As the proportion of congruent trials in a series of Stroop trials increases, the magnitude of the Stroop effect for incongruent trials increases (Jacoby, Lindsay, & Hessels, 2003; Lindsay & Jacoby, 1994; Tzelgov, Henik, & Berger, 1992; Cheesman & Merikle, 1986; Logan & Zbrodoff, 1979). In essence, we are lulled into complacency when we encounter a block of mostly congruent trials, and thus, deal less effectively with incongruent stimuli than when we are alert. Botvinick et al. (2001) explained this effect in the context of their conflict monitoring model by hypothesizing that the ACC–DLPFC network is sensitive to the *global* amount of conflict encountered over an extended period. That is, in a context involving a high degree of conflict overall (i.e., a low proportion of congruent trials), more control is recruited on a given incongruent trial compared to a context with a low degree of conflict (i.e., a high proportion of congruent trials). Thus, according to this model, the overall amount of conflict determines the magnitude of the Stroop effect. Botvinick et al. demonstrated the plausibility of this approach in their conflict monitoring model of the Stroop effect, and Carter et al. (2000) provided neural evidence supporting this account.

A different theoretical account of the proportion-congruent effect places control at a more *local* level, whereby the amount of control elicited in response to an item is based on the conflict associated with that item. Behavioral evidence for item-based control has emerged from experiments demonstrating that the proportion-congruent effect occurs even when specific stimuli, rather than blocks of stimuli, are associated with low or high conflict (e.g., the stimulus "blue" appears in blue ink 80% of the time or 20% of the time; Jacoby et al., 2003;

University of California, Berkeley

© 2010 Massachusetts Institute of Technology

Downloaded from ... on 18 May 2021

see also Notebaert & Verguts, 2008). Blais et al. (2007) adapted Botvinick and colleagues' conflict monitoring model, and showed that item-level control is sufficient to explain the proportion-congruent effect. That is, the model need not be sensitive to the global amount of conflict to demonstrate a proportion-congruent effect. Thus, despite the fact that the ACC–DLPFC conflict monitoring loop has been associated with the *overall* amount of conflict in a series of trials, it may actually be responding to the conflict associated with an individual *item* (e.g., Bugg, Jacoby, & Toth, 2008; Notebaert & Verguts, 2008; Verguts & Notebaert, 2008; Blais et al., 2007).

Previous research has typically overlooked the possibility that item-level conflict might elicit cognitive control; indeed, the typical cognitive control experiment confounds global and local sources of conflict. For example, if 80% of the trials within a Stroop block are congruent, each item is also congruent on 80% of the trials. Here, we devised a stringent behavioral test of the item-level account of cognitive control, in which we independently assessed the item-level and global contributions to the proportion-congruent effect in the Stroop task (see also Bugg et al., 2008, Experiment 1).

The general structure of the experiment is depicted in Figure 1A. Within each block of trials, the overall proportion of congruent trials was either 30%, 50%, or 70%; this constituted the *global* manipulation of conflict. Critically, embedded within each block was an item-level proportion-congruent manipulation. In the 30% congruent block, two items were congruent (i.e., color names printed in the corresponding ink color) on 10% of the trials, and two items were congruent on 50% of the trials. In the 50% condition, all four items were congruent on 50% on the trials. In the 70% block, two items were congruent on 50% of the trials, and two items were congruent on 90% of the trials. As is clear from Figure 1B, the proportion manipulation *within* the 30% and 70% blocks (i.e., 10% vs. 50% and 50% vs. 90% congruent) indexes the *local* conflict effect, whereas the comparison of the 50% conditions across the three blocks indexes the global conflict effect.

This design differs in several ways from the standard Stroop task; our goal was to design a variant of this task to study an item-level, potentially automatic, form of cognitive control. The current design enabled us to assess the presence or absence of both global and local effects.

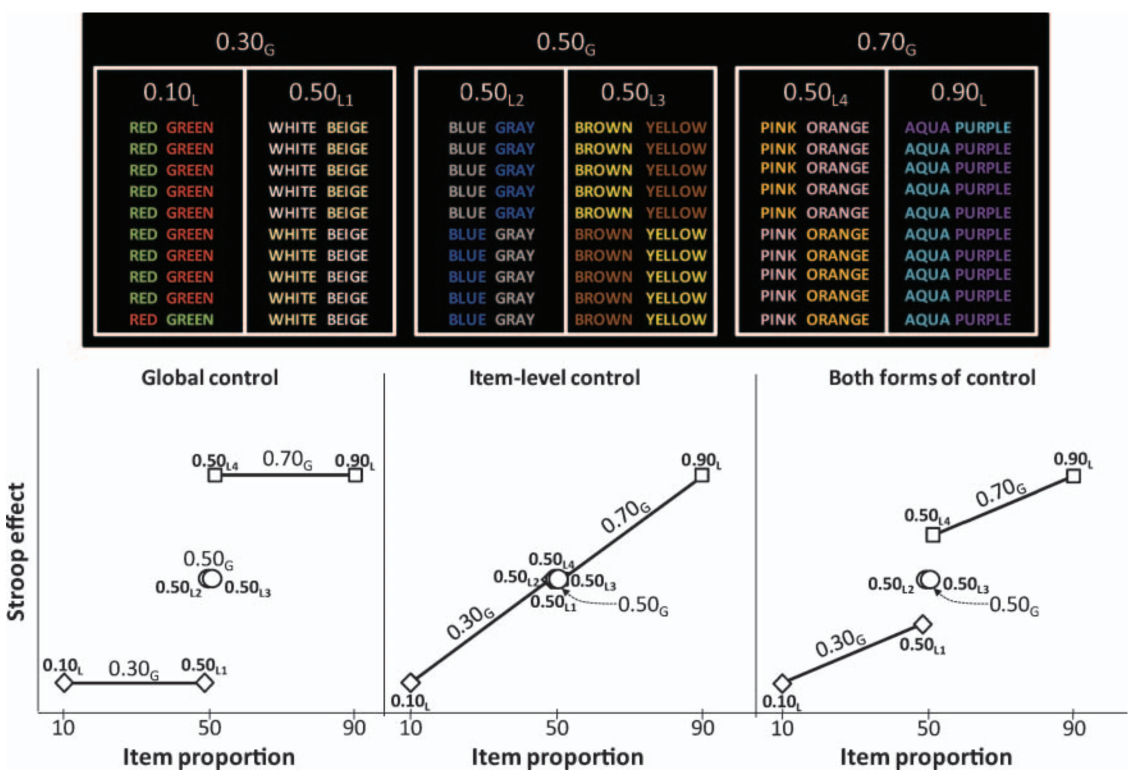

Figure 1. (Top) An outline of the composition of trials. Within a block, the overall (global) proportion of congruent trials was either 30%, 50%, or 70% congruent trials. Nested within each block was an item-level proportion-congruent manipulation (i.e., a local manipulation). In the 30% congruent block, two items were congruent on 10% of the trials, and two items were congruent on 50% of the trials. In the 50% block, all items were congruent on 50% of the trials. In the 70% block, two items were congruent on 50% of the trials, and two items were congruent on 90% of the trials. Comparing the 50% items across lists allows us to assess the influence of global effects. Comparing the items within a list (e.g., 10% vs. 50% within the 30% list) allows us to assess the influence of item effects. (Bottom) Plotted here are the predicted outcomes, depending on whether only global control contributes to the proportion-congruent effect (left), only item-level control contributes to the proportion-congruent effect (center), or both forms of control contribute to the proportion-congruent effect (right). G = global/block effects; L# = local/item effects.

Downloaded from ... on 18 May 2021

The left panel of Figure 1B shows the pattern of data that one would expect if proportion effects were driven entirely by global effects, a pattern of results that would be inconsistent with prior behavioral studies providing evidence of item-specific proportion-congruent (ISPC) effects. The middle panel of Figure 1B shows the pattern of data that one would expect if proportion effects were driven entirely by local effects. Note that, in this case, data for the four local 0.50 conditions, which differ based on the global context in which they are encountered, should lie atop one another. Finally, the right panel of Figure 1B shows the pattern of data that one would expect if proportion effects were driven by both global and local effects.

EXPERIMENT 1: BEHAVIORAL STUDY

Methods

Participants

Eighteen undergraduates (12 women) from the University of California, Berkeley, served as participants. Participants were paid $12 per hour. Informed consent was acquired from all participants in accordance with the Institutional Review Board at the University of California at Berkeley.

Task

Each trial began with a fixation marker for 500 msec. The color word then appeared and remained on the screen for up to 1500 msec or until a response was made. Although the duration of the target was short, missed responses occurred on less than 0.1% of trials across participants. The total trial duration was 2 sec. An additional ITI of 0, 2, 4, 6, 8, 10, or 12 sec followed the color word. The subjective experience is of a constantly present fixation cross that is replaced by the color word. Response to the item caused the fixation marker to appear immediately. The trial order and ITI durations were optimized for an event-related fMRI study by optseq2 (http://surfer.nmr.mgh.harvard.edu/optseq/).

There were 12 colors divided into groups that three raters (the authors and a colleague in the lab) agreed were clearly distinguishable from one another yielding three sets of four colors [RGB values follow in parentheses: pink (255, 192, 203), green (0, 176, 80), brown (139, 69, 19), yellow (255, 255, 0); gray (127, 127, 127), red (255, 0, 0), blue (0, 112, 192), beige (210, 180, 140); and white (255, 255, 255), orange (255, 192, 0), aqua (102, 205, 170), purple (112, 48, 160)].

These three sets of four colors were paired with each of the 30%, 50%, and 70% congruent blocks. The order of the blocks and color sets were fully counterbalanced across participants. Importantly, in all three blocks, there was a set of items (120 trials) that were congruent on 50% of trials. In the 30% block, the other two items were congruent for 10% of the trials (i.e., 12 congruent items and 108 incongruent items). In the 50% block, both sets of items were congruent on 50% of the trials. Finally, in the 70% block, the other two items were congruent on 90% of the trials (i.e., 108 congruent items and 12 incongruent items). In total, there were 240 trials per block, yielding a total of 720 trials.

Four color patches were located at the bottom of the screen throughout the experiment. Participants were told that these patches corresponded spatially to four keys on a standard keyboard (d, f, j, and k). The first and third buttons were assigned to one proportion condition and the second and fourth to the other. Each experimental block was preceded by 20 trials, five of each of the four colors, in which participants identified the color of a large rectangle presented at fixation to learn the set of four colors that appeared in each block.

Results

Figure 2A shows the mean response time for each condition. As shown in Figure 3, the proportion effect is attributable entirely to an item-level effect, consistent with the prediction in the center panel of Figure 1B. Formal tests confirm this pattern of results, as described below.

Interpreting a more tradition analysis of this experiment looking at List proportion (30% block vs. 70% block) × Item proportion (low congruence vs. high congruence) × Congruency (congruent vs. incongruent) is difficult. By definition, the item proportion and the list proportion are not independent as the item composition creates the list. Thus, if we find a main effect of list or any interactions with list, it is unclear whether the effect is driven by global or local processes or both. Instead, we opt to report more direct tests for global and local processes as noted below.

To assess the presence of an item-level proportion-congruent effect, independent of the list-level proportion-congruent effect, correct RT data from the 30% and 70% blocks were analyzed separately using a 2 (item proportion congruent) × 2 (congruency) repeated measures ANOVA. The critical interaction between proportion congruent and congruent was present for both the 30% block [$F(1, 17) = 13.07$, $MSE = 2297$, $p < .005$; partial $\eta^2 = .435$] and the 70% block [$F(1, 17) = 7.03$, $MSE = 2216$, $p < .05$; partial $\eta^2 = .292$]. The Stroop effect was smaller for the 10% items (−35 ± 16 msec) than the 50% items (47 ± 10 msec) in the 30% block and smaller for the 50% items (37 ± 9 msec) than the 90% items (96 ± 20 msec) in the 70% block. In sum, there is an ISPC effect independent of the list-level proportion effect.

To assess the presence of a list-level specific proportion-congruent effect, independent of the item-level proportion-congruent effect, the size of Stroop effect for the 50% items in each of the three (30%, 50%, and 70%) blocks was compared. All six comparisons were nonsignificant (Fs < 1.6, partial $\eta^2 < .087$). The strongest test for a

Downloaded from http://direct.mit.edu/jocn/article-pdf/22/12/2758/1769783/jocn.2009.21365.pdf by guest on 18 May 2021

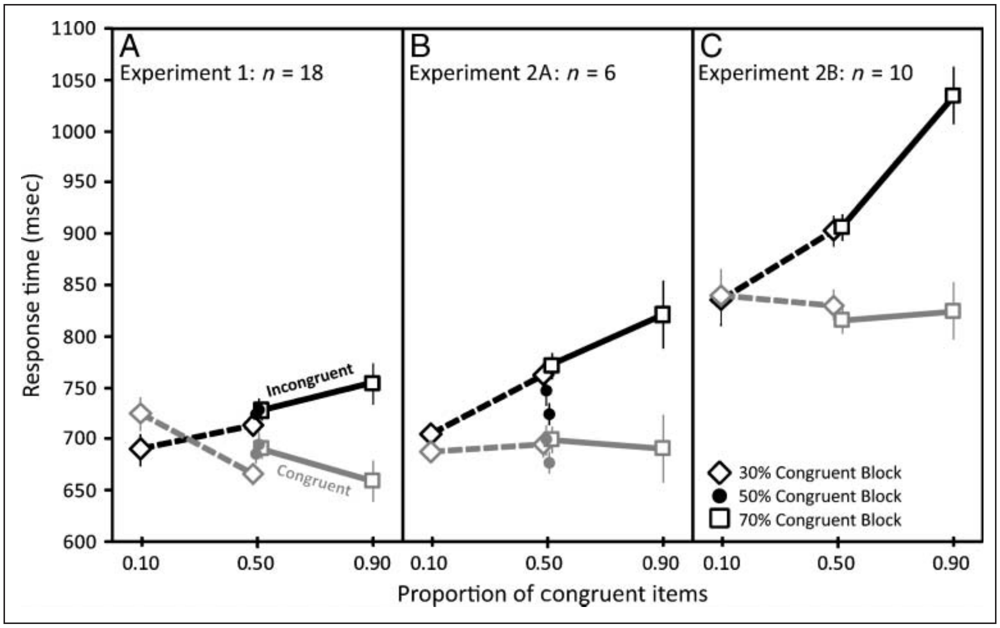

Figure 2. Response time data from Experiments 1 and 2. Error bars represent the standard error of the difference between the congruent and incongruent conditions at each item-level proportion congruence level such that nonoverlapping error bars indicate a significant Stroop effect at $\alpha < .05$.

list-level proportion-congruent effect, independent of an item-level effect is the comparison between the 50% items and the 30% and 70% blocks. This contrast had a very small effect size (partial $\eta^2 = .043$) and was 10 ± 12 msec in the wrong direction (larger in the 30% block than in the 70% block). In sum, there is no list-level proportion-congruent effect independent of the item-level proportion-congruent effect.

Discussion

The results of Experiment 1 are consistent with the hypothesis that an item-level control mechanism influences the size of the Stroop effect. Specifically, the size of the Stroop effect increased as a function of the proportion of congruent trials at the item level. Furthermore, the results strongly suggest that a list-level control mechanism did not operate in this task because list-wide proportion congruence has no effect on the magnitude of Stroop interference. These results exactly replicate Bugg et al. (2008, E1), who showed that the Stroop effect was driven by item-specific levels of proportion congruence, and not list-wide levels of proportion congruence across both college-aged and older adults in a vocal analogue of this task.

This finding is inconsistent with several past reports (e.g., Logan, Zbrodoff, & Williamson, 1984; Lowe & Mitterer, 1982), which show list-wide proportion congruence effects. In traditional studies, the list-wide proportion congruence manipulation is perfectly confounded with the item-specific manipulation. For example, if the list-level proportion is 80% congruent, then each color within the list is presented congruently on 80% of trials. Here, we manipulated the list-wide proportion congruence while holding constant the item-specific proportion congruence for a subset of the items. This design allows us to elucidate the distinct contributions of list-level and item-level control. The clear result from this experiment is that item-level control exerts an influence on Stroop performance, but that list-level control does not (see

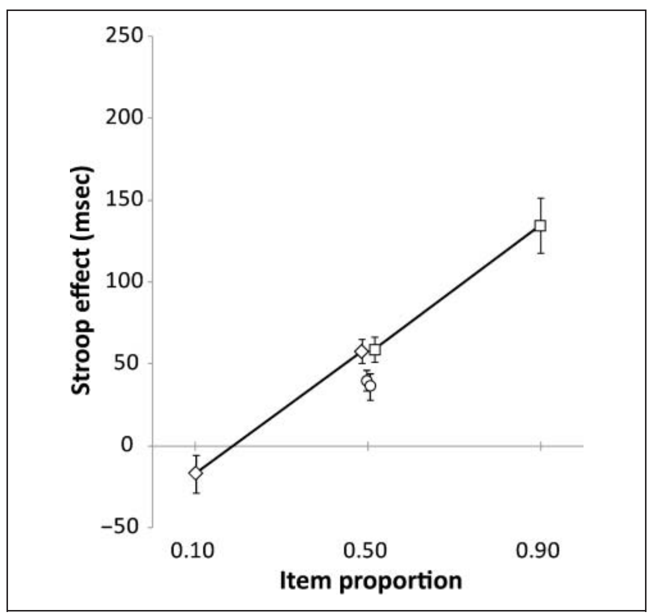

Figure 3. The average Stroop effect across Experiments 1 and 2 ($n = 34$). The size of the Stroop effect is plotted as a function of the proportion of congruent trials at the item level. The error bars represent one *SEM* for the Stroop effect at that proportion level. Note: Ten participants in Experiment 2 did not see the 50% congruent block; therefore, the circles represent data from 24 participants rather than the total sample of 34; that these data points are off the line is consistent with the fact that RTs in Experiment 2B are longer than Experiment 1 and Experiment 2A and that longer RTs have long been known to be associated with increased Stroop effects.

Downloaded from ... on 18 May 2021

also Bugg et al., 2008). These findings call into question the assumption that proportion-congruent effects in previous studies were driven entirely by list-level effects. Further, these findings call into question the assumption that the ACC–DLPFC loop implements a list-wide control mechanism (e.g., Botvinick et al., 2001). Experiment 2 assesses the possibility that the ACC–DLPFC control network is sensitive to a congruency manipulation operating at the level of a single item.

EXPERIMENT 2: fMRI STUDY

Methods

Participants

Sixteen adults (10 women) from the San Francisco Bay Area served as participants. Participants were paid $15 per hour plus $10 for transportation. All participants were recruited from the San Francisco Bay Area. Informed consent was acquired from all participants in accordance with the Institutional Review Board at the University of California at Berkeley.

Data Acquisition

The experiment was run in E-Prime 1.2 on a Mac Pro computer running Windows XP SP2. Scanning was performed with a standard whole-head coil on a 3-Tesla Siemens Trio scanner at the Neuroscience Imaging Center at the University of California, San Francisco. Visual stimuli were displayed on a 40-inch LCD monitor with a resolution of 1360 × 768 that participants viewed through a mirror. fMRI data were acquired with a gradient-echo echo-planar pulse sequence (TR = 2 sec, TE = 25 msec, 33 axial slices, 1.8 × 1.8 × 3 mm, 15% interslice gap). Before each scan, three volumes were discarded to allow for T1-equilibration effects. High-resolution T1-weighted anatomical images were collected. Head motion was restricted through the use of a pillow and foam inserts that surrounded the head.

fMRI Analysis

Data were processed with SPM5 (Wellcome Department of Cognitive Neurology, London). Images were corrected for differences in timing of slice acquisition, followed by rigid-body motion correction. Structural and functional volumes were spatially normalized to T1 and echo-planar imaging templates, respectively. The normalization algorithm used a 12-parameter affine transformation together with a nonlinear transformation involving cosine basis functions. During normalization, the volumes were resampled to 3-mm^3 voxels. Templates were based on the MNI305 stereotaxic space. Functional volumes were spatially smoothed with an 8-mm full-width at half-maximum isotropic Gaussian kernel.

Statistical analyses were performed on individual participants' data by using the general linear model fMRI time-series data that were modeled by the series of events convolved with a canonical hemodynamic response function. Error trials and misses were modeled separately and excluded from the analyses. The correct trial functions were used as covariates in a general linear model, along with a basic set of cosine functions that high-pass filtered the data (120 Hz) and a covariate for session effects. The least-squares parameter estimates of height of the best-fitting canonical hemodynamic response function for each condition were used in pairwise contrasts. The resulting contrast images, computed on a subject-by-subject basis, were submitted to group analyses. At the group level, contrasts between conditions were computed by performing one-tailed *t* tests on these images, treating participants as a random effect.

Task

Six participants completed the task exactly as described in Experiment 1 (Experiment 2A). For reasons explained below, the remaining 10 participants completed a slightly different version (Experiment 2B).

In Experiment 2B, each trial began with a fixation marker for 250 msec. The color word then appeared, and remained on the screen for up to 1750 msec, or until a response was made. A subset of eight of the colors was divided into groups that three raters (the authors and a colleague in the lab) agreed were clearly distinguishable from one another, yielding two sets of four colors: pink, green, purple, and yellow; red, blue, orange, and brown. The trial structure is shown in Figure 1A; the full design matrix is included in Supplementary Tables 1 and 2.

To maximize the number of trials and minimize movement artifacts, we chose to use a manual response. Concerns regarding this change in the Stroop paradigm are minimized by a prior fMRI study showing that the same region of ACC was involved in resolving conflict for both vocal and manual responses (Barch et al., 2001). Participants used a fORP button box to respond in the scanner. The first and third buttons were assigned to one proportion condition, and the second and fourth to the other.

Prior to each experimental block, participants learned the appropriate response key mappings for each of four colors. They performed a total of 52 practice trials in which they merely pressed one of four buttons to indicate the color of a large rectangle presented at fixation.

Preliminary fMRI results comparing all trials to the null condition in Experiment 2A revealed weaker than expected activation in visual cortex, a region that is typically robustly engaged in such comparisons. We wanted to ensure that this diminished response resulted from the fact that the response options were visible on the screen throughout the null period, resulting in greater baseline visual cortical activity than if a simple fixation cross-hair had been presented during the null period. Therefore, in

Downloaded from http://direct.mit.edu/jocn/article-pdf/22/12/2758/1769813/jocn.2009.21365.pdf by guest on 18 May 2021

Experiment 2B, the response options were not presented visually on the screen. This modification meant that participants had to memorize the response mappings. To reduce potential proactive memory effects, the number of colors was decreased from 12 to 8 by eliminating the 50% block. Two additional changes were made to the task in Version B. First, the duration of an individual run was decreased from 10 min 46 sec to 6 min 6 sec, but now participants received two runs for each conditions rather than one (the order of runs was either 30–30–70–70 or 70–70–30–30). Second, which pair of items went with which response keys was randomly selected for each participant. Third, as a result of these changes, we anticipated slightly longer reaction times for Version B so participants were given 1750 msec to respond instead of 1500 msec. These changes were sufficient in alleviating the concern regarding weak activation in visual cortex.

Results

Behavioral Analyses

As indicated in Figure 2B and C, the pattern of behavioral results was similar across the two versions of the task. Furthermore, when looking across all 16 subjects, the proportion effect is attributable entirely to an item-level effect, consistent with the prediction in the center panel of Figure 1B. As a formal test of whether there was an item-level proportion-congruent effect that is independent of the list-level proportion-congruent effect, the correct response time data from the 30% and 70% blocks were analyzed separately using a 2 (Experiment Version A vs. B) × 2 (item proportion congruency) × 2 (congruency) mixed ANOVA. Because none of the interactions with experiment version were significant, we collapsed across both versions.

Across Experiments 2A and 2B, the critical ISPC effect was present in both the 30% block [$F(1, 15) = 12.43$, $MSE = 976$, $p < .003$; partial $\eta^2 = .453$] and the 70% block [$F(1, 15) = 19.58$, $MSE = 1437$, $p < .001$; partial $\eta^2 = .566$]. The Stroop effect was smaller for the 10% items than the 50% items (8 ± 12 msec vs. 63 ± 9 msec) in the 30% block and smaller for the 50% items than the 90% items (78 ± 8 msec vs. 161 ± 23 msec) in the 70% block. All of these effects are statistically significant for Versions A and B when analyzed separately.

To test for a list-level effect that is independent from the item-level effect shown above, we compared the size of the Stroop effect for the 50% items from each of the 30% and 70% blocks. The Stroop effect was 15 ± 10 msec smaller ($p > .15$, partial $\eta^2 = .13$) for the 50% items in the 30% block than for the 50% items in the 70% block. Although Versions A and B do not differ statistically, this 15-msec increase is driven entirely by the 10 subjects in Version B (4 ± 15 msec vs. 18 ± 19 msec in Versions A and B, respectively, $p > .50$). If this difference is real (i.e., this is a Type II error), it provides an indication for a list-wise effect, but it is numerically very small in comparison to the item-level effect (as seen in Figure 2). We speculate that forcing subjects to hold on to the key labels in working memory may cause subjects to devote more effort to the task. It is possible that this increased level of engagement is necessary to produce a list-level effect.

The error rates were low (<2.5%) overall, and the only effect that was consistent was the overall Stroop effect [2.2% vs. 1.1% for incongruent vs. congruent trials in Version A, $F(1, 5) = 4.91$, $p = .078$; 4.6% vs. 1.3% in Version B, $F(1, 9) = 8.17$, $p < .01$].

Combined Behavioral Analysis

When looking across all 34 subjects from Experiments 1, 2A, and 2B, it becomes apparent that the proportion effect in this task is entirely attributable to the item-level effect. The ISPC effect was present in both the 30% block [$F(1, 33) = 23.0$, $MSE = 2056$, $p < .001$; partial $\eta^2 = .411$] and the 70% block [$F(1, 33) = 26.4$, $MSE = 1839$, $p < .001$; partial $\eta^2 = .445$]. The Stroop effect was smaller for the 10% items than the 50% items (−17 ± 12 msec vs. 58 ± 7 msec) in the 30% block and smaller for the 50% items than the 90% items (59 ± 8 msec vs. 135 ± 17 msec) in the 70% block. There is no behavioral evidence to support a list-level effect that is independent from the item-level effect in these data. Indeed, the Stroop effect was 1 ± 9 msec smaller ($p > .50$, partial $\eta^2 = .001$) for the 50% items in the 30% block than for the 50% items in the 70% block. Focusing on the 24 subjects who also did a 50% block, of the other five pairwise comparisons between the local 50% conditions, only two comparisons had $.09 < p < .20$ (0.50_{L1} vs. 0.50_{L2} and 0.50_{L1} vs. 0.50_{L3}). In sum, there is no compelling evidence to support a global proportion effect that is independent of the item-level proportion-congruent effect.

fMRI Analyses

To assess exactly what form(s) of control the ACC–DLPFC loop is responsive to, we identified regions of interest (ROIs) in ACC and DLPFC from a contrast used in prior Stroop studies (e.g., Carter et al., 2000). This whole-brain contrast consisted of a *T*-contrast comparing the Incongruent > Congruent contrast in the 70% block with the Incongruent > Congruent contrast in the 30% block (i.e., regions for which the Stroop effect is larger in the 70% block than the 30% block; Figure 4B). Based on the traditional lines of research, this contrast should identify the "strategic" component of control (e.g., Logan, Zbrodoff, & Williamson, 1984). However, based on our behavioral results, we argue that they may, in fact, be implementing item-specific control (see Figure 1A–C).

To isolate activation clusters that were limited to ACC and DLPFC, we computed the whole-brain contrast at $p < .0005$, uncorrected, cluster size ≥10. The peak coordinate for our ACC cluster for this contrast [−6 18 30; Figure 4]

Downloaded from http://direct.mit.edu/jocn/article-pdf/22/12/2758/1768826/jocn.2009.21365.pdf by guest on 18 May 2021

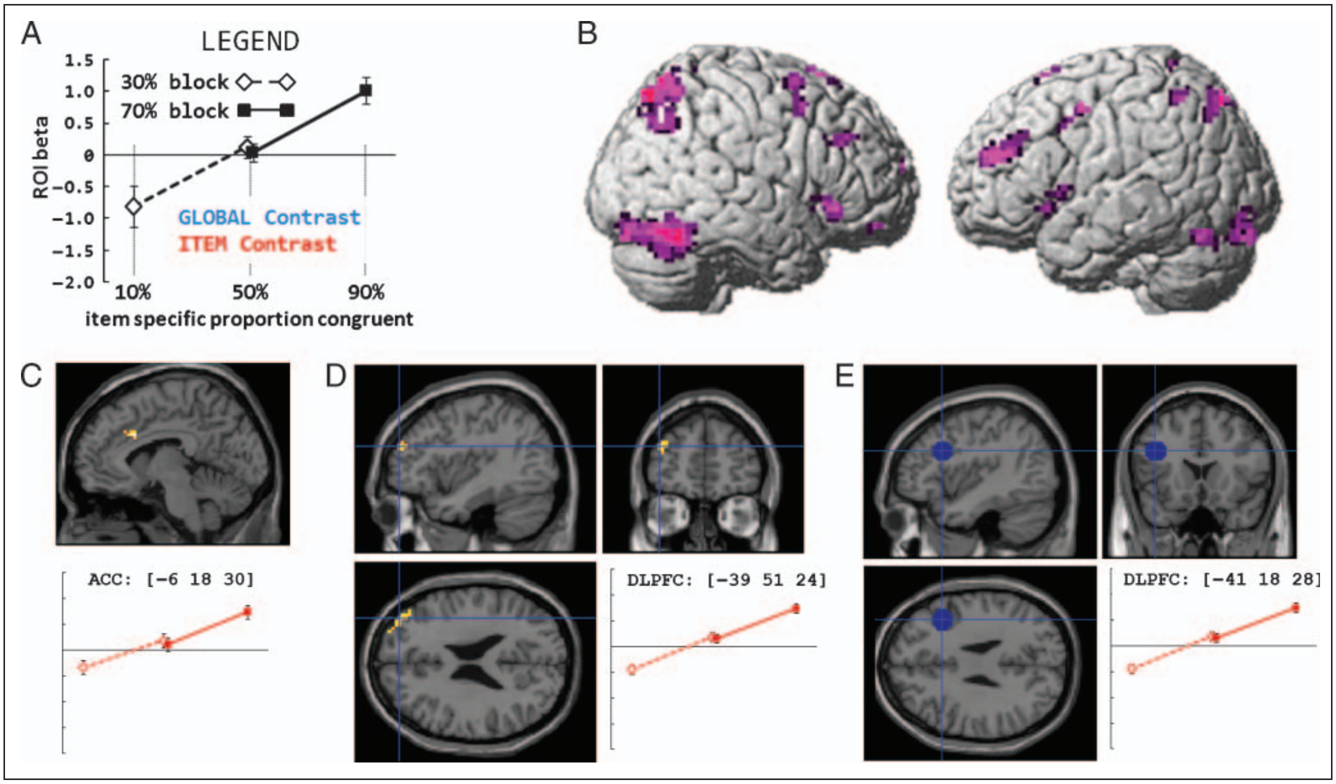

Figure 4. fMRI results. (A) The legend for the graphs displayed in B and C. Each point represents the Incongruent > Congruent contrast ("Stroop effect") parameter estimates. The global contrast, shown in blue, is the comparison between the Incongruent > Congruent contrasts for the 50% condition embedded within the 30% and 70% blocks. The local contrast, shown in red, is the average item-specific proportion-congruent effect in the 70% block [i.e., $(90\%_I–90\%_C)–(50\%_I–50\%_C)$] and the 30% block [i.e., $(50\%_I–50\%_C)–(10\%_I–10\%_C)$]. These colors apply to Figure 5. (B) Activation for the standard proportion contrast of Incongruent > Congruent in the 70% block compared to the Incongruent > Congruent contrast in the 30% block. $p < .005$, uncorrected, cluster size ≥10. (C and D) ACC and DLPFC regions identified from B at $p < .0005$, uncorrected, cluster size ≥10. (E) DLPFC as identified by MacDonald, Cohen, Stenger, & Carter (2000). DLPFC = dorsolateral prefrontal cortex; ACC = anterior cingulate cortex.

is well within the range of the Barch et al. (2001) meta-analysis (see Supplementary Table 3). However, because the peak for our DLPFC cluster [−39 51 24; Figure 4] is anterior to that of prior studies, we additionally constructed a 10-mm sphere centered on a DLPFC region known to be involved in response modulation [−41 18 28; MacDonald et al., 2000; Figure 4].

Formal analysis confirms the pattern that is evident in Figure 4 in which the item-specific effect dominates any potential global effect. There is an ISPC effect in the 70% block, such that the Stroop effect is larger for 90% congruent items than for 50% congruent items. There is also an ISPC effect in the 30% block, such that the Stroop effect is larger for 50% congruent items than for 10% congruent items. These ISPC effects were obtained for all three ROIs (paired t test, $p < .05$, Bonferroni corrected), indicative of a local control mechanism. Comparing the two 50% conditions across blocks yields no differences for any of these three ROIs (smallest $p = .534$, uncorrected), indicative of a lack of list-level control.

ACC and DLPFC have been argued to implement list-level control effects (Botvinick et al., 2001). However, using a design that can dissociate item-level from list-level effects, we find evidence only for item-specific control. This is not to say that these and/or other frontal regions are incapable of implementing a more global form of control. However, this study provides clear evidence that the network is capable of implementing control at the level of specific items, a level that is generally considered to be outside the realm of strategic control.

The findings presented in Figure 4 are all the more compelling given that the contrast used to identify the cognitive control regions (comparing the Stroop effect for the 70% congruent block vs. the 30% congruent block, similar to Carter et al., 2000 who used 80% and 20% congruent blocks) is heavily biased toward identifying regions involved in global control. However, as noted below, we also sought to identify more clearly the sets of regions exhibiting local or global effects, using more targeted whole-brain contrasts (Figure 5).

To identify regions uniquely engaged in local control, a whole-brain contrast was conducted identifying the average of the ISPC effect [$(\text{Incongruent}_{90\%\_Con} > \text{Congruent}_{90\%\_Con}) > (\text{Incongruent}_{50\%\_Con} > \text{Congruent}_{50\%\_Con})$] in the 70% block and [$(\text{Incongruent}_{50\%\_Con} > \text{Congruent}_{50\%\_Con}) > (\text{Incongruent}_{10\%\_Con} > \text{Congruent}_{10\%\_Con})$] in the 30% block. That is, this contrast compares the Stroop effect for the specific items associated with a higher

Downloaded from ... on 18 May 2021

Figure 5. See Figure 4A for the legend of these graphs. All significant clusters are reported in Supplementary Table 3. $p < .005$, uncorrected, cluster size ≥10. MFG = middle frontal gyrus; SFG = superior frontal gyrus; ACC = anterior cingulate cortex; BA = Brodmann's area.

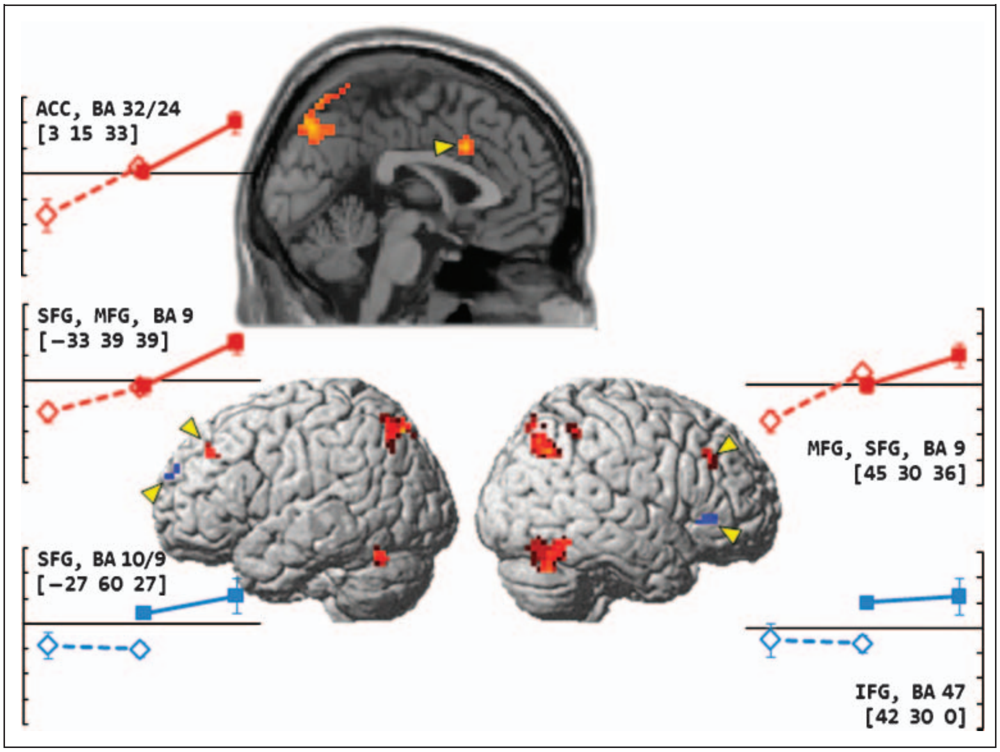

proportion congruency (90% vs. 50%, and 50% vs. 10%). As shown in red in Figure 5, the set of activations revealed by this contrast is similar to that shown previously in Figure 4.

To identify regions sensitive to the global-level manipulation of proportion of congruent trials in a block, a whole-brain contrast was conducted comparing $\text{Incongruent}_{50\%\_\text{Con}} > \text{Congruent}_{50\%\_\text{Con}}$ in the 30% block versus $\text{Incongruent}_{50\%\_\text{Con}} > \text{Congruent}_{50\%\_\text{Con}}$ in the 70% block. This contrast, shown in blue in Figure 4, yielded activation clusters in right ventrolateral PFC (VLPFC; BA 47) and left anterior DLPFC (superior frontal gyrus; BA 10/9).

The Role of Associative Learning

An issue that is heavily debated in the behavioral literature concerns whether the ISPC effect results from a combination of control and associate learning (the control hypothesis; e.g., Bugg et al., 2008; Jacoby et al., 2003), or whether it is the result of purely associative learning (the contingency hypothesis; Schmidt & Besner, 2008). To adjudicate between these hypotheses, Schmidt and Besner (2008) advocate a Contingency by Congruency analysis. Both the control hypothesis and the contingency hypothesis predict main effects of congruency and contingency, but only the control hypothesis predicts that these factors will interact because "incongruent trials should be more affected by attention, given that the majority of the Stroop effect is interference, with little or no facilitation from congruent trials" (Schmidt & Besner, 2008, p. 516).

To conduct this analysis, we combined the two 50% conditions from the 30% and 70% blocks which yielded three contingency conditions (*high* with 108 trials, *medium* with the average of 60 trials, and *low* with 12 trials).[1] This Contingency (high vs. medium vs. low) by Congruency analysis was conducted for Experiments 1, 2A, and 2B, which yielded main effects of contingency and congruency for each of these experiments (All Fs > 7, ps $< .012$). These factors did not interact in Experiment 1 ($F < 1$). Critically, however, they did interact in Experiments 2A [$F(2, 10) = 4.60$, $p < .05$] and 2B [$F(2, 18) = 4.25$, $p <$ 05]. We are unsure as to why Experiment 1 failed to show this interaction. Nonetheless, this interaction was observed in Experiments 2A and 2B, which supports our claim that ACC–DLPFC supports item-specific control.

Discussion

These behavioral and neural results indicate that the ACC–DLPFC network modulates attention based on the relative conflict associated with a *specific* item. The data provide no evidence to suggest that the ACC–DLPFC network modulates attention based on *overall* conflict. Specifically, when item proportion was controlled, the ACC–DLPFC network was not affected by the overall proportion of congruent trials. This result signals the need for a critical departure from current thinking regarding cognitive control as indexed by proportion-congruent effects. These findings strongly support the item-specific conflict monitoring model (Blais et al., 2007). We do not claim that our results fully explain participants' performance

Downloaded from [illegible] on 18 May 2021

on the standard Stroop paradigm—rather, that our paradigm provides evidence that there are situations under which cognitive control is implemented at the level of individual stimuli.

The demonstration here that the ACC–DLPFC network was insensitive to global amounts of conflict does not mean that cognitive control never operates at the global level. Even in this dataset, two regions in lateral PFC were sensitive to global-level effects; incidentally, these activations appear to have no impact on behavior, given that there were no global effects on performance. Rather, our claim is that the ACC–DLPFC network does not, in this task, rely on or implement representations consistent with global control. Both item-level and global effects are likely to come into play in most cognitive control tasks, with the balance depending on specific task demands. We would expect a larger contribution from global effects on tasks that involve deterministic rather than probabilistic rules (Nieuwenhuis, Schweizer, Mars, Botvinick, & Hajcak, 2007; Bunge, 2004; Logan et al., 1984).

These results are consistent with the general premise that the conflict monitoring system is a "dumb system" that modulates attention on the basis of the input that the system receives (Botvinick, Cohen, et al., 2004). They are, however, inconsistent with Botvinick et al.'s (2001) specific implementation of this hypothesis: Conflict signals do not relay information to general task-level demands such as "ignore the word," but rather to specific task demands such as "ignore the word green." This finding has implications for the specificity of representations that elicit cognitive control via ACC–DLPFC circuitry.

Further research is needed to examine the nature of the memory representations that underlie the ISPC effect (e.g., see Verguts & Notebaert, 2008). What is clear from the current study is that associations with individual stimuli influence behavior and brain function. The form of learning that underlies this phenomenon is rapid: Participants exhibit the ISPC effect within as little as 16 trials (Jacoby et al., 2003, E3). We hypothesize that these item-specific associations are learned implicitly and that they are updated continuously in a dynamic setting (see Crump, Vaquero, & Milliken, 2008; Pasupathy & Miller, 2005).

These item-specific effects must be based, at least in part, on an associative learning mechanism such as contingency (S–R) learning, whereby participants produce the response most frequently associated with a particular word (Schmidt & Besner, 2008; Jacoby et al., 2003; Melara & Algom, 2003). The body of evidence to date does not allow one to make strong conclusions about the relative contributions of cognitive control versus associative mechanisms to the item-specific effect and the neural responses that accompany it. Importantly, however, the present findings provide compelling evidence for item-specific control mediated by brain regions that have been implicated previously in strategic, goal-directed behavior.

The present results are inconsistent with an account in which the ACC–DLPFC network modulates attention based on global attributes of the present task, such as the overall proportion of congruent trials. Rather, the data are consistent with an account in which the ACC–DLPFC network modulates attention based on the local attributes of the task, such as the conflict associated with a specific item. These findings mark a new way of thinking about the representations upon which cognitive control operates.

Acknowledgments

This research was supported by a postdoctoral fellowship grant to C. B. from the Natural Sciences and Engineering Research Council of Canada (NSERC).

Reprint requests should be sent to Chris Blais, University of California, Berkeley, Helen Wills Neuroscience Institute, 132 Barker Hall MC 3190, Berkeley, CA 94720, or via e-mail: cblais@berkeley.edu.

Note

1. An analysis which ignores the medium condition yields a similar pattern of results.

REFERENCES

Barch, D. M., Braver, T. S., Akbudak, E., Conturo, T., Ollinger, J., & Snyder, A. (2001). Anterior cingulate cortex and response conflict: Effects of response modality and processing domain. *Cerebral Cortex, 11,* 837–848.

Blais, C., Robidoux, S., Risko, E. F., & Besner, D. (2007). Item specific adaptation and the conflict monitoring hypothesis: A computational model. *Psychological Review, 114,* 1076–1086.

Botvinick, M., Braver, T. S., Barch, D. M., Carter, C. S., & Cohen, J. D. (2001). Conflict monitoring and cognitive control. *Psychological Review, 108,* 624–652.

Botvinick, M., Braver, T. S., Yeung, N., Ullsperger, M., Carter, C. S., & Cohen, J. D. (2004). Conflict monitoring: Computational and empirical studies. In M. I. Posner (Ed.), *Cognitive neuroscience of attention* (pp. 91–102). New York: Guilford Press.

Botvinick, M., Cohen, J. D., & Carter, C. S. (2004). Conflict monitoring and anterior cingulate cortex: An update. *Trends in Cognitive Sciences, 8,* 539–546.

Bugg, J. M., Jacoby, L. L., & Toth, J. P. (2008). Multiple levels of control in the Stroop task. *Memory & Cognition, 36,* 1484–1494.

Bunge, S. A. (2004). How we use rules to select actions: A review of evidence from cognitive neuroscience. *Cognitive, Affective, and Behavioral Neuroscience, 4,* 564–579.

Carter, C. S., MacDonald, A. M., Botvinick, M., Ross, L. L., Stenger, V. A., Noll, D., et al. (2000). Parsing executive processes: Strategic vs. evaluative functions of the anterior cingulate cortex. *Proceedings of the National Academy of Sciences, U.S.A., 97,* 1944–1948.

Cheesman, J., & Merikle, P. M. (1986). Distinguishing conscious from unconscious perceptual processes. *Canadian Journal of Psychology, 40,* 343–367.

Crump, M. J. C., Vaquero, J. M. M., & Milliken, B. (2008). Context-specific learning and control: The roles of awareness,

Downloaded from http://direct.mit.edu/jocn/article-pdf/22/12/2758/1769983/jocn.2009.21365.pdf by guest on 18 May 2021

task-relevance, and relative salience. *Consciousness and Cognition, 17,* 22–36.
Jacoby, L. L., Lindsay, D. S., & Hessels, S. (2003). Item-specific control of automatic processes: Stroop process dissociations. *Psychonomic Bulletin & Review, 10,* 638–644.
Lindsay, D. S., & Jacoby, L. L. (1994). Stroop process dissociations: The relationship between facilitation and interference. *Journal of Experimental Psychology: Human Perception and Performance, 20,* 219–234.
Logan, G. D., & Zbrodoff, N. J. (1979). When it helps to be misled: Facilitative effects of increasing the frequency of conflicting stimuli in a Stroop-like task. *Memory & Cognition, 7,* 166–174.
Logan, G. D., Zbrodoff, N. J., & Williamson, J. (1984). Strategies in the color–word Stroop task. *Bulletin of the Psychonomic Society, 22,* 135–138.
Lowe, D. G., & Mitterer, J. O. (1982). Selective and divided attention in a Stroop task. *Canadian Journal of Psychology, 36,* 684–700.
MacDonald, A. W., Cohen, J. D., Stenger, V. A., & Carter, C. S. (2000). Dissociating the role of dorsolateral prefrontal and anterior cingulate cortex in cognitive control. *Science, 288,* 1835–1838.
Melara, R. D., & Algom, D. (2003). Driven by information: A tectonic theory of Stroop effects. *Psychological Review, 110,* 422–471.
Nieuwenhuis, S., Schweizer, T. S., Mars, R. B., Botvinick, M. M., & Hajcak, G. (2007). Error-likelihood prediction in the medial frontal cortex: A critical evaluation. *Cerebral Cortex, 17,* 1570–1581.
Notebaert, W., & Verguts, T. (2008). Cognitive control acts locally. *Cognition, 106,* 1071–1080.
Pasupathy, A., & Miller, E. K. (2005). Different time courses of learning-related activity in the prefrontal cortex and striatum. *Nature, 433,* 873–876.
Schmidt, J. R., & Besner, D. (2008). The Stroop effect: Why proportion congruent has nothing to do with congruency and everything to do with contingency. *Journal of Experimental Psychology: Learning, Memory, & Cognition, 34,* 514–523.
Tzelgov, J., Henik, A., & Berger, J. (1992). Controlling Stroop effects by manipulating expectations for color words. *Memory & Cognition, 20,* 727–735.
Verguts, T., & Notebaert, W. (2008). Hebbian learning of cognitive control: Dealing with specific and nonspecific adaptation. *Psychological Review, 115,* 518–525.

Downloaded from http://direct.mit.edu/jocn/article-pdf/21/12/2758/1759883/jocn.2009.21365.pdf by guest on 18 May 2021