

Corticostriatal Output Gating during Selection from Working Memory

Christopher H. Chatham,^{1,*} Michael J. Frank,¹ and David Badre^{1,*}

¹Cognitive, Linguistic, and Psychological Sciences, Brown Institute for Brain Science, Brown University, Providence, RI 02912, USA

*Correspondence: chathach@gmail.com (C.H.C.), david_badre@brown.edu (D.B.)

<http://dx.doi.org/10.1016/j.neuron.2014.01.002>

SUMMARY

Convergent evidence suggests that corticostriatal interactions act as a gate to select the input to working memory (WM). However, not all information in WM is relevant for behavior simultaneously. For this reason, a second “output gate” might advantageously govern which contents of WM influence behavior. Here, we test whether frontostriatal circuits previously implicated in input gating also support output gating during selection from WM. fMRI of a hierarchical rule task with dissociable input and output gating demands demonstrated greater lateral prefrontal cortex (PFC) recruitment and frontostriatal connectivity during output gating. Moreover, PFC and striatum correlated with distinct behavioral profiles. Whereas PFC recruitment correlated with mean efficiency of selection from WM, striatal recruitment and frontostriatal interactions correlated with its reliability, as though such dynamics stochastically gate WM’s output. These results support the output gating hypothesis, suggesting that contextual representations in PFC influence striatum to select which information in WM drives responding.

INTRODUCTION

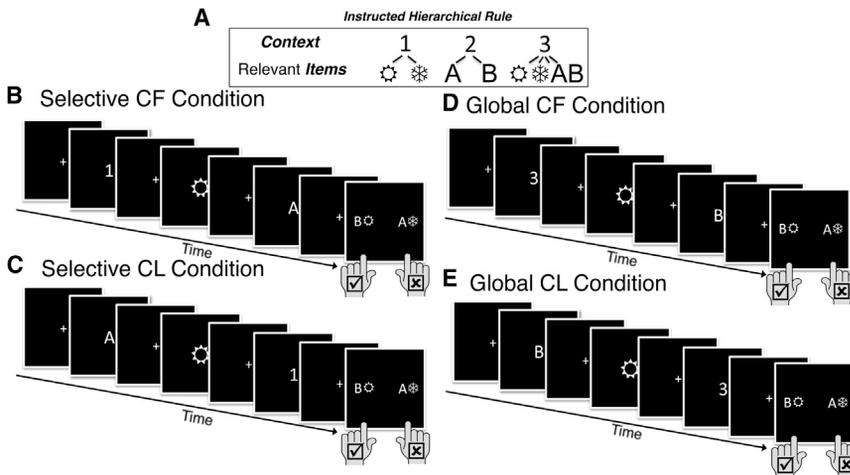
Cognitive control requires a balance between two incompatible demands: flexibly updating goals versus robustly maintaining them. One solution to this dilemma is to separate the neural mechanisms for working memory (WM) maintenance from those that update the information that is to be maintained, i.e., a WM “gate.” This computational division of labor is thought to have anatomical correlates, with prefrontal cortex (PFC) supporting maintenance and basal ganglia (BG) supporting gating (Braver and Cohen, 2000; Frank et al., 2001; Gruber et al., 2006; Frank and O’Reilly, 2006; Cools et al., 2007). From this perspective, disinhibition of cortico-striato-thalamic loops enables the selective updating of task-relevant information into WM. Once maintained, information supported by PFC is available to exert a top-down bias on posterior neocortex (Desimone and Duncan, 1995). This type of selective control over the input to WM, termed “input gating,” relies on dopaminergic corticostriatal systems

(Frank and O’Reilly, 2006; Moustafa et al., 2008; Cools et al., 2010; Murty et al., 2011; McNab and Klingberg, 2008).

However, not everything in WM will be relevant for behavior at any one point in time. Rather, it is also adaptive to control which representations within WM can influence attention and behavior and when. Such selection from within WM or “singling out” of WM representations (Oberauer and Hein, 2012) is resource demanding and PFC dependent (e.g., Rowe et al., 2000; Bunge et al., 2002; Hester et al., 2007; Tamber-Rosenau et al., 2011). Nevertheless, relative to control over the input to WM, little evidence exists regarding this type of selection process and its relationship to cognitive control.

One hypothesis is that selection from within WM can be conceived of as a gating function analogous to that described above for WM updating. From this perspective, an “output gate” may control the flow of information within WM between an actively maintained but inert state to one that is capable of exerting a top-down influence on behavior. In other words, for any given WM representation, when the output gate is closed, that representation would be maintained but would not have a top-down influence. Conversely, when the output gate is opened, the maintained representation provides a top-down contextual signal. Output gating is a shared solution to the problem of selection from within WM across many computational models (Hochreiter and Schmidhuber, 1997; Hazy et al., 2007; Brown et al., 2004; Kriete and Noelle, 2011; Chatham et al., 2011; Eliasmith et al., 2012; Frank and Badre, 2012; Collins and Frank, 2013; Huang et al., 2013).

Like input gating, output gating may be hypothesized to arise from cortico-striato-thalamic loops, wherein candidate contextual representations maintained in PFC act as input to dorsal striatum, which in turn amplifies one of these representations via its pallido-thalamic disinhibitory loop (Hazy et al., 2007; Kriete and Noelle, 2011; Chatham et al., 2011; Collins and Frank, 2013; Kriete et al., 2013). This putative output gating dynamic for selection of a WM representation is an extension of more established interactions between the cortex and dorsal striatum during selection of a candidate motor plan (e.g., Mink 1996; Graybiel, 1998; Gurney et al., 2001; Brown et al., 2004). According to this general class of accounts, some potential stimulus-response action plan represented by the premotor cortex is “gated out” by striatum and thereby becomes the motor plan executed by the primary motor cortex. However, the hypothesis that similar frontostriatal interactions could support selection of a WM representation currently lacks support. The present study seeks to fill this gap.



appears last—CL Condition in (C) and (E)—subjects must have updated both items in WM and can only use these cues for output gating. In the behavioral version only (not illustrated), these response mappings were presented at the bottom of the screen containing the final token.

To test output gating of WM, we focus on hierarchical control tasks requiring the use of conditional rules (e.g., the shape of a stimulus cues whether to attend to the size or color of a stimulus, and size or color then determines which response must be made). In the brain, simpler rules that directly map a stimulus to a response tend to recruit more caudal frontal cortex than rules involving higher-order contingencies (Koechlin et al., 2003; Badre and D'Esposito, 2007; Badre et al., 2009, 2010), with a potentially similar rostral-to-caudal gradient of function in striatum owing to its parallel rostrocaudal organization (Badre et al., 2010; Badre and Frank, 2012; Verstynen et al., 2012; Mestres-Missé et al., 2012). Of note, such conditional rules rely on higher order contexts (e.g., shape) to select which lower order context (e.g., the color or size) should drive action. Thus, following such rules might rely on output gating, if these lower order representations must be selected from within WM. Indeed, a recent computational model predicts that performance advantages arise from rostrocaudal frontostriatal organization specifically if rostral frontal areas modulate the output gating of more caudal frontostriatal circuits (Frank and Badre, 2012). Although this model fits behavioral and fMRI data during hierarchical rule learning (Badre and Frank, 2012), its predictions regarding the priority of output gating dynamics for hierarchical frontostriatal interactions are untested.

To this end, we used a second-order control task to separately manipulate demands on output or input gating, while also controlling for differences in WM load. Subjects based their responses on one of two possible letters and/or two possible “wingding” characters depending on the identity of a number cue (Figure 1). The number acted as a higher-level context by specifying which set of lower-level items (letters, wingdings, or both) would be relevant to the response for that trial. The final stimulus event for each trial also included a set of response mappings. Subjects pressed a button corresponding to the side of the response mappings (left versus right) on which the relevant item(s) appeared (Figure 1A).

Of note, each stimulus was presented in an unpredictable serial order on each trial. Thus, higher order context was pre-

Figure 1. Task Rules and Example Trials

(A) Numbers acted as higher-order context, specifying which of the lower-level items (two possible letters and two possible wingdings) would be relevant on each trial. The numbers 1 and 2 indicate that only the wingdings or letters, respectively, would be relevant for responses; these are “selective” contexts. The number 3 is termed a global context because it indicates that both the letter and the wingding would be relevant. (B–E) All trials conclude with response mappings, to which subjects must indicate (using a left or right button press) where the relevant item (or items, in the case of the global context) from that trial appears on the mappings screen. (Correct answer above is always “left”). Critically, the order of stimuli is rearrangeable. When a selective or global context appears first—CF Condition in (B) and (D)—subjects can use it to selectively input only relevant items into WM. By contrast, if context

sented either before (context first [CF]) or after (context last [CL]) the lower order contexts had appeared. Under CF conditions, subjects could use the number to select which of the subsequent items to update into WM. Thus, the CF conditions allow subjects to use an “input gating” strategy. Conversely, during CL conditions, subjects would need to input each lower order item as it was presented, because they could not know which was going to be relevant. Then, on presentation of the context in the last position, subjects would need to select from among the items maintained in WM that which would be allowed to influence the response. Hence, the CL condition requires subjects to use an “output gating” strategy.

Because the successful use of an input gating strategy would also serve to reduce WM load when the context was presented first and only the letter or wingding was relevant (termed the CF, selective [CF-S] condition), we included a “global” context cue (e.g., the digit “3” in Figure 1A) that specified that both lower-level items were relevant for determining the correct response. Thus, in contrast to the “selective” context cues (e.g., the digits “1” or “2” in Figure 1A) that marked only a subset of lower-level items as relevant, a “global” context cue required subjects to utilize two items in memory, regardless of whether it was presented first (CF, global [CF-G]) or last (CL, global [CL-G]). These global trials therefore acted as a strong load-matched control condition for the CL, selective (CL-S) condition, wherein only one of the two items in memory was relevant for behavior and therefore had to be selectively output gated.

This design allowed us to test the hypothesis that cortico-striatal output gating mechanisms support selection from within WM. Specifically, by contrasting load-matched CL with “item last” conditions, we could test the prediction that cortico-striatal systems support selection from within WM. Moreover, contrasting CL with CF conditions, we can assess whether regions previously identified with second-order hierarchical control are indeed preferentially activated and coupled with the striatum during second-order control over output relative to input gating. We focus here both on regional differences in BOLD activation

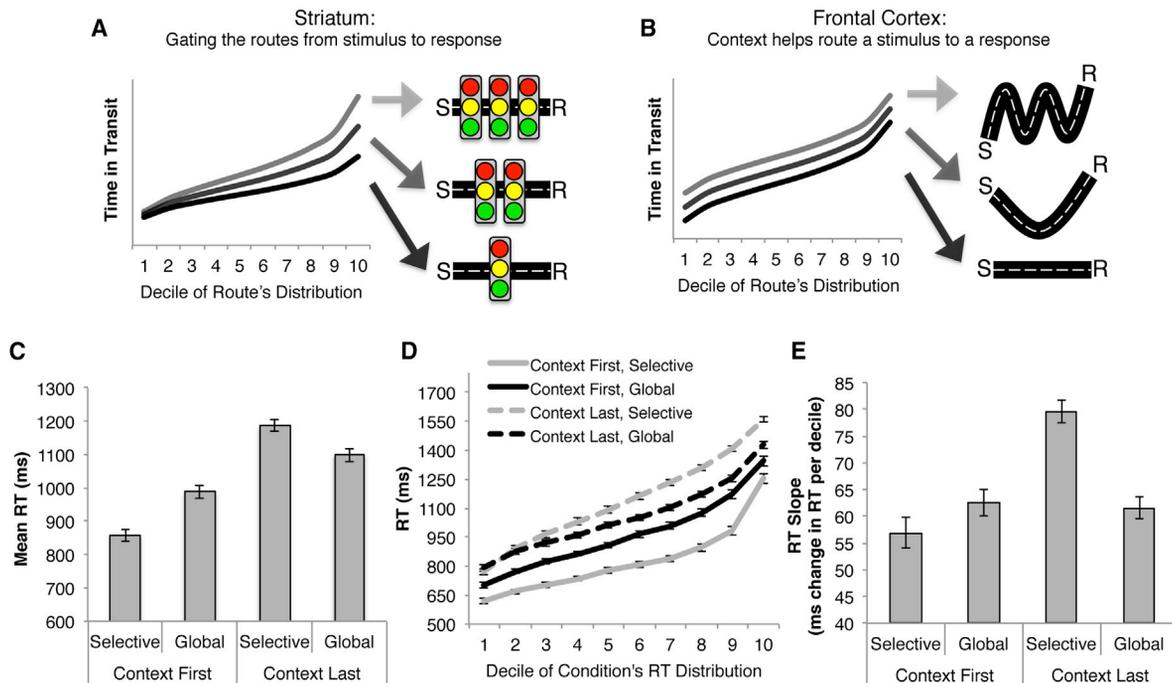


Figure 2. Behavioral Signatures of Output Gating and a Clarifying Analogy

(A) The influence of output gating on execution of a WM-dependent stimulus-response (S-R) mapping can be understood by allusion to the influence of traffic on driving along a route. For a given route (condition), the shortest trips (trials) are the least likely to have involved red lights (output gating); trips of increasing duration are more likely to have been “gated” by red lights. Thus, the rate at which trip durations (RTs) increase when rank ordered can index the prevalence of lights (output gating).

(B) By contrast, more general factors like the efficiency of the route (e.g., a WM-mediated S-R mapping) will manifest as a constant offset across ranks and be captured by mean differences.

(C) As this analogy implies, mean RT was increased in the CL condition relative to the CF condition (e.g., where the stimulus-response mapping was more indirect by virtue of requiring selection from WM). This difference was apparent even when comparing the CL conditions to the load-matched CF-G condition, although it was particularly pronounced in the CL-S condition (putatively involving selective output gating). However, these shifts in mean RT were also accompanied by changes in the shape of the RT distribution.

(D) Similar to the disproportionate influence of traffic lights on the longer trips taken along a route, the CL-S condition showed disproportionate increases in response latency across deciles.

(E) Within-participant estimates of the slope of RT as a function of decile confirm the disproportionately greater RTs observed at later portions of the RT distribution in the CL-S condition, relative to the load-matched global conditions. Error bars reflect the SEM.

See also Figure S1.

between these conditions and changes in functional connectivity across frontostriatal circuits.

Furthermore, this design allows us to test the maintenance versus gating division of labor between frontal cortex and striatum hypothesized to underlie output gating. In particular, the impact of each component to output gating should be distinguishable in separable behavioral correlates that can, in turn, be related to our brain measures. To illustrate this logic, we draw an analogy to the effects on driving duration of how direct a route is versus how frequently one encounters red lights along the way (Figure 2).

Consider each trial in our task as a trip driven along a route from stimulus to response via representations in WM. Taking a more direct route to a location reduces all transit times along that route equally, relative to less direct routes. Similarly, maintaining a good rule or “policy” in WM—e.g., one that most directly and effectively maps a stimulus to the response—will also yield more efficient performance. Thus, during CL

conditions, to the degree that strong prefrontal representations of context support more efficient responses, activation in the frontal cortex should predict speeded response times (RTs), and this speeding should be reflected by a constant shift across the entire RT distribution.

In contrast, stochastic influences like traffic will fundamentally reshape, rather than simply shift, the distribution. Consider that traffic lights can, at best, impose no delay to the fastest transit time (i.e., when all lights are green) and otherwise impose delays proportional to the total number of lights (i.e., when all lights are red). Thus, the gating of traffic by red lights will directly correlate with the rate by which transit times increase when trips are rank-ordered from fastest to slowest—that is, the slope of latency across ranks. Similarly, to the degree that the striatum relates to gating of a particular stimulus-response mapping, striatal activation and/or cortico-striatal interactions should relate to this slope estimate during conditions that put greater demands on selective output gating (e.g., CL-S).

Our results show evidence consistent with the dynamics of a corticostriatal output gating mechanism for selection from within WM. Specifically, we observed (1) that demands on output gating (whether global or selective) are associated with differential recruitment of the dorsal anterior premotor cortex (PrePMd), the cortical area most strongly identified with second-order control in prior work (Koechlin et al., 2003; Badre and D'Esposito, 2007; Badre et al., 2009, 2010; Badre and Frank, 2012); (2) that coupling between PrePMd and striatum increases specifically when output gating is required; and (3) that PrePMd, the striatum, and their functional coupling correlate with separate behavioral signatures of output gating that are potentially consistent with a maintenance versus gating division of labor within this circuit. That is, PrePMd BOLD specifically and uniquely predicts mean shifts in RT during selective output gating, whereas BOLD in caudate and its coupling with PrePMd uniquely predict the slope of the RT distribution during selective output gating.

RESULTS

Behavioral Dissociations

The behavioral results revealed that the CL conditions were associated with characteristic performance costs distinguishing them both from the CF conditions and from costs related to WM load. Figure 2B plots the effect of experimental conditions on mean RT; similar condition effects were evident in terms of error rates although performance was satisfactory on the task overall (12% errors; Figure S1 available online). Specifically, the CL conditions yielded increased RT relative to the CF conditions overall, $F(1, 21) = 250.54$, $p < 0.001$. This difference is important, as subjects could conceivably have always waited until the final stimulus to decide which lower order item to select. To the contrary, the behavioral facilitation on CF relative to CL trials suggests that subjects took advantage of the input gating strategy available to them on CF trials and that CL trial imposed a cost associated with selection from WM.

These mean RT costs were particularly evident in the CL-S condition, $F(1, 21) = 6.02$, $p = 0.02$, even when comparing this to the CF-G and CL-G conditions with equivalent WM loads (Figure 2A; t values > 4.07 , and p values > 0.001). These effects were not due to insufficient practice with each contextual token, as the differences were stable with experience (see Figures S1A–S1C). Thus, a robust performance decrement is associated with the CL-S condition relative to CF conditions and is dissociable from costs related to greater WM load.

Only partial information about the nature of this difference is conveyed by mean RT. Indeed, and as implied by gating effects in our traffic analogy, separable effects of condition were evident in the RT distribution's shape (Figure 2C; Balota and Yap, 2011; see also Ratcliff and Frank, 2012). Relative to the CF-S condition, the greater WM load in the CF-G condition was disproportionately apparent in the tail of the RT distribution. Relative to the CF-G condition, however, the load-matched CL conditions yielded dissociable effects on the shape of the distribution. Whereas the CL-G condition increased RT equally across the entire distribution relative to CF-G—akin to a mean shift due to demand on selecting information from WM—only the CL-S con-

dition yielded an additional elongation of the distribution. This elongation is hypothesized to relate to inefficiencies in selection from WM due to output gating (akin to the effect of traffic lights on the distribution of trip durations), an assumption we relate to fMRI data later.

These changes in shape were quantified as the rate at which RTs increased across deciles of the distribution (Figure 2E, "RT Slope"). The CL-S condition yielded a larger increase in RT slope than both the CL-G and the CF-G conditions, $t(21) = 6.54$, $p < 0.001$; $t(21) = 4.20$, $p < 0.001$; and both of these differences were greater than that observed between the CL-G and CF-G conditions, $t(21) = 3.24$, $p < 0.005$; $t(21) = 6.54$, $p < 0.001$. Thus, and in contrast to the combined effect of WM load on both the mean and shape of the RT distribution, the CL conditions yielded a change in the mean of the RT distribution, with an additional elongation of the distribution only in the CL-S condition. In contrast to the mean shift evident between CF-G and CL-G, the effect of CL-S on RT slope is presumably attributable to the unique, load-independent demands on selective output gating required by this condition.

Performance was comparable in the scanned version of the experiment, where we imposed a delay between the final stimulus event and the response mappings (discussed above Figure 1A). This delay was imposed to ensure that BOLD responses to stimuli appearing last were unrelated to any demands imposed by the active selection and execution of a motor response, which may involve mechanisms similar to output gating. Despite this change, accuracy remained high across all conditions (means of 90.9%–93.2%). The only reliable performance effect of condition was increased RT in the global conditions relative to both selective conditions, $F(1, 21) = 22.7$, $p < 0.001$, perhaps reflecting the increased WM load in the global conditions. See Figures S1E–S1H for behavioral results from the scanned session.

Univariate fMRI Contrasts

Widespread regions in the PFC were recruited during task performance, both on presentation of higher-level context (red regions in Figure 3) and lower-level items (Figure 3, blue regions; overlap in purple) relative to baseline, averaging across presentation orders (first versus last). A contrast of these effects revealed that a subset of lateral frontal cortex exhibited a stronger response to context information (i.e., numbers) relative to items (i.e., letters or wingdings; black outlined regions in Figure 3). The frontal peak of this response fell within the previously identified locus of second-order control, the left dorsal pre-motor cortex (Montreal Neurological Institute [MNI] coordinates: $-50, 8, 30$; Badre and D'Esposito, 2007).

To assess whether frontal cortex might be differentially involved in second-order control over output gating relative to input gating, we contrasted the BOLD responses to contexts appearing last versus first relative to the analogous contrast of items appearing last versus first. This higher-order contrast yielded increased recruitment throughout lateral frontal cortex, including the aforementioned PrePMd region as well as the adjacent inferior frontal sulcus (IFS) and dorsal premotor cortex (PMd), ruling out general influences of stimulus order or WM load (see Experimental Procedures for further explication of

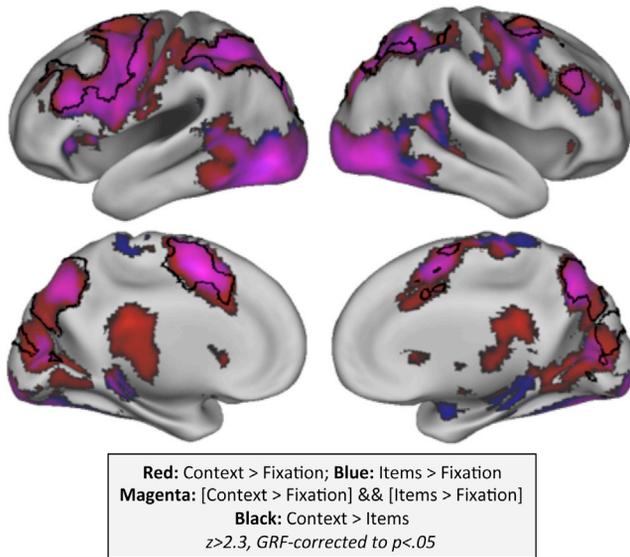


Figure 3. Univariate Recruitment to Context and Items

BOLD responses to context > fixation (red) and items > fixation (blue) regardless of position (first versus last) were largely overlapping (magenta) but significantly larger for context than for items (black outlined regions) within lateral PFC. $Z > 2.3$, corrected to $p < 0.05$ via GRF theory. The differential response to context versus items peaked in frontal cortex within the vicinity of the pre-PMd.

this logic). Also, like the simpler contrast of context versus items reported above, this differential response also peaked in frontal cortex within the vicinity of the PrePMd (MNI coordinates: $-52, 6, 38$; Figures 4A and 4B).

Given the proximity of this univariate effect to prior observations of PrePMd during hierarchical cognitive control tasks, we replicated the above analysis using an unbiased region of interest (ROI) defined from prior work (Badre and D'Esposito, 2007) and its right hemisphere homolog. A flexible model was used to characterize the hemodynamic response function (FLOBS; Woolrich et al., 2004) and estimated using FMRIB's fMRI Expert Analysis Tool (FEAT; see also Experimental Procedures for further details). The reconstructed time courses (Figure S2) showed a peak in the BOLD response at approximately 4 s poststimulus across all conditions. Average percent signal change (PSC) in the 2 s surrounding each peak was then subjected to a 2 (stimulus order: first versus last) \times 2 (stimulus type: context versus items) \times 2 (hemisphere: left versus right) repeated measures ANOVA. We found a greater BOLD response to context appearing last versus first, relative to the same contrast of items—i.e., the interaction of order and type; $F(1, 21) = 4.2$, $p = 0.05$ —with no effect of hemisphere (F values < 1 , p values > 0.3 ; Figure S2). This pattern also replicated in terms of PSC as estimated from the canonical hemodynamic response function (HRF) model, suggesting that the canonical HRF accurately captured this key aspect to the hemodynamic differences among our events.

Univariate Brain-Behavior Correlations

As CL-S was behaviorally distinguished from the other conditions on the basis of the mean and slope of its RT distribution,

with the latter being unique to the CL-S versus CL-G difference, we sought to assess whether activation in PFC and striatum related to these behavioral signatures. Of note, as these are statistics computed only at the whole participant level, we tested whether the mean and slope of RT would relate to individual differences in the recruitment of this area. Specifically, mean PSC was extracted from the PrePMd, as well as the adjacent IFS and PMd. Given the good agreement between the current foci and those identified in prior work, these ROIs were defined on the basis of this prior work (Badre and D'Esposito, 2007) as well as the homologous regions on the right hemisphere (although, we note that similar effects were observed using functionally defined ROIs).

The most robust correlation to emerge from this analysis was between mean RT in the CL-S condition and mean PSC in the left and right PrePMds during the selective contexts when presented last (i.e., CL-S events: $r = -0.46$, $p = 0.03$; $r = -0.52$, $p = 0.01$; Figures 4C and 4D).

The observed bilateral correlation was highly specific to this condition, measure, and region (as determined through semi-partial correlations; see also Experimental Procedures). (1) The relationship of PrePMd recruitment to mean RT during the CL-S condition was independent of the analogous measures in the other conditions ($r = -0.43$, $p < 0.05$; Figure 4E). (2) The correlation was also independent of the recruitment of both the more caudal PMd and the more rostral IFS in the same condition ($r = -0.55$, $p < 0.01$; and $r = -0.54$, $p < 0.01$, respectively), both of which failed to correlate with CL-S mean RT in either hemisphere (all p values > 0.5), with PMd significantly less correlated than PrePMd (Fisher's $z = 1.93$, $p = 0.05$). (3) The correlation was independent of the other behaviorally correlated neural measures we report later (partial $r = -0.58$, $p < 0.01$), which themselves failed to correlate with mean RT during the CL-S condition (p values > 0.14). Given this specificity of the link between mean RT during the CL-S condition and PrePMd recruitment to CL-S contexts, this brain-behavior correlation is unlikely to reflect generalized factors such as differences in sensorimotor abilities between individuals or differences in interference within WM caused by either WM load or simply by presenting context last.

Next, given our a priori focus on frontostriatal systems, we tested an anatomically defined ROI in the dorsal striatum (specifically, caudate; Figure S5). We reasoned that, if the caudate is involved in output gating, then activity in this region should reduce the RT cost associated with these demands. Indeed, caudate activation correlated with reductions in the CL-S RT slope, once again bilaterally (left: $r = -0.50$, $p < 0.02$; right: $r = -0.51$, $p = 0.01$; Figures 4F and 4G).

As with the correlations described earlier, the caudate's correlation with RT slope was also highly specific. (1) The relationship of RT slope in the CL-S condition to the effect of CL-S context events on caudate recruitment was independent of the corresponding brain and behavioral measures from all other conditions at a trend level ($r = -0.354$, $p < 0.11$; Figure 4H), reaching significance when bilateral caudate recruitment during CL-G was not partialled from that during CL-S ($r = -0.49$, $p < 0.02$) or when bilateral caudate recruitment was averaged across both CL conditions ($r = -0.54$, $p = 0.01$). Thus, CL-S RT slope is correlated with variance in bilateral caudate recruitment shared

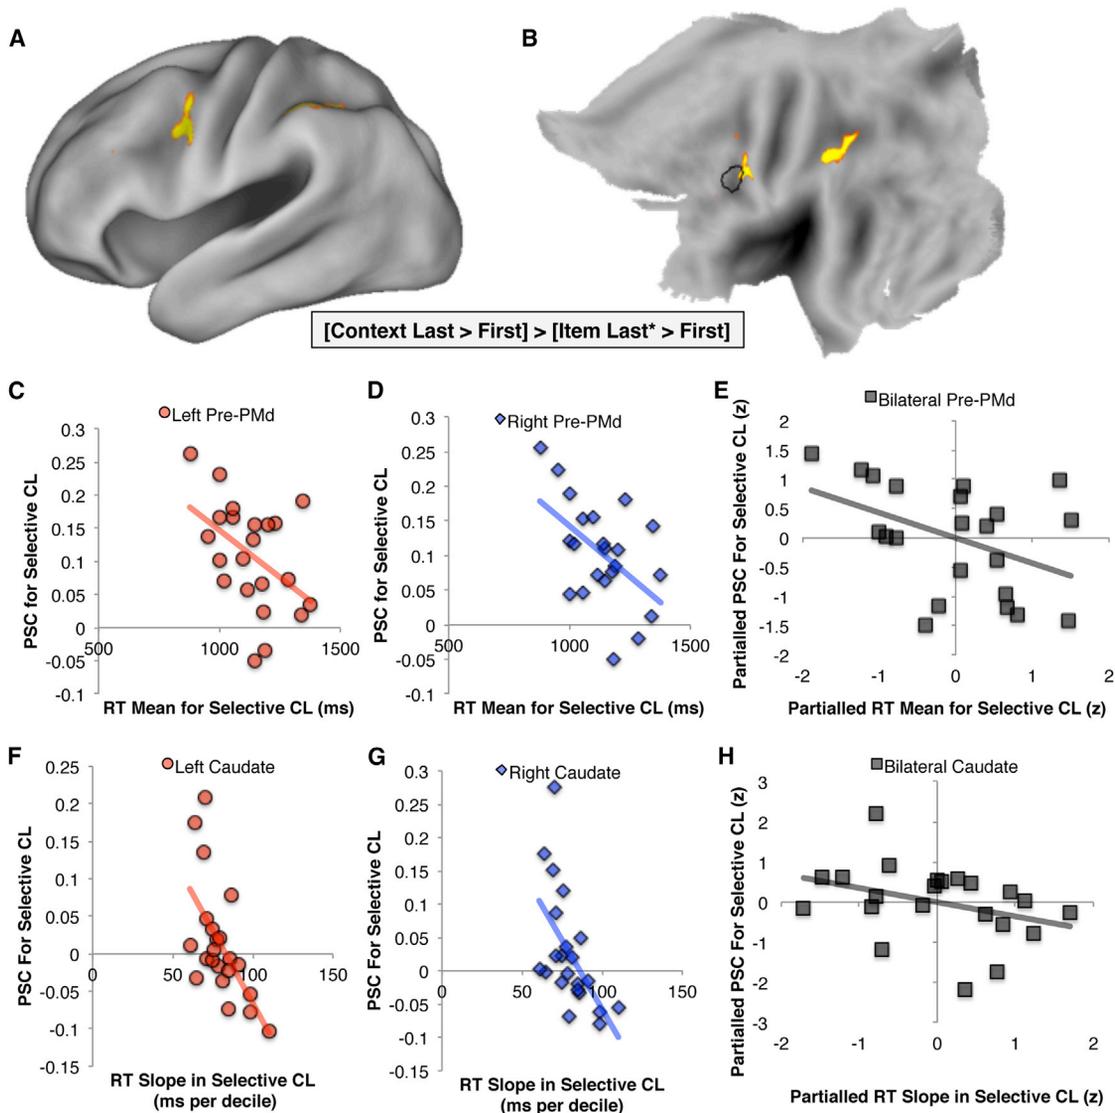


Figure 4. Univariate Brain-Behavior Correlations

(A) The differential BOLD response of frontal cortex to context appearing last versus first, as compared to load- and order-matched items (i.e., items appearing first or last* in the global context), peaked within the vicinity of the PrePMd (here shown following the initial voxelwise threshold of $z > 2.3$, and GRF correction to $p < 0.05$, then subject to a further increase in the voxelwise threshold equivalent to $p < 0.05 \times 10^{-7}$).

(B–D) In (B), this peak fell partially within an ROI centered on the PrePMd observed in prior work (black outlined region). PSC extracted from the outlined region (C) and the equivalent location on the right hemisphere (D) during the CL-S condition negatively correlated with mean RT in CL-S.

(E) This correlation was independent of PSC observed in the other conditions and of mean RT observed in these other conditions ($r = -0.43$, $p < 0.05$).

(F–H) Univariate recruitment of left (F) and right (G) caudate was negatively correlated with the slope of RT during the CL-S condition, and this was independent of caudate recruitment and RT slope in all cases except for caudate recruitment in CL-G, controlling for which the correlation was marginally significant (H).

*The “item last” trials used in this contrast are limited to those in the global condition, where WM load is matched to that when context appears last (see text). See also Figure S2.

across conditions where context appears last and, thus, demands selection from WM. (2) The relationship of RT slope in the CL-S condition to the effect of CL-S context events on caudate recruitment was also independent of the other behaviorally correlated neural measures we report ($r = -0.49$, $p < 0.03$). Thus, these observations again argue against an interpretation in terms of differences in general sensorimotor abilities or WM interference across subjects or conditions.

Corticostriatal Interactions and Correlations with Behavior

To test whether the CL conditions might drive differential corticostriatal interactions than the CF conditions, we conducted a whole-brain search for voxels differentially correlated with our caudate ROI's time series during the CL condition compared to the CF condition (i.e., a psychophysiological interaction; see Experimental Procedures). Caudate increased

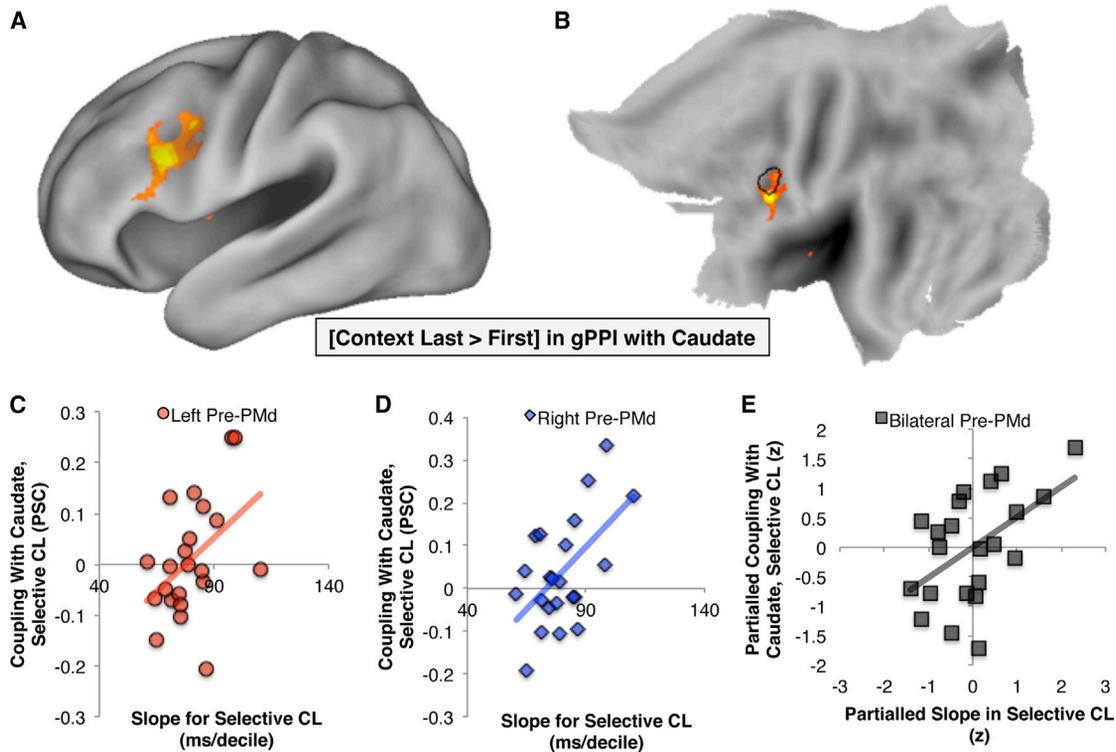


Figure 5. Coupling-Behavior Correlations

(A) Left lateral PFC increased its coupling with caudate during the CL conditions relative to CF conditions, and this again was within the vicinity of the PrePMd identified in prior work; indicated by black outline in (B–D). The PSC associated with this change in coupling was extracted from this a priori ROI (C) and its right hemisphere counterpart (D) for the CL-S condition. These measures robustly correlated with the slope of the RT distribution (see Figure 1D). (E) This correlation was independent of PrePMd's coupling with caudate in the other conditions and of the slope in RT observed in other conditions, as demonstrated through partial correlation.

its coupling with the left lateral PFC, including a subset of the PrePMd, when context appeared last relative to first (Figures 5A and 5B). These results indicate that corticostriatal interactions are pronounced in the CL condition relative to the CF condition (when putative demands on output gating are high).

We next assessed whether these corticostriatal interactions might predict performance during the CL-S condition. Coupling of the caudate with our a priori PrePMd ROI during the context events of the CL-S condition predicted the slope of the RT distribution in this same condition, bilaterally ($r = 0.45$, $p < 0.05$, and $r = 0.56$, $p < 0.01$, for left and right PrePMds, respectively; Figures 5C and 5D). Once again, this correlation was highly specific. It remained significant when controlling for the slope of the RT distribution in other experimental conditions, PrePMd-caudate coupling during context in other experimental conditions, and PMd-caudate coupling during context in the CL-S condition (r values ≥ 0.45 , p values < 0.05 ; Figure 5E). In addition, this correlation was significant even when partialing univariate PrePMd and caudate recruitment to the same events ($r = 0.54$, $p < 0.02$). Together, these effects once again rule out general factors like sensorimotor abilities or WM interference, e.g., due to differences in WM load or the presentation of CL.

DISCUSSION

We report that selection from within WM is supported by corticostriatal dynamics analogous to those involved in selective control over the input to WM, with lateral frontal and striatal components of this circuit relating to separable individual differences. As detailed below, our results support the output gating hypothesis by suggesting that the lateral frontal cortex may maintain contextual representations that influence the striatum to gate which subset of WM representations can influence behavior.

Our task was designed to isolate the effects of selective output gating, independent of those due to input gating or overall WM maintenance demands. Our behavioral results confirmed that these manipulations were effective. RT was decreased when a selective context was presented first (CF-S) relative to when context was presented last (CL-G and CL-S), suggesting that subjects took advantage of the input gating strategy available to them on CF-S trials. Moreover, CL-S trials were associated with additional behavioral costs compared with CL-G trials. As WM load is equated between these conditions, these additional costs on CL-S trials are directly attributable to the demand to select from within WM.

Following from the above logic, comparing CL versus CF conditions targeted selection from WM independent of WM

load. This comparison yielded two key results consistent with the output gating hypothesis: (1) Activation was greater for context presented last relative to first in lateral frontal and parietal cortices, peaking in PrePMd. (2) Connectivity between the dorsal striatum and lateral frontal cortex increased selectively during context presented last relative to first. This association of output gating with the CL conditions, and with the CL-S condition in particular, is further supported by the brain-behavior correlations unique to the CL-S condition—a condition whose behavioral costs uniquely predicted its neural correlates (as assessed in a separate session) within the pre-PMd, the dorsal striatum, and the coupling between these regions. Together, these observations are consistent with the hypothesis that frontostriatal gating dynamics support selection from WM, like those previously documented for input gating. Indeed, contexts used to support input gating (i.e., contexts presented first) recruited overlapping areas with those used to support output gating (i.e., contexts presented last), albeit with a difference in degree (see [Figure 3](#)). This overlap is consistent with the hypothesis that colocal cortico-striatal circuits support both types of gating and perhaps with the notion of spatially interdigitated “stripes” supporting them (e.g., [Hazy et al., 2007](#)).

A role for output gating during selection from WM could extend from the more well-established role of these mechanisms in action selection. Specifically, corticothalamic loops are separably modulated by “Go” and “NoGo” striatal pathways, with activity in Go pathways acting to disinhibit each loop and activity in NoGo pathways acting to inhibit that same loop ([Mink, 1996](#); [Graybiel, 1998](#); [Frank et al., 2001](#); [Hazy et al., 2007](#)). Classically, this opponent process is thought to modulate excitability in effector-specific thalamo-M1 loops and thereby underlie the role of Go and NoGo pathways in the execution and inhibition of specific motor responses (respectively). In addition, premotor representations are thought to influence which of several candidate motor responses constitutes the effective output of M1 (i.e., which representation is output gated) by biasing the Go and NoGo pathways modulating those thalamo-M1 loops. Here, by what might fundamentally be a similar mechanism, pre-premotor representations (e.g., representations of context in PrePMd) may influence which of several candidate premotor representations (e.g., candidate WM items; here, letters and windings) constitutes the most effective output of premotor cortex by biasing the Go and NoGo pathways modulating those thalamo-PMd loops. More concretely, a top-down contextual signal from PrePMd to Go cells might amplify the representation of letter over winding, so that the maintained letter becomes the most effective output of the PMd (see [Figure S3](#) for evidence of PMd involvement in the current task). In this way, our results support the output gating mechanisms of many models ([Hazy et al., 2007](#); [Kriete and Noelle, 2011](#); [Chatham et al., 2011](#); [Collins and Frank, 2013](#); [Kriete et al., 2013](#); [Huang et al., 2013](#)), including models of hierarchical cognitive control in particular ([Frank and Badre, 2012](#)).

Given the analogy of the output gating WM mechanism tested here to the circuit responsible for motor selection, it is important to note that the present imaging results cannot be accounted for by motor preparation. This crucial alternative was addressed in two ways in the fMRI design. First, the response phase was

separated in time from the presentation of the final stimulus event, including during CL conditions. This design feature ensured that we measured changes in BOLD related to processing of the contextual cue in CL conditions rather than any additional demands related to selection and execution of a motor response. Second, the response mappings themselves changed unpredictably on each trial. Thus, it was not possible for the participant to prepare a response prior to the delayed response cue and so confound our effects. For these reasons, the cortico-striatal correlates identified here are entirely attributable to the demands on selection from within WM rather than motor selection.

Although we observed increased frontostriatal coupling during CL conditions, we did not find greater univariate activation in the caudate in this condition relative to controls. Notably, prior work has shown a high degree of variability in the observation of striatal recruitment across different WM manipulations, including maintenance (e.g., [Postle and D’Esposito, 1999, 2003](#); [Marklund et al., 2007a, 2007b](#); [Narayanan et al., 2005](#); [Olsen et al., 2013](#)), encoding/updating (e.g., [Narayanan et al., 2005](#); [McNab and Klingberg, 2008](#); [Dahlin et al., 2008](#); [Nee and Brown, 2013](#)), manipulation (e.g., [Lewis et al., 2004](#); [Dodds et al., 2009](#)), and selection ([Tamber-Rosenau et al., 2011](#)). This inconsistency might reflect a dependence of the striatal BOLD response on the relative balance of Go and NoGo striatal cells recruited by each task. Perhaps parametric effects coding for both the direction and magnitude of effects, like those reported here with respect to RT slope—or by prediction errors derived from reinforcement learning (RL) models in much recent work—may be better suited to locate clear striatal BOLD effects than simple aggregate mean contrasts across conditions. However, simple RL dynamics do not seem to account for the effects we report here (cf. [Chatham and Badre, 2013](#)).

A second key observation from this study is that, although frontal cortex and striatum participate in a common circuit for output gating, they correlated with separable behaviors during selection from within WM. Specifically, analysis of the RT distribution across experimental conditions revealed a mean shift in RT when context was presented last, but WM load was matched (e.g., CL-G > CF-G), reflecting a general cost of selection from WM. This effect was relatively constant across deciles of the distribution, as expected of a general factor influencing the efficiency of stimulus-response mappings. However, this cost was dissociable from a change in the shape of the RT distribution quantified by RT slope, which accompanied demands on selective versus global output gating (CL-S condition). In brain-behavior correlations, individual differences in the activation of PrePMd were related to mean RT during the condition involving selective output gating (CL-S). By contrast, individual differences in striatal activation and striatal coupling with PFC correlated with the shape of the RT distribution during the CL-S condition, as expected of mechanisms involved in gating.

This dissociation is potentially consistent with a frontostriatal division of labor between maintenance and gating ([Frank et al., 2001](#); [Gruber et al., 2006](#); [Frank and O’Reilly, 2006](#); [Cools et al., 2007](#)). More specifically, lateral frontal cortex (PrePMd in this case) is hypothesized to maintain the context that influences what lower order information should be gated by the striatum.

Consequently, PrePMd sets the appropriate policy and so determines how rapidly selection can begin on average, thereby affecting the whole distribution equivalently (rather than its shape). Also consistent with this interpretation, PrePMd was generally more activated by context stimuli (i.e., numbers) than load- and order-matched item stimuli (i.e., letters and wingdings). This context preference was enhanced when the context occurred last and could therefore drive output gating during the selection from WM. Such effects can be broadly understood as reflecting the efficiency of a context-mediated mapping between stimulus and response, analogous to the directness of a route.

Of note, outflowing information in an output-gated system must wait until the gate is opened before it can influence behavior. Consequently, the selection process itself exerts further “waiting times” on responses beyond those associated with the policy’s efficiency itself. These wait times can only make RT longer and so are more likely to contribute to longer than shorter RTs. As in traffic, such wait times will elongate the distribution in a way that can be quantified by RT slope (as well as alternative nonparametric measures of spread; see [Table S2](#)). Although such changes in distribution shape have not been previously demonstrated for output gating or selection from WM specifically, simulation studies have demonstrated similar effects on RT arising from input gating failures ([Reynolds et al., 2006](#)). Thus, the correlation of individual differences in striatal activation and frontostriatal connectivity with the RT slope could reflect the efficiency with which a given policy (established by PFC) is implemented by striatally mediated gating mechanisms (see also [Figure S4](#) for other correlates of RT slope, including trial-to-trial variability in striatal BOLD). Although this view of the brain-behavior dissociation is consistent with current output gating models of selection from WM, further work is required to confirm this conclusion and distinguish it from alternatives.

Future models of output gating may be notably constrained by our results. The nonconstant magnitude of the CL-S versus CL-G difference across deciles of each distribution implies a greater influence of gating in CL-S in the same way that greater increases in transit time by rank order imply a greater number of traffic lights along a route. However, this effect is counterintuitive from a classical account of WM selection: more items must be selected from WM in CL-G than CL-S. One possibility is incongruity: all WM representations should influence behavior in the CL-G condition (a congruent “open gates” policy across items in WM), whereas in CL-S, only one item should influence behavior (an incongruent “one open, one closed” gating policy). Such congruency effects would mirror those seen in the motor domain. Alternatively, if an output gate is closed for just one item via striatal NoGo mechanisms, perhaps this NoGo pathway exerts behavioral slowing during selection from WM, as it does in action selection (e.g., [Kravitz et al., 2010](#)). These possibilities are not mutually exclusive, but both imply analogous roles for output gating during selection from WM and action selection.

One outstanding issue is how precisely output gating reshapes RT distributions. Here, nonparametric methods quantify changes in shape. While waiting times can be modeled with gamma distributions (e.g., [Chatham et al., 2012](#)), those due to

output gating are unlikely to arise from the gamma’s most widely used special case: the exponential distribution ([Balota and Yap, 2011](#)). The exponential may be poorly suited to model output gating, which may involve information accrual, contrasting with the “memoryless” nature of the exponential, and a modal influence on RT ([Wiecki and Frank, 2013](#)), contrasting with the mode of zero on the exponential. The centrality of parametric assumptions for modeling the latent component of RT ([Jones and Dzhafarov, 2013](#)) further motivates such caution.

Although here we show that selection from WM may rely on corticostriatal output gating mechanisms, such mechanisms may play a more general role in selection of internal representations. For example, striatal involvement in object switching has been reported ([Tamber-Rosenau et al., 2011](#)), but it is unknown whether frontostriatal coupling also accompanies such shifts. It is also unknown whether either striatal recruitment or coupling correlates with either the mean or the shape of the distribution of object switch costs. Likewise, striatal recruitment during selective retrieval from long-term memory ([Scimeca and Badre, 2012](#)) could reflect output gating of hippocampus itself, of cortical regions responsible providing retrieval cues to hippocampus, or of cortical regions responsible for postretrieval selection (e.g., [Badre et al., 2005](#); [Badre and Wagner, 2007](#); [Snyder et al., 2010, 2011](#)). This array of predictions, if confirmed, would suggest substantial generality to the current account of corticostriatal output gating interactions for selection in WM.

Finally, we demonstrate these effects in the domain of hierarchical control, where output gating confers computational advantages ([Frank and Badre, 2012](#); [Kriete and Noelle, 2011](#)). Together, our results suggest a critical role for more rostral PFC—and, most consistently, PrePMd—for output gating in second-order hierarchical control. In this way, both current and prior results associate PrePMd with second-order control across a variety of hierarchical control tasks ([Koechlin et al., 2003](#); [Badre and D’Esposito, 2007](#); [Badre et al., 2009, 2010](#); [Badre and Frank, 2012](#)). However, our effects are not perfectly unique to PrePMd; posterior parietal cortex is also more strongly recruited for CL than for CF conditions (e.g., [Figure 4B](#)). Similarly, while effects in PrePMd are consistently independent of the more caudal PMd and the most rostral RLPFC, PrePMd only differed from the adjacent IFS in its correlation with mean RT. Thus, while our results strongly support a role for PrePMd in second-order control over output gating, they should not be taken to show that PrePMd alone is responsible for performing the second order rule.

This observation informs active debates on frontal organization, where two recent studies report unexpectedly rostral frontal BOLD for second-order control tasks ([Reynolds et al., 2012](#); [Crittenden and Duncan, 2012](#)). We note that this is also true both in this study ([Figure 3](#)) and in prior studies from our lab ([Badre et al., 2010](#)). In contrast, selective effects along the rostro-caudal axis emerged only from parametric manipulations of the number of mappings (e.g., stimulus-response mappings in first-order control) resolved at each level of control ([Koechlin et al., 2003](#); [Badre and D’Esposito, 2007](#)). Indeed, [Crittenden and Duncan \(2012\)](#) replicated this effect, reporting selectively increased recruitment in caudal PMd (as compared to other ROIs from our prior work) when lower-order response competition was manipulated most

similarly to that prior work. Thus, the question for theory becomes why a parametric manipulation of choice at each level of control might expose these rostrocaudal differences, whereas certain other manipulations yield less consistent and functionally selective results. Previous attempts to reconcile these results have focused on demands on sustained context maintenance and differential encoding versus retrieval demands, although our results do not clearly support these accounts either (Figure S3). One possibility is that parametric manipulations better control demands on the gating processes that govern directed influences across frontal subregions—a possibility that motivates tight control over gating in future studies of rostrocaudal frontal gradients.

In support of this possibility, Nee and Brown (2013) reported results more compatible with prior accounts of rostrocaudal gradients when gating demands were more closely controlled. In this study (and also in Nee and Brown, 2012), rostral prefrontal areas (albeit in regions rostral to PrePMd during a third-order control task) showed stronger recruitment by demands to input gate a higher-level context (versus retaining the prior higher-level context), whereas the more caudal PMd was more strongly recruited by the demand to input gate a lower-level context (versus retaining the prior lower-level context). Nee and Brown (2013) failed to observe corticostriatal connectivity during input gating of the lower-level context; we likewise failed to obtain significant increases in corticostriatal coupling during the input gating of lower-level items. Instead, corticostriatal coupling increased during the output gating of lower-level items. This difference is consistent with our conclusion that hierarchical corticostriatal interactions may be particularly crucial during output gating.

EXPERIMENTAL PROCEDURES

Subjects

Twenty-two right-handed adults (aged 18–35; 8 female, 14 male) with normal or corrected-to-normal vision completed the experiment. All spoke English natively, were screened for both neurological medications/conditions and MRI contraindications, and provided informed consent in accordance with the Research Protections Office at Brown University. All subjects first performed a behavioral version of the task and then a modified version of the task while undergoing fMRI.

Behavioral Procedure

Three tokens were presented on each trial. Each token was of a distinct type: digits, letters, and symbols. The tokens appeared in random order, but digits acted as a context for a hierarchical rule. Digits 1 and 2 specified “Selective” contexts. For example, 1 might specify that the symbol observed on that trial would be relevant for a response, and 2 would specify that the letter from that trial would be relevant. The 3 always acted as a “Global” context, indicating that both the letter and the symbol would be relevant for the participant’s response. These three contexts were equiprobable, as was the likelihood of a digit at stimulus position in a trial. These trial characteristics were randomized in the behavioral experiment but optimized for power in the fMRI experiment.

Response mappings were presented at the bottom left and right of the screen as the final stimulus event for each trial. Only one mapping contained the relevant item or items for that trial, as determined by the hierarchical rule. Subjects indicated which side (left versus right) contained the relevant item(s) using a button press (Figure 1A). For reasons described below, during the behavioral experiment, the mappings were presented simultaneously with the last token, whereas during fMRI scanning, the response mappings were separated in time from the last item by a jittered delay.

Half of all trials included a response lure, where the incorrect response option contained one item that had been seen on that trial. Subjects were instructed that lure trials would occur and that every trial would therefore require them to proactively attend to all items specified as relevant by the context (or the lack thereof). Analysis of lure trials confirmed that subjects maintained both items, since lure accuracy was well above that expected if subjects randomly attended to only one item, both in the behavioral (CF-G: 93.6%, $t(21) = 10.85$, $p < 0.001$; CL-G: 90.9%, $t(21) = 6.26$, $p < 0.001$) and scanned experiments (CF-G: 93.2%, $t(21) = 8.44$, $p < 0.001$; CL-G: 91.4%, $t(21) = 8.66$, $p < 0.001$).

Our task’s logic is that context can drive input gating only when it appears first (CF); by contrast, context can only be used to drive output gating when it appears last (CL). One complication is that when context appears last, subjects will experience a larger WM load, relative to when a Selective context appears first. The inclusion of the global context remedies this problem. Regardless of whether the global context appears first or last, all items from that trial are behaviorally relevant. These global trials were thus matched in terms of WM load with the CL-S condition.

In the fMRI version of the task, between-position contrasts were also used to match load. For example, the BOLD responses to CL were compared to the BOLD responses to items appearing last in the global context. In both cases, two items from the trial are being maintained; these “global item last” trials thus constitute the appropriate load- and order-matched control for CL.

Following detailed instructions (see section S1 in the Supplemental Experimental Procedures), subjects underwent training and practice, followed by 180 trials of behavioral testing, and finally scanning.

MRI Procedure

Whole-brain imaging was performed on a Siemens 3T TIM Trio MRI system with a 32-channel head coil. A high-resolution T1 multi-echo MPRAGE was collected from each participant. Four functional runs each consisting of 303 volumes used a fat-saturated gradient-echo echo-planar sequence (repetition time = 2 s, echo time = 28 ms, flip angle = 90°, 38 interleaved axial slices, 192 mm field of view with voxel size of 3 mm³). Head motion was restricted with padding, and visual stimuli were rear-projected and viewed with a mirror attached to the head coil. Subjects responded using an MRI-compatible button box.

Data Processing

Behavioral Data

The first ten trials of the behavioral experiment were excluded from analysis, as were incorrect trials and trials with RTs ≥ 5 SD from a participant’s mean RT.

Only results from one of two qualitatively distinct subtypes of the CF-S condition are presented here. The presented subtype consists of those trials where the contextually irrelevant item appears as the second event in the trial (the other subtype consists of those where the contextually irrelevant item appears last). This matches all conditions for the demand to attend to the central portion of the final stimulus display.

Imaging Data

Data were processed using a combination of Statistical Parametric Mapping and FMRIB’s Software Library (FSL). First, SPM8 tools (artglobal and tsdiffana) were used for artifact detection, and slice timing correction was then performed. The first six volumes of each run were discarded to allow the scanner to reach steady state. Data were motion corrected using rigid transformations in MCFLIRT to the middle acquisition of each run. Runs with movement of more than 2 mm were excluded from analysis ($n = 1$). Grand-mean intensity normalization of the entire four-dimensional data set was performed with a single multiplicative factor, and the data were subjected to a temporal highpass filter (Gaussian-weighted least-squares straight line fitting, with $\sigma = 32.5$ s), and the data were smoothed at 8 mm full-width at half-maximum. The middle acquisition of each run was then registered to each participant’s brain-extracted MPRAGE using a linear 7DOF transform, and the MPRAGE was registered to the MNI standard brain using a linear 12DOF transform.

Statistical Analysis

Our generalized linear model (GLM) was estimated using FEAT, version 5.98 (FSL, <http://www.fmrib.ox.ac.uk/fs/>), on the basis of explanatory variables

(EVs) coding for the following event types: Digit 1 or 2 (i.e., selective contexts) or 3 (i.e., the global context) appearing in the first position, middle position, or last position; and items appearing in the first or middle position. Separate EVs were used to capture items appearing in the last position under the selective context and global contexts. The duration of each event in all these EVs was set to the actual stimulus duration in the scanner (500 ms). Separate EVs were used to capture responses; the duration of each event in these EVs was set to the observed RT on each trial. Additional EVs of no interest were also included in the GLM, including those corresponding to incorrect trials and to 6 degrees of freedom of motion. All EVs except those corresponding to motion were convolved with a standard HRF, high-pass filtered in the same way as the functional data, and then used as regressors (including temporal derivatives) in the GLM.

The following linear contrasts were constructed: the general effect of context appearing first, context appearing last, and the difference between these; the analogous effects of items appearing first, items appearing last (in the global context, as noted above), and the difference between these. A higher order contrast compared the effect of order on context (first versus last) with the effect of order on items that were matched for WM load (items first versus items in the CL-G condition, as discussed above). Alternative ways of addressing potential WM load differences between context and other items in the last position (e.g., involving a parametric WM load regressor) yielded qualitatively identical effects. The location and extent of clusters of activation for critical contrasts in this design can be found in [Table S1](#).

For the generalized psychophysiological interaction (PPI) analyses, the anatomically defined caudate of the Harvard-Oxford subcortical atlas was warped to subject space, thresholded at 50% probability, and then used as a mask for extracting the time series of this region from the filtered functional data. These values were used as a regressor in the design matrix without convolution or temporal derivatives and allowed to interact with each of the other regressors estimated above. This procedure effectively removes variance due to other experimental factors and demonstrably improves PPI model fit ([McLaren et al., 2012](#)). Interaction regressors for contexts appearing last and first were then subjected to a linear contrast (contexts appearing last > contexts appearing first). Qualitatively similar increases in coupling during CL relative to CF are observed using the non-generalized form of PPI, although such “nongeneralized” PPIs can suffer from unstable estimation ([McLaren et al., 2012](#)). Such consistency suggests robustness to this key result.

For both the univariate and the generalized PPI analyses, a second-level fixed-effects analysis combined the results of these first-level models across runs, and a third-level mixed-effects analysis combined these results across subjects (using FSL's local analysis of mixed effects; FLAME1). For all whole-brain analyses, the resulting voxelwise *z* statistics were initially thresholded at $z = 2.3$ and cluster corrected to $p < 0.05$ using Gaussian random field (GRF) theory.

To assess brain-behavior correlations, peak PSC from a 10-mm-radius sphere centered on the IFS, PrePMd, and PMd ROIs of [Badre and D'Esposito \(2007\)](#); MNI coordinates: $-50, 24, 24$; $-36, 8, 34$; and $-30, -12, 66$, respectively) and their right hemisphere homologs were extracted using the featquery tool from the conditions of interest and correlated with the behavioral measures. These ROIs are illustrated in [Figure S5](#). For partial correlations, PSC was first averaged across the ROIs of each hemisphere. Average PSC from the other conditions was then used to predict PSC during the CL-S condition via linear regression; these standardized residuals were then correlated with the standardized residuals of the analogous regression predicting behavior in the CL-S condition compared to that in the other conditions. Such semi-partial correlations reveal whether there is a relationship between the PrePMd PSC that is unique to the CL-S condition and the behavior that is unique to the CL-S condition (see also section S1 in the [Supplemental Experimental Procedures](#)).

Sustained BOLD following contexts was modeled with regressors beginning at the onset of the digits lasting until the response. After convolution with the HRF and its derivative, betas for these regressors were estimated by a flexible model of event-related transients ([Woolrich et al., 2004](#)). Similar results were obtained if sustained regressors are estimated with canonical HRFs.

All *p* values are two-tailed, and error bars indicate SEM. Correlations are uncorrected for attenuation due to reliability or restricted range.

SUPPLEMENTAL INFORMATION

Supplemental Information includes Supplemental Experimental Procedures, five figures, and two tables and can be found with this article online at <http://dx.doi.org/10.1016/j.neuron.2014.01.002>.

ACKNOWLEDGMENTS

We thank the Badre and Frank Labs for helpful comments and suggestions. This study was supported by awards from the National Institute of Neurological Disease and Stroke (R01 NS065046), the Alfred P. Sloan Foundation, and the James S. McDonnell Foundation.

Accepted: December 17, 2013

Published: February 19, 2014

REFERENCES

- Badre, D., and D'Esposito, M. (2007). Functional magnetic resonance imaging evidence for a hierarchical organization of the prefrontal cortex. *J. Cognit. Neurosci.* **19**, 2082–2099.
- Badre, D., and Wagner, A.D. (2007). Left ventrolateral prefrontal cortex and the cognitive control of memory. *Neuropsychologia* **45**, 2883–2901.
- Badre, D., and Frank, M.J. (2012). Mechanisms of hierarchical reinforcement learning in cortico-striatal circuits 2: evidence from fMRI. *Cereb. Cortex* **22**, 527–536.
- Badre, D., Poldrack, R.A., Paré-Blagoev, E.J., Insler, R.Z., and Wagner, A.D. (2005). Dissociable controlled retrieval and generalized selection mechanisms in ventrolateral prefrontal cortex. *Neuron* **47**, 907–918.
- Badre, D., Hoffman, J., Cooney, J.W., and D'Esposito, M. (2009). Hierarchical cognitive control deficits following damage to the human frontal lobe. *Nat. Neurosci.* **12**, 515–522.
- Badre, D., Kayser, A.S., and D'Esposito, M. (2010). Frontal cortex and the discovery of abstract action rules. *Neuron* **66**, 315–326.
- Balota, D.A., and Yap, M.J. (2011). Moving beyond the mean in studies of mental chronometry: the power of response time distributional analyses. *Curr. Dir. Psychol. Sci.* **20**, 160–166.
- Braver, T.S., and Cohen, J.D. (2000). On the control of control: The role of dopamine in regulating prefrontal function and working memory. In *Attention and Performance XVIII: Control of Cognitive Processes*, S. Monsell and J. Driver, eds. (Cambridge: MIT Press), pp. 713–737.
- Brown, J.W., Bullock, D., and Grossberg, S. (2004). How laminar frontal cortex and basal ganglia circuits interact to control planned and reactive saccades. *Neural Netw.* **17**, 471–510.
- Bunge, S.A., Hazeltine, E., Scanlon, M.D., Rosen, A.C., and Gabrieli, J.D.E. (2002). Dissociable contributions of prefrontal and parietal cortices to response selection. *Neuroimage* **17**, 1562–1571.
- Chatham, C.H., and Badre, D. (2013). Working memory management and predicted utility. *Front. Behav. Neurosci.* **7**, 83.
- Chatham, C.H., Herd, S.A., Brant, A.M., Hazy, T.E., Miyake, A., O'Reilly, R.C., and Friedman, N.P. (2011). From an executive network to executive control: a computational model of the n-back task. *J. Cogn. Neurosci.* **23**, 3598–3619.
- Chatham, C.H., Claus, E.D., Kim, A., Curran, T., Banich, M.T., and Munakata, Y. (2012). Cognitive control reflects context monitoring, not motoric stopping, in response inhibition. *PLoS ONE* **7**, e31546.
- Collins, A.G.E., and Frank, M.J. (2013). Cognitive control over learning: creating, clustering, and generalizing task-set structure. *Psychol. Rev.* **120**, 190–229.
- Cools, R., Sheridan, M., Jacobs, E., and D'Esposito, M. (2007). Impulsive personality predicts dopamine-dependent changes in frontostriatal activity during component processes of working memory. *J. Neurosci.* **27**, 5506–5514.

- Cools, R., Miyakawa, A., Sheridan, M., and D'Esposito, M. (2010). Enhanced frontal function in Parkinson's disease. *Brain* 133, 225–233.
- Crittenden, B.M., and Duncan, J. (2012). Task difficulty manipulation reveals multiple demand activity but no frontal lobe hierarchy. *Cereb. Cortex*. Published online November 6, 2012. <http://dx.doi.org/10.1093/cercor/bhs333>.
- Dahlin, E., Neely, A.S., Larsson, A., Bäckman, L., and Nyberg, L. (2008). Transfer of learning after updating training mediated by the striatum. *Science* 320, 1510–1512.
- Desimone, R., and Duncan, J. (1995). Neural mechanisms of selective visual attention. *Annu. Rev. Neurosci.* 18, 193–222.
- Dodds, C.M., Clark, L., Dove, A., Regenthal, R., Baumann, F., Bullmore, E., Robbins, T.W., and Müller, U. (2009). The dopamine D2 receptor antagonist sulpiride modulates striatal BOLD signal during the manipulation of information in working memory. *Psychopharmacology (Berl.)* 207, 35–45.
- Eliasmith, C., Stewart, T.C., Choo, X., Bekolay, T., DeWolf, T., Tang, Y., and Rasmussen, D. (2012). A large-scale model of the functioning brain. *Science* 338, 1202–1205.
- Frank, M.J., and O'Reilly, R.C. (2006). A mechanistic account of striatal dopamine function in human cognition: psychopharmacological studies with cabergoline and haloperidol. *Behav. Neurosci.* 120, 497–517.
- Frank, M.J., and Badre, D. (2012). Mechanisms of hierarchical reinforcement learning in cortico-striatal circuits 1: computational analysis. *Cereb. Cortex* 22, 509–526.
- Frank, M.J., Loughry, B., and O'Reilly, R.C. (2001). Interactions between frontal cortex and basal ganglia in working memory: A computational model. *Cogn. Affect. Behav. Neurosci.* 1, 137–160.
- Graybiel, A.M. (1998). The basal ganglia and chunking of action repertoires. *Neurobiol. Learn. Mem.* 70, 119–136.
- Gruber, A.J., Dayan, P., Gutkin, B.S., and Solla, S.A. (2006). Dopamine modulation in the basal ganglia locks the gate to working memory. *J. Comput. Neurosci.* 20, 153–166.
- Gurney, K., Prescott, T.J., and Redgrave, P. (2001). A computational model of action selection in the basal ganglia. I. A new functional anatomy. *Biol. Cybern.* 84, 401–410.
- Hazy, T.E., Frank, M.J., and O'Reilly, R.C. (2007). Towards an executive without a homunculus: computational models of the prefrontal cortex/basal ganglia system. *Philos. Trans. R. Soc. Lond. B Biol. Sci.* 362, 1601–1613.
- Hester, R., D'Esposito, M., Cole, M.W., and Garavan, H. (2007). Neural mechanisms for response selection: comparing selection of responses and items from working memory. *Neuroimage* 34, 446–454.
- Hochreiter, S., and Schmidhuber, J. (1997). Long short-term memory. *Neural Comput.* 9, 1735–1780.
- Huang, T.R., Hazy, T.E., Herd, S.A., and O'Reilly, R.C. (2013). Assembling old tricks for new tasks: a neural model of instructional learning and control. *J. Cogn. Neurosci.* 25, 843–851.
- Jones, M., and Dzhafarov, E.N. (2013). Unfalsifiability and mutual translatability of major modeling schemes for choice reaction time. *Psychol. Rev.* Published online September 30, 2013. <http://dx.doi.org/10.1037/a0034190>.
- Koechlin, E., Ody, C., and Kouneiher, F. (2003). The architecture of cognitive control in the human prefrontal cortex. *Science* 302, 1181–1185.
- Kravitz, A.V., Freeze, B.S., Parker, P.R., Kay, K., Thwin, M.T., Deisseroth, K., and Kreitzer, A.C. (2010). Regulation of parkinsonian motor behaviours by optogenetic control of basal ganglia circuitry. *Nature* 466, 622–626.
- Kriete, T., and Noelle, D.C. (2011). Generalisation benefits of output gating in a model of prefrontal cortex. *Connect. Sci.* 23, 119–129.
- Kriete, T., Noelle, D.C., Cohen, J.D., and O'Reilly, R.C. (2013). Indirection and symbol-like processing in the prefrontal cortex and basal ganglia. *Proc. Natl. Acad. Sci. USA* 110, 16390–16395.
- Lewis, S.J., Dove, A., Robbins, T.W., Barker, R.A., and Owen, A.M. (2004). Striatal contributions to working memory: a functional magnetic resonance imaging study in humans. *Eur. J. Neurosci.* 19, 755–760.
- Marklund, P., Fransson, P., Cabeza, R., Larsson, A., Ingvar, M., and Nyberg, L. (2007a). Unity and diversity of tonic and phasic executive control components in episodic and working memory. *Neuroimage* 36, 1361–1373.
- Marklund, P., Fransson, P., Cabeza, R., Petersson, K.M., Ingvar, M., and Nyberg, L. (2007b). Sustained and transient neural modulations in prefrontal cortex related to declarative long-term memory, working memory, and attention. *Cortex* 43, 22–37.
- McLaren, D.G., Ries, M.L., Xu, G., and Johnson, S.C. (2012). A generalized form of context-dependent psychophysiological interactions (gPPI): a comparison to standard approaches. *Neuroimage* 61, 1277–1286.
- McNab, F., and Klingberg, T. (2008). Prefrontal cortex and basal ganglia control access to working memory. *Nat. Neurosci.* 11, 103–107.
- Mestres-Missé, A., Turner, R., and Friederici, A.D. (2012). An anterior-posterior gradient of cognitive control within the dorsomedial striatum. *Neuroimage* 62, 41–47.
- Mink, J.W. (1996). The basal ganglia: focused selection and inhibition of competing motor programs. *Prog. Neurobiol.* 50, 381–425.
- Moustafa, A.A., Cohen, M.X., Sherman, S.J., and Frank, M.J. (2008). A role for dopamine in temporal decision making and reward maximization in parkinsonism. *J. Neurosci.* 28, 12294–12304.
- Murty, V.P., Sambataro, F., Radulescu, E., Altamura, M., Iudicello, J., Zolnick, B., Weinberger, D.R., Goldberg, T.E., and Mattay, V.S. (2011). Selective updating of working memory content modulates meso-cortico-striatal activity. *Neuroimage* 57, 1264–1272.
- Narayanan, N.S., Prabhakaran, V., Bunge, S.A., Christoff, K., Fine, E.M., and Gabrieli, J.D. (2005). The role of the prefrontal cortex in the maintenance of verbal working memory: an event-related fMRI analysis. *Neuropsychology* 19, 223–232.
- Nee, D.E., and Brown, J.W. (2012). Rostral-caudal gradients of abstraction revealed by multi-variate pattern analysis of working memory. *Neuroimage* 63, 1285–1294.
- Nee, D.E., and Brown, J.W. (2013). Dissociable frontal-striatal and frontal-parietal networks involved in updating hierarchical contexts in working memory. *Cereb. Cortex* 23, 2146–2158.
- Oberauer, K., and Hein, L. (2012). Attention to information in working memory. *Curr. Dir. Psychol. Sci.* 21, 164–169.
- Olsen, A., Ferenc Brunner, J., Evensen, K.A.I., Garzon, B., Landrø, N.I., and Håberg, A.K. (2013). The functional topography and temporal dynamics of overlapping and distinct brain activations for adaptive task control and stable task-set maintenance during performance of an fMRI-adapted clinical continuous performance test. *J. Cogn. Neurosci.* 25, 903–919.
- Postle, B.R., and D'Esposito, M. (1999). Dissociation of human caudate nucleus activity in spatial and nonspatial working memory: an event-related fMRI study. *Brain Res. Cogn. Brain Res.* 8, 107–115.
- Postle, B.R., and D'Esposito, M. (2003). Spatial working memory activity of the caudate nucleus is sensitive to frame of reference. *Cogn. Affect. Behav. Neurosci.* 3, 133–144.
- Ratcliff, R., and Frank, M.J. (2012). Reinforcement-based decision making in cortico-striatal circuits: mutual constraints by neurocomputational and diffusion models. *Neural Comput.* 24, 1186–1229.
- Reynolds, J.R., Braver, T.S., Brown, J.W., and Van der Stigchel, S. (2006). Computational and neural mechanisms of task switching. *Neurocomputing* 69, 1332–1336.
- Reynolds, J.R., O'Reilly, R.C., Cohen, J.D., and Braver, T.S. (2012). The function and organization of lateral prefrontal cortex: a test of competing hypotheses. *PLoS ONE* 7, e30284.
- Rowe, J.B., Toni, I., Josephs, O., Frackowiak, R.S., and Passingham, R.E. (2000). The prefrontal cortex: response selection or maintenance within working memory? *Science* 288, 1656–1660.
- Scimeca, J.M., and Badre, D. (2012). Striatal contributions to declarative memory retrieval. *Neuron* 75, 380–392.

- Snyder, H.R., Hutchison, N., Nyhus, E., Curran, T., Banich, M.T., O'Reilly, R.C., and Munakata, Y. (2010). Neural inhibition enables selection during language processing. *Proc. Natl. Acad. Sci. USA* *107*, 16483–16488.
- Snyder, H.R., Banich, M.T., and Munakata, Y. (2011). Choosing our words: retrieval and selection processes recruit shared neural substrates in left ventrolateral prefrontal cortex. *J. Cogn. Neurosci.* *23*, 3470–3482.
- Tamber-Rosenau, B.J., Esterman, M., Chiu, Y.C., and Yantis, S. (2011). Cortical mechanisms of cognitive control for shifting attention in vision and working memory. *J. Cogn. Neurosci.* *23*, 2905–2919.
- Verstynen, T.D., Badre, D., Jarbo, K., and Schneider, W. (2012). Microstructural organizational patterns in the human corticostriatal system. *J. Neurophysiol.* *107*, 2984–2995.
- Wiecki, T.V., and Frank, M.J. (2013). A computational model of inhibitory control in frontal cortex and basal ganglia. *Psychol. Rev.* *120*, 329–355.
- Woolrich, M.W., Behrens, T.E.J., and Smith, S.M. (2004). Constrained linear basis sets for HRF modelling using Variational Bayes. *Neuroimage* *21*, 1748–1761.