Role of prefrontal cortex and the midbrain dopamine system in working memory updating

Kimberlee D'Ardenne[a,b,c], Neir Eshel[b,d], Joseph Luka[b], Agatha Lenartowicz[b,e], Leigh E. Nystrom[b], and Jonathan D. Cohen[b,e,1]

[a]Department of Chemistry, [b]Princeton Neuroscience Institute, [d]Department of Molecular Biology, and [e]Department of Psychology, Princeton University, Princeton, NJ 08544; and [c]Virginia Tech Carilion Research Institute, Roanoke, VA 24016

This Feature Article is part of a series identified by the Editorial Board as reporting findings of exceptional significance.

Edited by Edward E. Smith, Columbia University, New York, NY, and approved September 19, 2012 (received for review October 11, 2011)

Humans are adept at switching between goal-directed behaviors quickly and effectively. The prefrontal cortex (PFC) is thought to play a critical role by encoding, updating, and maintaining internal representations of task context in working memory. It has also been hypothesized that the encoding of context representations in PFC is regulated by phasic dopamine gating signals. Here we use multimodal methods to test these hypotheses. First we used functional MRI (fMRI) to identify regions of PFC associated with the representation of context in a working memory task. Next we used single-pulse transcranial magnetic stimulation (TMS), guided spatially by our fMRI findings and temporally by previous event-related EEG recordings, to disrupt context encoding while participants performed the same working memory task. We found that TMS pulses to the right dorsolateral PFC (DLPFC) immediately after context presentation, and well in advance of the response, adversely impacted context-dependent relative to context-independent responses. This finding causally implicates right DLPFC function in context encoding. Finally, using the same paradigm, we conducted high-resolution fMRI measurements in brainstem dopaminergic nuclei (ventral tegmental area and substantia nigra) and found phasic responses after presentation of context stimuli relative to other stimuli, consistent with the timing of a gating signal that regulates the encoding of representations in PFC. Furthermore, these responses were positively correlated with behavior, as well as with responses in the same region of right DLPFC targeted in the TMS experiment, lending support to the hypothesis that dopamine phasic signals regulate encoding, and thereby the updating, of context representations in PFC.

cognitive control | context updating | task switching | cognitive neuroscience

How is it that someone can spend hours playing a video game, seemingly oblivious to everything else in the environment, but then instantly jump to answer the phone as soon as it rings, and become equally absorbed in the telephone conversation? One of the greatest challenges in cognitive neuroscience is to understand the brain mechanisms that allow us to steadfastly pursue goals and yet flexibly and adaptively switch between these as circumstances demand.

It has long been recognized that prefrontal cortex (PFC) is critical for this ability, often referred to as cognitive control (1, 2). There is a growing consensus that PFC contributes to cognitive control by maintaining in working memory representations of information (e.g., instructions or rules) that are necessary to direct attention to and successfully perform goal-directed tasks (3–6). These representations are thought to serve as context for, and guide the execution of, processes carried out by other brain systems responsible for performing the task (7–10).

The conflicting needs to maintain and yet appropriately update context representations present opposing perils—a sort of Scylla and Charybdis—that must be navigated to avoid both distraction and perseveration. On the one hand, PFC representations must persevere throughout performance of a task, so that attention and behavior remain directed toward the goal and are not derailed by internal or external distraction. On the other hand, PFC representations cannot be so resistant to modification that they perseverate when task demands or goals change. Understanding the mechanisms that allow us to persist in pursuing a goal in the face of distraction and yet flexibly adapt behavior when circumstances change is critical not only for understanding the remarkably adaptive nature of normal human behavior, but also how and why the mechanisms involved break down in disease, yielding to distractibility and/or perseveration.

Several studies have documented the updating of context information in PFC (11–18). However, few studies have addressed the mechanisms by which this is accomplished. One hypothesis for how the brain updates context is with the use of a gating mechanism. This proposes that the updating of context representations depends on a gating signal that regulates the inputs to PFC and thereby the encoding of new context representations in working memory. In the absence of this signal, inputs have a weak influence on PFC, allowing representations that are currently active to persist. When the gating signal occurs, inputs to the PFC are enhanced. This allows new representations to be activated (encoded) and replace the ones previously maintained, thereby updating context information in working memory. The new context information is then maintained until the next gating signal occurs. The gating signal is presumed to occur only when there is an indication that conditions have changed and a new task or goal should be pursued.

Theoretical work has demonstrated that, in general, gating mechanisms are an effective way to regulate the updating of working memory (19) and, in particular, the updating of task and/or goal representations [both forms of context information (20)]. Several models have been proposed for how a gating mechanism might be implemented in the brain, all of which assign an important role to dopamine (21–23). The involvement of dopamine presents a potential solution to an important challenge for the gating hypothesis: to explain how the timing of the gating signal is learned. Phasic dopamine signals have been proposed to implement a form of reinforcement learning (24). According to this theory, the phasic release of dopamine acts as a learning signal that is used to predict when rewards will occur. Consistent with this theory, there is growing evidence that dopamine neurons fire in response to events associated with reward prediction errors—that is, unexpected events that are associated with a subsequent reward (25, 26). These

Author contributions: K.D., N.E., A.L., L.E.N., and J.D.C. designed research; K.D., N.E., and J.L. performed research; K.D., N.E., J.L., L.E.N., and J.D.C. analyzed data; and K.D., N.E., J.L., A.L., L.E.N., and J.D.C. wrote the paper.

The authors declare no conflict of interest.

This article is a PNAS Direct Submission.

See Commentary on page 19878.

[1]To whom correspondence should be addressed. E-mail: jdc@princeton.edu.

This article contains supporting information online at www.pnas.org/lookup/suppl/doi:10.1073/pnas.1116727109/-/DCSupplemental.

are also precisely the conditions under which a gating signal (and consequent updating of context representations) should occur—when an unexpected event signals the opportunity for reward, and behavior should be redirected to obtain it. Thus, the release of dopamine when a gating signal should occur can strengthen the likelihood that this signal will occur again under similar circumstances in the future, providing an adaptive, self-organizing mechanism for learning the timing of gating signals.

The computational plausibility of the gating hypothesis has been established in at least two different types of models. One proposes that dopamine release simultaneously implements the gating signal in PFC and the learning signal used to train the system when this should occur (Fig. 1) (21). This model is simple and exploits the idea that both the gating and learning effects of dopamine could be implemented by the same physiological mechanism: gain control (27–29). A different model proposes that dopamine is used to train the timing of the gating signal but that the gating mechanism itself is implemented by the basal ganglia (22). Although these models differ with regard to the source of the gating signal, both make the same novel prediction: that, at least while performing a relatively new task, gating signals and updating of working memory representations in PFC should be accompanied by the phasic release of dopamine. Recently findings have been reported that are consistent with this hypothesis, including the involvement of dopaminergic mechanisms in training on updating in a working memory task (30, 31), and in response to cued preparation for a control-demanding task (32, 33). However, to our knowledge, no studies have tested the causal involvement of dorsolateral PFC (DLPFC) in updating of context representations, or the association of phasic dopaminergic signals with this PFC process.

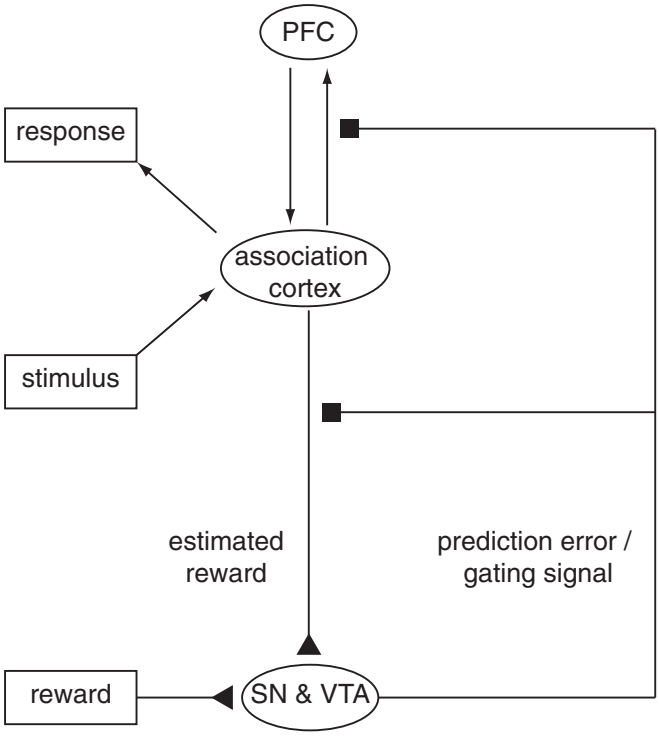

Fig. 1. Midbrain dopamine neurons in the SN and VTA broadcast signals that may be used to gate information into the PFC. Dopamine neurons have been shown to encode reward prediction errors through phasic changes in firing rate (25, 75). The reward prediction error is the difference between rewards received and rewards predicted (24). The theory that dopamine plays a role in updating working memory information in PFC (21) posits that the phasic increases in firing rate that encode reward prediction errors not only modulate the sensory inputs predicting reward (76, 77) but also simultaneously adjust the gain of inputs to the PFC (27). Modulation of these inputs permits the selective updating of representations in PFC. The reward prediction error and the gating signal work in concert because stimuli linked to the PFC representations needed to procure reward (e.g., representations of task rules and goals) themselves also predict reward. Arrows between PFC and association cortex indicate connectivity between cortical areas, triangles represent excitatory input to the SN and VTA, and squares represent gain modulation by dopamine at target sites.

Here we sought to test these predictions in three steps. First, we used functional MRI (fMRI) to identify regions of PFC associated with the representation of context information in a simple context-dependent working memory task. Next we used single-pulse transcranial magnetic stimulation (spTMS), guided by our fMRI findings and those of an event-related EEG study (18), to interfere with the encoding and thereby updating of context representations in PFC to test their causal relationship to task performance. Finally we used fMRI methods for brainstem imaging (34) to test for a correlation between the phasic activation of dopaminergic nuclei and activation of the same areas of PFC identified with the encoding of context representations in the preceding fMRI and TMS experiments.

Simple Task Requiring Updating Context Representations in Working Memory. For all of the experiments, participants performed the AX-CPT task (7, 35, 36), a variant of the Continuous Performance Test (CPT) (37), in which participants are presented with a sequence of letters and required to respond to a subset of these ("probes") with a button press (*SI Text, Note A*). The sequences comprised three types of trials: context-dependent, context-independent, and control (Fig. 2). Context-dependent trials consisted of pairs of letters, the first of which (the "cue") determined how to respond to the second (the probe; e.g., press button 1 to an X after an A, but press button 2 to an X after a B). In context-independent trials the response to the probe was the same, irrespective of the preceding cue (e.g., press button 1 to a W irrespective of whether it followed C or D). In these trials, as in context-dependent trials, participants could use the cue to prepare the response mappings for the upcoming probes, but this was not necessary because the response could be determined from the probe alone. Finally, in control trials, only a cue appeared that required no response. Thus, in comparison with control trials, cues in the context-dependent and context-independent conditions could be used to prepare for the upcoming probes. However, in context-dependent trials, this was a necessity: they required that participants encode a representation of the response mapping that was needed to respond correctly to the forthcoming probes. We predicted that this would be reliably associated with a gating signal, used to encode and update the context representation in PFC. To test this prediction, we first used fMRI to identify regions of brain activity associated with the updating of context and then used these regions as targets for spTMS. Finally, we conducted a follow-up fMRI study to determine whether these processes were associated with activity in midbrain dopaminergic nuclei, as predicted by the gating hypothesis.

Results and Discussion

Experiment 1: Association of DLPFC Activity with Context Updating. Twelve participants performed the AX-CPT task in the MRI scanner while whole-brain functional data were collected. Behavioral measures did not differ significantly across conditions (*SI Text, Note B*). Group analysis of the fMRI data revealed bilateral activation of DLPFC associated with the updating of context, with right DLPFC showing greater activity than left (Fig. 3*A*). This activation of DLPFC replicates previous findings using tasks with similar demands for the updating of context (9, 10), although the lateralization of this effect varies across studies and tasks. However, none of these studies has provided evidence that activity in DLPFC is causally related to the updating of context representations. To test this, we conducted a second experiment in the

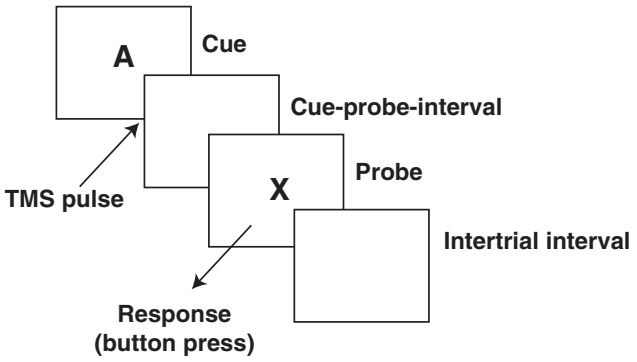

Trial Type	Cue	Probe	Response
Context Dependent	A	X Y	1 2
	B	X Y	2 1
Context Independent	C	W Z	1 2
	D	W Z	1 2
Control	H, K, L, M	none	none

Fig. 2. Behavioral task. The AX-CPT task consisted of three types of trials: context-dependent, context-independent, and control. Control trials (not shown in the figure) consisted of display of a cue letter (500 ms) followed by an intertrial interval. Context-dependent and context-independent trials (shown) consisted of display of the cue (500 ms), a blank screen comprising the cue-probe interval, presentation of the probe (1,000 ms), the participant's response to the probe, and finally the intertrial interval. Letters shown in the figure and table are examples (*Materials and Methods* gives details of actual stimulus assignments and interstimulus intervals). All responses involved pressing one of two buttons on a response box. In the TMS experiment, control trials were not included, and pulses were delivered at either 10, 100, 150, or 200 ms after the onset of the cue.

same participants, in which we sought to disrupt the encoding and thereby updating of context through the use of spTMS.

Experiment 2: Interference of Context Updating by spTMS in DLPFC. Participants in experiment 1 returned the following week and underwent spTMS while performing context-dependent and context-independent trials of the AX-CPT task. TMS pulses were localized individually for each participant, using the fMRI data from experiment 1. To do so, we analyzed each participant's data individually to identify the areas of DLPFC associated with the updating of context (*SI Text, Note C*). We found such activation in all 12 participants (Fig. 3*B*). TMS pulses were targeted at the coordinates of the most active voxel in either the left or right DLPFC for each participant (Fig. 3*C* and Table S1).

Single TMS pulses were delivered at 10, 100, 150, or 200 ms after the onset of the cue letter. Pulse delivery times were chosen on the basis of a previous study using scalp electrical recordings, and the same AX-CPT task (*SI Text, Note A*), to identify the timing of context updating in PFC (18). That study revealed a positivity that was selectively enhanced in context-dependent trials, beginning ~150 ms after cue presentation, and source localized to right DLPFC. This was interpreted as being consistent with the occurrence of a gating signal in right DLPFC after cues in the context-dependent condition of the task. On the basis of these findings, we predicted that a TMS pulse delivered to right DLPFC within 200 ms after cue presentation would disrupt the gating signal, interfering with the encoding of context information, and thereby impact responses to the probe in context-dependent trials. To test this prediction, responses to probes (reaction time and accuracy) were analyzed using repeated-measures ANOVAs, with participants as a random effect and trial type (context-dependent or context-independent), stimulation site (left or right DLPFC), and pulse time (10, 100, 150, or 200 ms) as within-subject factors.

As in experiment 1, mean accuracy was high, and there were no significant effects of trial type, stimulation site, or timing, nor any interaction among these factors (*SI Text, Note B*). However, TMS did affect reaction time (Fig. 4). As predicted, there was a significant three-way interaction between trial type, stimulation site, and pulse time [$F(3,33) = 4.19$; $P = 0.01$], but no main effects or two-way interactions (Fig. 4). To confirm the source of the three-way interaction, we performed ANOVAs separately for the left and right DLPFC stimulation sites. We found that spTMS to left DLPFC had no significant effect on reaction times ($P > 0.50$). However, as predicted, when applied to right DLPFC, there was a main effect of pulse time after context-dependent cues [$F(3,33) = 2.96$; $P = 0.05$] but not after context-independent cues [$F(3,33) = 0.35$; $P = 0.79$]. Specifically, participants were slower to respond to probe letters when a TMS pulse was applied to right DLPFC 150 ms after the onset of context-dependent cues (Fig. 4*A*) compared with other intervals (*SI Text, Note D*).

In summary, the results of this experiment revealed a selective effect of spTMS on performance: only pulses delivered to right DLPFC 150 ms after context-dependent cues significantly influenced behavior. Pulses delivered at other times or in a different location, or when the response to the probe did not rely on the cue, did not have a measurable effect. Furthermore, note that even when the pulse did affect performance, it was delivered at

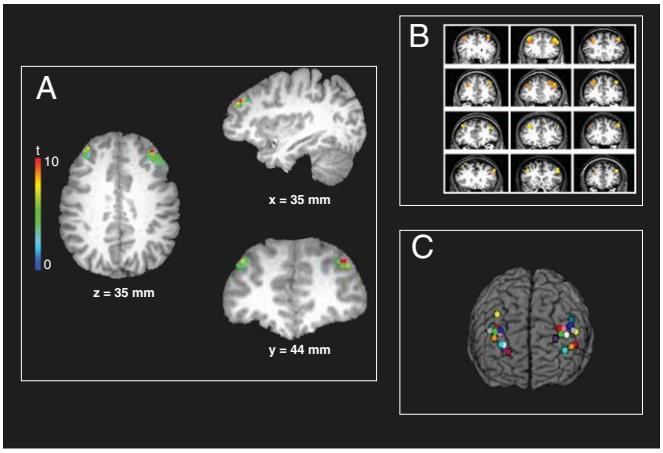

Fig. 3. Whole-brain fMRI results and TMS stimulation sites for each participant. (*A*) Results of the group analysis of the fMRI data from experiment 1, identifying areas associated with context updating in the AX-CPT [$n = 12$; $P < 0.05$ corrected for multiple comparisons; left DLPFC region is 292 mm^3 (9 contiguous voxels) and the right DLPFC region is 1199 mm^3 (37 voxels)]. The statistical map is overlaid on the T1-weighted structural image in Talairach space from a representative participant. (*B*) Sites of activity in individual participants used to target TMS pulses in experiment 2. In all participants, this contrast identified activation in left and right DLPFC that was centered on BA 9 ($n = 12$; $P < 10^{-6}$). Statistical maps for each participant are displayed on their T1-weighted structural images. (*C*) TMS stimulation sites in the right and left DLPFC for each participant in experiment 2. For each participant, the most active voxel within each area shown in *B* was chosen as the stimulation site. The locations of the stimulation sites are shown on a 3D reconstruction of a canonical Talairach brain. All images are displayed in radiological convention (left in the image corresponds to participant's right side). *Materials and Methods* gives details of statistical analyses.

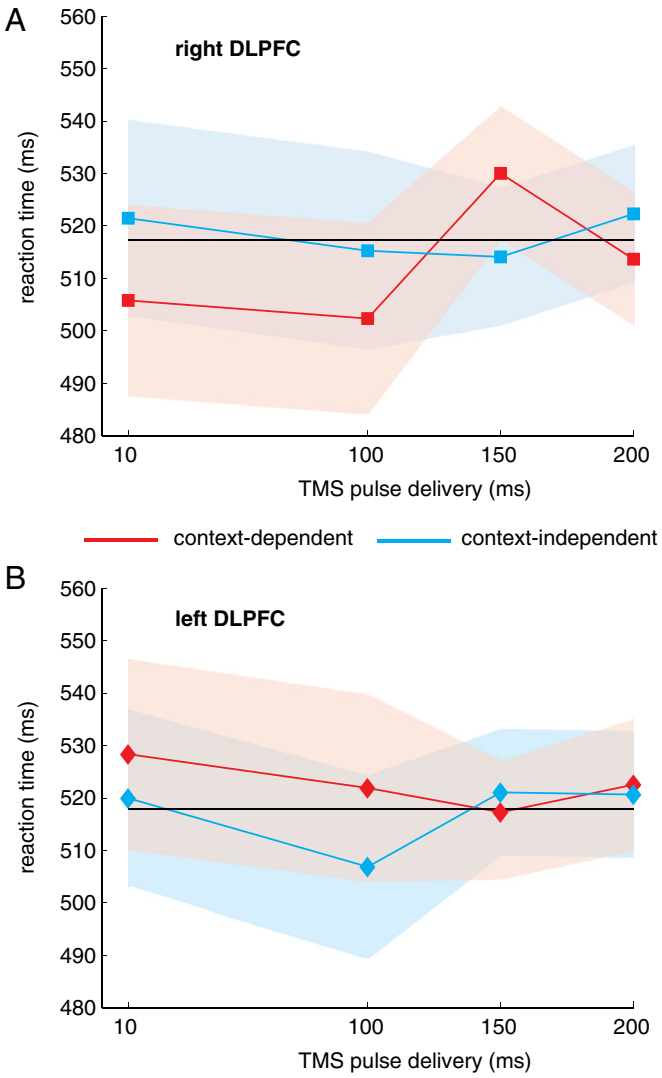

Fig. 4. TMS disrupts performance during context-dependent trials. (*A*) TMS pulses applied to right DLPFC at 150 ms after cue onset significantly slowed reaction time compared with pulses applied at 10 ms after cue onset [$t(11) = 2.26$; $P < 0.05$] or at 100 ms [$t(11) = 2.40$; $P < 0.05$]. There was a nonsignificant trend at 200 ms [$t(11) = 1.87$; $P < 0.1$]. No other pulse time comparisons showed a significant effect on reaction time during context-dependent trials. (*B*) No pulse time comparisons showed a significant effect on reaction time during any trial type in left DLPFC. In *A* and *B*, shaded regions represent 95% confidence intervals, and the black line shows the overall mean reaction time.

least 1 and sometimes 2 s before the appearance of the probe and corresponding response, and that pulses delivered contralateral to the hand of the response did not have an impact on performance. Thus, it is unlikely that the effect of the pulse was due to a direct effect on motor processing. Rather, the TMS pulse seems to have influenced the use of the cue to update context information needed to respond to the subsequent probe.

These findings raise several questions. First, if spTMS interfered with context updating, why did this prolong reaction times but not impair accuracy in context-dependent trials? One possibility is that it may have partially or only transiently disrupted working memory maintenance, and that the recovery from this induced participants to take longer to respond and thereby preserve accuracy. However, this seems unlikely for at least two reasons. First, the occurrence of the pulse at 150 ms should have left ample time for recovery (a minimum of 1 s) before probe presentation. Second, the effect should then have been at least as great or greater in the 200-ms condition, which occurred closer in time to the probe and response; however, this was not the case. An alternative, consistent with the gating hypothesis, is that the TMS pulse disrupted context encoding in working memory as predicted and that this forced participants to access context information from episodic memory when the probe appeared. Although further work is necessary to test this conjecture, it is consistent with the current literatures on cognitive control (e.g., ref. 38 and 39) and prospective memory (e.g., refs. 40–43). These lines of research have converged in identifying two types of strategies that people use for storing information when it is needed for later use: proactive (or "activation-based"), in which task-relevant information is encoded and maintained actively in working memory to prepare for and ensure an efficient response to a later stimulus; and reactive (or "retrieval-based"), which relies on the storage and retrieval of task-relevant information from episodic memory [presumably served by regions other than PFC, such as hippocampus (43)]. The latter forgoes the effort required to encode and maintain information in working memory, but at the expense of the time required to retrieve it from episodic memory when it is needed. In our study, the TMS pulse may have disrupted the use of the proactive strategy that relies on working memory, forcing participants to use the reactive strategy that relies on episodic memory. This would preserve accuracy, but at the cost of the additional time required to retrieve context information from episodic memory at the time of the probe, thus prolonging reaction time. This explanation is also consistent with the correlation between reaction time findings and brainstem activity reported below.

A second question is why context-dependent cues were associated with some activity in left DLPFC in the fMRI study (albeit less than on the right), whereas spTMS interfered with context updating only when it was applied to right DLPFC. It is possible that the left DLPFC is also necessary for context updating but follows a different time course and was thus unaffected by a disruption within the tested time window (15, 44). However, the lateralization we observed in the TMS study was predicted by the source localization of an event-related potential (ERP) observed in our previous EEG study, which was associated with the updating of context in the same task (18). It is also consistent with the additional observation we report below, of a correlation between brainstem dopaminergic activity and right (but not left) DLPFC responses to cues in context-dependent trials. This suggests that, at least for this task, right DLPFC may be critical for encoding and/or updating context information, whereas left DLPFC may perform some other control-related function(s), such as maintenance (45, 46).

Finally, spTMS at 150 ms after the cue disrupted context-dependent performance, but this effect was diminished for the 200-ms interval. On the surface, this may seem to be at odds with the observation in our EEG study that the context-dependent scalp potential peaks between 180 and 200 ms after the cue (18). However, a TMS pulse is most likely to interfere with an ERP when it precedes the ERP (47). Thus, a pulse delivered at 150 ms after the cue, immediately preceding or coinciding with the onset of the gating-related ERP, would be more likely to have an interfering effect than one delivered at 200 ms, when the ERP has already peaked. These results are also consistent with single-cell studies (48–50) and other ERP studies (17, 18, 51, 52) that have demonstrated the involvement of PFC in selecting behaviorally relevant information within the time frame of ~150–200 ms after stimulus onset.

Experiment 3: Association of Brainstem Dopaminergic Nuclei Responses with Context Updating. The findings from the first two experiments implicate a region of right DLPFC in context encoding and updating. In a final experiment, we sought to test whether activity in right DLPFC is associated with phasic dopamine release, as

predicted by models of a gating mechanism (21–23). We did so by collecting high-resolution and cardiac-gated fMRI data from the brainstems of participants (n = 24) completing the same task used in the whole-brain fMRI session of experiment 1 (Fig. 2). Behavioral findings were similar to those observed in experiment 1 (*SI Text, Note B*). Functional image acquisition was centered on the midbrain and oriented to include as much as possible of both the brainstem dopaminergic nuclei [substantia nigra (SN) and ventral tegmental area (VTA)] and the regions of DLPFC identified in experiment 1 (Fig. 5*A*).

A phasic increase in the blood oxygen level-dependent (BOLD) response was found after cues in a midbrain region including SN and VTA, which was significantly greater for context-dependent cues than context-independent cues (Fig. 5*B*). In contrast, the BOLD response to probe letters was not significantly different across trial types [$t(23) = 0.90$; $P = 0.38$]. There was also no significant BOLD response to control trial cue letters.

To further establish the specificity of midbrain BOLD time courses to context updating, we examined their relationship to the time course of signals measured in DLPFC. Both the midbrain (Fig. 5*C*) and right DLPFC (Fig. 5*D*, solid lines) exhibited phasic BOLD responses to context-dependent but not context-independent cues, whereas responses in left DLPFC to both cue types did not differ from baseline (Fig. 5*D*, dashed lines). Furthermore, midbrain activity for the contrast of context-dependent > context-independent cues was positively and significantly correlated with right DLPFC activity ($r = 0.24$; $P = 0.03$) but not with left DLPFC ($r = 0.15$; $P = 0.13$). Consistent with this finding, an examination of the relationship between midbrain and DLPFC activity separately for the two cue types revealed that midbrain activity again correlated with right DLPFC for context-dependent cues ($r = 0.25$; $P = 0.04$) but not for context-independent cues nor with left DLPFC for either cue type. The specificity of these functional correlations to the context-dependent condition, and their lateralization to right DLPFC, are consistent with the spTMS results reported above (Fig. 4).

The correlation of midbrain activity with right DLPFC in the context-dependent condition showed variance across participants. One source of this variance may have been the differential use, by different participants, of the two storage strategies discussed above. Some may have elected more frequently to use the proactive strategy that relies on working memory, and therefore should more consistently have exhibited a gating signal, whereas

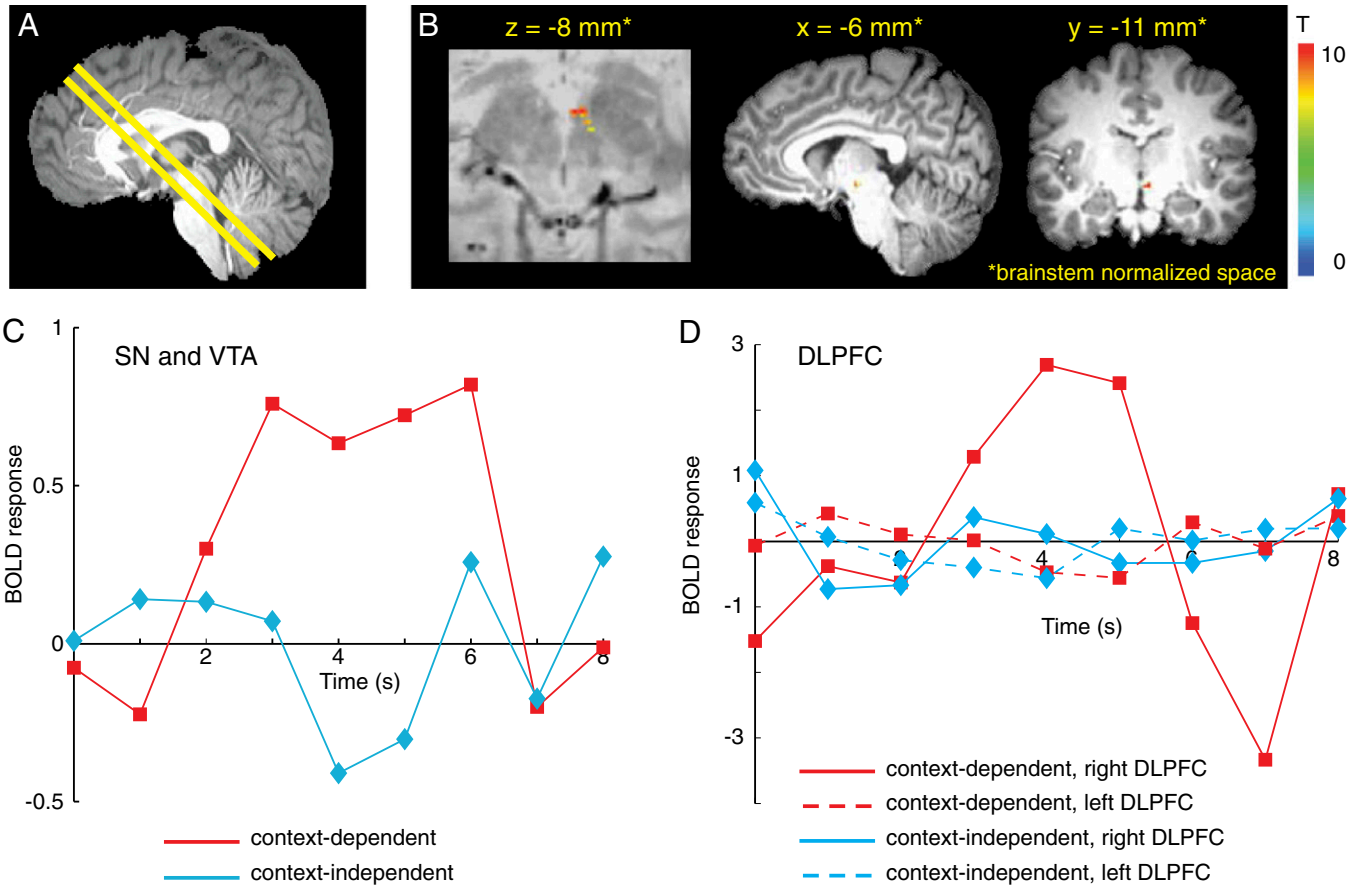

Fig. 5. BOLD responses in the midbrain SN and VTA were associated with context updating. (*A*) For brainstem fMRI, the midbrain was identified in the central sagittal slice of the T1-weighted structural image, and an oblique slab comprising axial/coronal slices was centered on as much of the SN and VTA as possible (constrained by the number of slices allowed by the participant's heart rate) and then tilted to include the regions of DLPFC shown in Fig. 3*B*. (*B*) A random effects general linear model analysis revealed that the BOLD response in the SN and VTA for context-dependent cues was greater than the BOLD response for context-independent cues (n = 24; $P < 0.05$, corrected for multiple comparisons). No other brain regions within the brainstem exhibited this effect. Statistical maps are overlaid on a proton-density weighted image (*Left*) and a T1-weighted image (*Right*) in brainstem-normalized space (74). (*C*) Average BOLD event-related time courses in SN and VTA (*B*) during context-dependent and context-independent trials in the AX-CPT task. The BOLD response to context-dependent cue letters (red) was greater than the BOLD responses to context-independent cue letters (blue). Cue letter presentation occurred at time t = 0. (*D*) Average BOLD event-related time courses from right and left DLPFC ROIs during context-dependent and context-independent trials. In right DLPFC, the BOLD response to context-dependent cue letters (solid red) was greater than the BOLD response to context-independent cue letters (solid blue). Left DLPFC (red and blue dashed lines) showed no BOLD responses to cue letters. BOLD time courses are in arbitrary MR units.

others may have opted more frequently to use the reactive strategy, relying on episodic memory and thus not as consistently exhibited a gating signal. As discussed above, the reactive strategy should be associated with slower responding in the context-dependent condition. Thus, we predicted that the midbrain signal after the cue should correlate with reaction time in the context-dependent condition. This was confirmed ($r = 0.55$, $P = 0.03$; Fig. 6*A*). Finally, we also found a significant reduction of the SN and VTA responses to context-dependent (vs. context-independent) cues on error trials [$t(23) = 12.50$; $P = 10^{-12}$; Fig. 6*B*], consistent with its association with a gating signal involved in proactive control.

Our findings corroborate the predictions made by a model in which a gating signal regulates the encoding of new information in PFC and thereby the updating of context information in working memory. In a previous ERP study we identified a right-lateralized frontal signal that was selectively enhanced 150–200 ms after context-dependent cues (18). Here, under similar experimental conditions, we used (*i*) fMRI to localize such activity to a region of right DLPFC; (*ii*) spTMS to disrupt this signal and selectively impair performance in context-dependent trials; and (*iii*) specially adapted methods for brainstem imaging to demonstrate that phasic activity of dopaminergic nuclei correlates with activity of this same region of right DLPFC, as well as context-dependent behavior. The specificity of the findings to the context-dependent conditions of the task strongly suggests that the effects were not related to simpler or more general processes—such as sensory encoding, motor responding, or overall motivation—but rather to the encoding and updating of context information.

Previous work, both theoretical and empirical, has suggested that both PFC and dopamine play important roles in working memory function and cognitive control (9, 10, 21, 22, 53–56). However, the data presented here go beyond previous findings by causally implicating a region of DLPFC in the updating of context information in working memory and relating this directly to phasic dopamine signals (*SI Text, Note E*). The latter is consistent with models proposing that dopamine plays an important role in gating. However, it does not adjudicate between different hypotheses concerning the precise role that dopamine plays—that is, whether dopamine signals themselves implement gating in PFC or whether they serve to train a gating mechanism that is implemented by other brain systems such as the basal ganglia that, in turn, regulate access to PFC (*SI Text, Note F*). Future studies that examine changes in these signals over the course of learning, or that combine similar behavioral methods with more invasive measurements in nonhuman species, are needed to further refine our understanding of the implementation of the gating mechanism in the brain. Such studies could be complemented by other forms of intervention (e.g., pharmacological in humans, and/or lesions in nonhuman species) to establish the causal role of dopaminergic mechanisms in gating.

We should also note that our study addressed working memory for context information of a particular type (specific to the demands of the task that was used) and thus identified a restricted area of PFC. We presume that other areas of prefrontal cortex are responsible for representing other types of information but are also subject to dopamine-dependent gating signals. Indeed, recent work has suggested that different regions of prefrontal cortex may be responsible for representing different information at varying levels of abstractness or complexity (9, 10, 57, 58), and that basal ganglia mechanisms may be responsible for restricting gating signals selectively to those regions required for the performance of a particular task (22).

Our findings provide evidence that DLPFC is causally involved in representing and actively maintaining context information, and that the updating of this information is associated with phasic dopamine signaling. Although further work is needed to refine our understanding of these mechanisms, the present findings provide an initial window into their operation and their role in generating the remarkable flexibility that is characteristic of the human capacity for cognitive control.

Materials and Methods

Participants. The Princeton University and the Baylor College of Medicine institutional review boards approved for human participants the three experiments discussed in this article. All participants were right-handed and were screened to rule out neurological disorders and other contraindications for MRI or spTMS. For experiments 1 and 2, participants were recruited from the Princeton University community. For experiment 3, participants were recruited from the Texas Medical Center community and the Houston, TX area. In the first two experiments (fMRI followed by spTMS approximately 1 wk later), 19 participants were imaged and 7 were excluded for excessive head motion in the scanner (motion greater than 3 mm in any direction; $n = 4$),

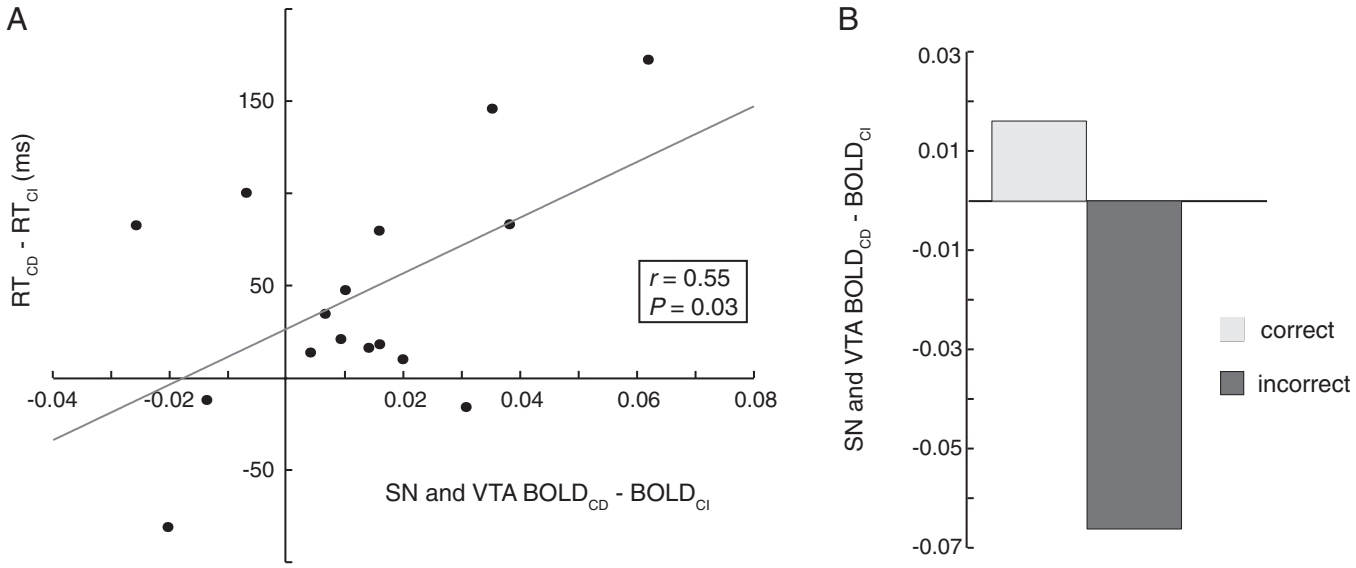

Fig. 6. The SN and VTA BOLD responses to context updating were correlated with reaction time and were reduced on error trials. (*A*) Correlation across participants of differences in reaction time (RT) for context-dependent vs. context-independent trials with corresponding differences in the BOLD responses in SN and VTA (Fig. 5*B*). RT differences are in milliseconds; BOLD signal differences are in arbitrary MR units. (*B*) This same midbrain region exhibited significantly diminished context updating effect on BOLD response during error compared with correct trials [$t(23) = 12.50$; $P = 10^{-12}$].

difficulty tolerating TMS (n = 2), or excessive head movement during TMS (n = 1). The remaining 12 participants (6 men) varied in age from 19 to 28 y. For experiment 3, 26 participants were imaged and 2 were excluded for excessive head motion (defined as motion greater than 1.5 mm in any direction). The remaining 24 participants (11 men) varied in age from 20 to 53 y. All but one participant (who was ambidextrous) were right handed. Written consent was obtained from all participants.

Stimulus Paradigm. In all experiments, participants were told they were to perform a task designed to study working memory. The task was a variant of the AX-CPT (7, 35, 36), in which participants viewed a continuous stream of letters and were required to respond to specific letter pairings (constituting a "trial") with a button press. There were three types of trials: context-dependent, context-independent, and control (Fig. 2). In context-dependent trials, one of two "cues" (e.g., A or B) was followed by a "probe" (e.g., X or Y). Upon viewing the probe, participants had to respond with a button press (either "1" by pressing the button under their index finger or "2" by pressing the button under their middle finger). The correct button press was determined by the pairing of the cue and probe letters. For example, an A followed by an X would require pressing button 1, whereas an A followed by a Y would require pressing button 2. Conversely, a B followed by an X would require pressing button 2, and a B followed by a Y would require pressing button 1. Therefore, for context-dependent trials, execution of the correct response to the probe required that the participant encode and represent the response rule associated with the cue.

In context-independent trials, participants again viewed one of two cues (e.g., C or D) and then a probe (e.g., W or Z). In these trials, the probe always mapped to the same response (a 1 button press for the probe letter W and a 2 button press for the probe letter Z) irrespective of the cue. Because the probe provided the information needed to respond appropriately, context-independent trials did not require updating working memory upon viewing the cue. Nevertheless, doing so could be helpful to prepare the response mappings for the upcoming probe.

Control trials, which were used in the MRI studies but not in the spTMS experiment, were composed of only a cue (e.g., H, K, L, or M) and no probe nor any response. Thus, they did not require any updating of context.

Experiment 1: Whole-brain fMRI AX-CPT. Before entering the scanner, participants were given task instructions and completed several practice sessions on a computer to learn the letter pairings for context-dependent and context-independent trials. When participants achieved at least 90% accuracy on the practice sessions, they began the experiment. Visual stimuli were presented using E-Prime software (Psychology Software Tools) and were displayed on a rear-projection screen. Participants viewed the screen through a mirror that was attached to the head coil. High-density foam padding was used to stabilize the participant's head and to minimize head motion during the experiment. Stimuli were white capital letters on a black background. Cue and probe letters for each condition were pseudorandomly selected from a set of 12 letters (A, B, C, D, H, K, L, M, W, X, Y, and Z) and then held constant for that participant. The cue probe interval was jittered randomly among 3,250, 5,250, or 7,250 ms.

Experiment 2: spTMS AX-CPT. Only context-dependent and context-independent trials were used for this experiment, to keep the session as short as possible (*SI Text, Note A*) (Fig. 2). These were randomly intermixed within blocks. For a similar reason, shorter cue probe and intertrial intervals were used, both of which alternated randomly between 1,000 and 2,000 ms. All stimuli were white capital letters presented on a black background, using a script written in MATLAB 5.2 (MathWorks) and the Psychophysics toolbox (www.psychtoolbox.org). Cue and probe letters were pseudorandomly selected for each condition from a set of eight letters (A, B, C, D, W, X, Y, and Z), with the provision that the assignments for each participant differed from the ones used in the fMRI session (experiment 1). Participants were given practice with the new letter assignments and, when they achieved 90% accuracy, began the TMS session.

Experiment 3: Brainstem fMRI AX-CPT. The brainstem fMRI experiment used the same task presentation and response collection protocols as experiment 1 described above.

Data Acquisition. ***Experiment 1: Whole-brain fMRI.*** The goal of the whole-brain fMRI experiment was to identify regions associated with the updating of context representations and localize these within individual participants as targets for the application of spTMS. An MRI-compatible button response system was used to record behavioral performance. Participants responded with their right thumb or index finger and were instructed to respond as quickly as possible. All fMRI sessions comprised six blocks, each containing 48 randomly intermixed trials, including 24 control trials, 12 context-dependent trials, and 12 context-independent trials (task structure was modeled after ref. 18). All images were collected on Princeton Neuroscience Institute's 3T Siemens Allegra head-dedicated MRI scanner. Anatomical images were collected using a T1-weighted magnetization prepared rapid acquisition gradient echo (MPRAGE) sequence (1-mm^3 voxels), and whole-brain functional images were collected using a T2*-weighted echo-planar imaging sequence [EPI; repetition time (TR)/echo time (TE) 2,000/30 ms, flip angle (FA) 75°, 64 × 64 matrix, 3 × 3-mm^2 in-plane voxels, 30 × 3.6-mm-thick slices]. The presentation of task events was synchronized with TR onsets. Six task runs (191 scans each) of fMRI data were acquired with a brief rest delay (at least 30 s) between runs. The first three scans of each run were discarded to allow longitudinal magnetization to reach a steady state.

Experiment 2: spTMS. Participants in experiment 1 returned to the laboratory 1 wk later for the TMS session. The goal of the session was to use single pulses of TMS delivered to the regions of DLFPC identified in the fMRI study, and timed according to previous ERP results (18) to disrupt the encoding of context. Participants performed three blocks of trials with spTMS to left DLPFC, and three with spTMS to right DLPFC, with the order of blocks counterbalanced across participants. Each block included 48 trials, divided evenly between context-dependent and context-independent trials.

On each trial, TMS was applied at one of four delays—10, 100, 150, or 200 ms—after cue presentation. Twice as many trials used the 150 and 200 ms delays as the 10- and 100-ms delays. Thus, of the 144 trials (three blocks) in each target region (i.e., left or right PFC), there were 24 trials each for the 10- and 100-ms delays and 48 trials each for the 150- and 200-ms delays. The order of stimuli and delays was randomized across trials and participants. Responses were collected using a standard keyboard, with the numbers 1 and 2 used for responses. Both reaction time and accuracy were recorded.

The Brainsight system of frameless stereotaxy (Rogue Research) was used to target stimulation sites on an individual basis. Participants sat in a height-adjusted chair and placed their chins in a chinrest. Using the Polaris infrared system (Northern Digital), the participants' heads were registered with their MRI scans by marking the position of scalp landmarks that were visible on the scans (i.e., nasion, inion, tip of the nose, and the intertragal notches of the left and right ears). After registration, infrared tracking was used to target the left and right DLPFC stimulation sites. A 70-mm figure-eight TMS coil (Magstim) was positioned tangentially to the scalp over these areas, so that an imaginary line drawn from the center of the coil and perpendicular to the plane of the coil casing would intersect the targeted coordinates. The coil was clamped in place, and the infrared camera continued to monitor its position relative to the participant's brain in real time during the task. If the participant's head moved out of position, the coil was repositioned between blocks. This ensured that spTMS was always applied within 5 mm of the targeted site.

Pulses were applied with the Magstim 200 monopulse stimulator. For each participant, the output of the stimulator was set to 60% of the maximum output. If the stimulation led to persistent and distracting facial twitches, however, the output was reduced in steps of 1% until the twitching stopped. For these participants, the average output was 56%. Although many studies have used individual motor thresholds to determine the stimulator output (e.g., 59–61), we decided against this practice because there is little evidence to suggest that the threshold over one brain region is a reliable indicator of the threshold over any other region (62, 63). In line with several recent studies (64–69), we believed that a standard output would be the least arbitrary method. Furthermore, the use of within-subject, within-site controls (i.e., two trial types and four stimulation times) ensured that our results would not be confounded by differences in effective stimulation strength between areas or participants.

Experiment 3: Brainstem fMRI. All images were acquired on one of two 3T Siemens Allegra head-dedicated MRI scanners using circularly polarized head volume coils. The scanners were located in the Human Neuroimaging Laboratory at Baylor College of Medicine.

High-resolution (0.25-mm^3 voxels) T1-weighted whole-brain structural images were acquired with an MPRAGE pulse sequence at the beginning of the scanning session. All functional data were acquired using a high-resolution T2*-weighted EPI sequence (TE 41 ms, FA 60°, 128 × 128 matrix, 1.5 × 1.5-mm^2 voxels in-plane, 1.9-mm-thick slices) that was cardiac gated (cf. 34). The participant's pulse, monitored with a finger pulse-oximeter that interfaced with the scanner, was used to trigger the scanner during functional imaging. The pulse-oximeter was placed on the left middle finger of each participant.

The midbrain was identified on the central sagittal slices of the high-resolution structural, and a slab comprising axial/coronal slices (each slice 1.9 mm thick; Fig. 5*A*) was centered on the midbrain and tilted to include as much as possible of the SN, VTA, and regions of the DLPFC identified in the whole-brain fMRI study (Fig. 3*A*). The number of slices in the axial/coronal

slab was determined according to the participant's heart rate. The volume acquisition time (comparable to the TR for non-cardiac–gated imaging) and the maximum length of the acquisition window (how long the scanner will wait for a heartbeat to trigger the next image acquisition) were also determined according to the participant's heart rate. The volume acquisition time was always set to be as fast as possible, depending on the number of slices used. The mean number of slices was 8 ± 0.95 (SD), with a maximum of 10 slices and a minimum of 6 slices across participants. The mean volume acquisition time was 804 ms, and the mean acquisition window was 854 ms. All scanner trigger times during functional scanning were recorded and used in data analysis.

After functional scanning, a non–cardiac-gated whole-brain functional image (TR/TE 2,500/41 ms, FA 60°, 128 × 128 matrix, 25 slices, 6 mm thick, 30% gap, four volumes) with the same center and orientation as the cardiac-gated functional images was acquired to facilitate registration of the whole-brain structural image to the functional data. Finally, a proton-density weighted image (TR/TE 6,000/16 ms, FA 149°, 0.75 × 0.75-mm^2 voxels in-plane, 1.9-mm-thick slices, echo spacing 15.6 ms) was acquired using the same slices (orientation and thickness) as the functional images for localization of the SN and VTA in the midbrain (70; cf. 34).

Data Analysis. ***Experiment 1: Whole-brain fMRI preprocessing.*** All whole-brain fMRI data were preprocessed using analysis of functional neuroimages (AFNI, 71). Functional images were first corrected for slice-timing offset and motion. The motion correction parameters were used to determine whether head movement exceeded 3 mm in any direction. Participants with head motion in excess of 3 mm were not included in any analyses. In addition, motion correction parameters were used as regressors of noninterest in a multiple linear regression analysis. Data were next spatially smoothed with a 6-mm FWHM Gaussian kernel and mean subtracted. Data from all participants were then spatially normalized by aligning them in Talairach space.

Experiment 1: Whole-brain fMRI general linear model analysis. For each participant, design matrices were created in which each experimental event was considered as an impulse stimulus that generated a hemodynamic response function of unknown amplitude. Experimental event times were convolved with a γ function to generate regressors. The set of modeled experimental events included the context-dependent cue, context-independent cue, control cue, context-dependent probe, and the context-independent probe. Context-dependent and context-independent trials in which participants made an incorrect response or responded too slowly were modeled separately from trials in which participants executed the correct response. In addition to regressors reflecting the motion correction parameters, regressors of noninterest that modeled baseline drift (scanning run mean, linear, and quadratic trends) were also included in the model. The full model was regressed to the data, with participants as a random effect, giving the best fitting amplitude values for each experimental event.

To identify areas associated with context updating, statistical maps were created contrasting the BOLD response to context-dependent vs. context-independent cues. Because numerous human neuroimaging studies using variants of our task have detected significant activations in bilateral DLPFC (e.g., refs. 11 and 12), we restricted our analysis to an a priori mask composed of Brodmann Area (BA) 9. The BA 9 mask was generated using the AFNI's preloaded maps of all BAs, was ~24,000 mm^3 in volume (~730 voxels), and was applied in Talairach space. The statistical significance of results in the DLPFC for the group analysis was determined using the AFNI program AlphaSim, which implements the cluster-size threshold procedure as a protection against type 1 error (72). On the basis of AlphaSim results, we determined that a P value of 0.05 corrected for multiple comparisons was achieved with a minimum cluster size of nine contiguous voxels ($3 \times 3 \times 3.6$ mm^3) each significant at $P < 0.008$ (two-tailed t test). The regions of DLPFC activity shown in Fig. 3*A* pass this combined threshold.

Experiment 2: spTMS. To identify targets for TMS stimulation, the fMRI data for each participant from experiment 1 were reanalyzed on an individual basis. These within-participant analyses involved a substantial diminution of statistical power. Therefore, to maximize sensitivity for detecting regions associated with context updating, we combined context-dependent and context-independent trials and contrasted these with control trials (*SI Text, Note C*). This yielded statistical maps in participant-native space that we then thresholded ($P < 10^{-6}$, uncorrected for multiple comparisons) to identify the peak coordinates of the regions of activity in DLPFC (Fig. 3 *B* and *C*). These coordinates were then used as targets for the TMS pulses.

Reaction time (RT) and accuracy were analyzed using repeated-measures ANOVAs with subject as a random effect and stimulation site (left or right DLPFC), trial type (context-dependent or context-independent), and pulse time (10, 100, 150, or 200 ms) as within-subject factors. The omnibus ANOVA was conducted first, and significant interactions were analyzed as planned contrasts using lower-level ANOVAs and paired t tests (two-tailed: $\alpha = 0.05$) to assess the statistical significance of differences between conditions. Incorrect trials were excluded for investigation of RTs. All statistical analyses were performed using JMP 7.0.2 (SAS Institute) or MATLAB 7.3 (MathWorks).

Experiment 3: Brainstem fMRI preprocessing. Before any preprocessing was carried out, data were corrected for T1 variations that occur with cardiac-gated fMRI (73) using software written in MATLAB (MathWorks). Corrected data were then preprocessed and analyzed using AFNI. To account for the variable time between image acquisitions, regressors were calculated at high temporal resolution and then resampled at the image acquisition times. Functional images were corrected for slice-timing effects using Fourier interpolation and motion corrected to the fifth volume of the first run also using Fourier interpolation. The motion correction parameters were used to determine whether participant head motion exceeded 1.5 mm in any direction. Participants with head motion in excess of 1.5 mm were not included in the analysis. In addition, motion correction parameters were used as regressors of noninterest in a multiple linear regression analysis. Next, data were spatially smoothed with a 3-mm FWHM Gaussian kernel and mean subtracted. For all participants, the most superior and inferior slices were excluded from analysis as a precaution against those slices shifting into previously nonexcited brain areas because of head motion.

The T1-weighted whole-brain structural image was aligned to the functional data and then transformed into Talairach space. All Talairached T1-weighted whole-brain structural images were subsequently brainstem-normalized using an algorithm that maximizes the overlap between participants' brainstems (74). The transforms to Talairach space and to brainstem-normalized space were then applied to the functional data.

Experiment 3: Brainstem-fMRI general linear model analysis. Data were analyzed according to the procedures described above for the whole-brain fMRI data in experiment 1. The statistical significance of results in the region of the SN and VTA was determined using the AFNI program AlphaSim, which implements the cluster-size threshold procedure as a protection against type 1 error (72). On the basis of AlphaSim results, we determined that for these data a corrected P value of 0.05 was achieved with a minimum cluster size of six contiguous voxels each significant at $P < 0.002$ (two-tailed t test). The cluster comprising the SN and VTA shown in Fig. 5*B* consists of at least nine contiguous voxels (cluster volume is 39 mm^3), each significant at $P < 0.002$.

Group results for each regression coefficient contrast were calculated in both Talairach space (for DLPFC) and brainstem-normalized space (for midbrain areas) (74). Active regions in the brainstem were visualized on brainstem-normalized anatomical images.

Event-related time courses were generated by using the "iresp" function in AFNI's 3dDeconvolve program. Onset times of cues were written into a vector, and the best-fitting impulse response function for a given time interval was estimated. BOLD time courses were generated for the response to the context-dependent and context-independent cues through the duration of the trial. Average event-related BOLD time courses (shown in Fig. 5 *C* and *D*) were generated for the SN and VTA region shown in Fig. 5*B* and for the regions of DLPFC as defined below.

Experiment 3: DLPFC region of interest analysis. Regions of interest (ROIs) within DLPFC were generated in Talairach space using AFNI and limited to AFNI's BA 9 mask as described for experiment 1 above. They were defined by drawing a sphere (~400 mm^3 in volume) around the centroid of each area of activity identified in the group analysis of the fMRI data from experiment 1 (Fig. 3*A*). Analysis was restricted to the 13 of 24 participants who had sufficient data in the DLPFC from experiment 3 (Fig. 5*A*).

Using MATLAB, correlation coefficients were computed between BOLD time series from the midbrain SN and VTA region (Fig. 5*B*) and each DLPFC ROI. To determine statistical significant across participants, correlation coefficients were converted to z statistics, and a two-tailed t test was performed. We examined two sets of correlations between the midbrain and DLPFC ROIs: (*i*) the BOLD response to context-dependent cues alone; and (*ii*) the contrast of BOLD response to context-dependent vs. context-independent cues.

Experiment 3: Reaction time analysis. For each participant we calculated the difference in mean reaction time for context-dependent and context-independent trials. We used this difference as an index of the extent to which each participant updated context in working memory in the context-dependent condition (with the context-independent condition as a control for overall differences in reaction time across participants). Using MATLAB, we then correlated these differences with an index of the gating signal in each participant: the regression coefficient of the contrast of the BOLD signal in SN and VTA for context-dependent cue > context-independent cues. A scatter plot of these reaction time differences, and the corresponding β values from the SN and VTA, are shown in Fig. 6*A*.

ACKNOWLEDGMENTS. We thank E. Barkley-Levenson for help with the whole-brain fMRI and TMS experiments; A. Conway for statistical advice on TMS data analysis; D. Knoch for helpful comments on the TMS sections of the manuscript; and K. Bartley, C. Bracero, M. Ross, and J. King for help with the brainstem fMRI experiment. This work was supported by the Regina and John Scully Center for the Neuroscience of Mind and Behavior; National Institutes of Health Grant 5R01 MH052864 and T32 MH065214 (to J.D.C.) and F32 DA027306 (to K.D.); by Princeton University's R. W. Crecca '46 senior thesis research funds (N.E.); and by The Kane Family Foundation (K.D. post-doctoral support; P. Read Montague, Principal Investigator).

1. Duncan J (1986) Disorganisation of behaviour after frontal lobe damage. *Cogn Neuropsychol* 3:271–290.
2. Shallice T (1988) *From Neuropsychology to Mental Structure* (Cambridge Univ Press, Cambridge, UK).
3. Fuster JM, Alexander GE (1971) Neuron activity related to short-term memory. *Science* 173(3997):652–654.
4. Goldman-Rakic PS, Leranth C, Williams SM, Mons N, Geffard M (1989) Dopamine synaptic complex with pyramidal neurons in primate cerebral cortex. *Proc Natl Acad Sci USA* 86(22):9015–9019.
5. Wallis JD, Anderson KC, Miller EK (2001) Single neurons in prefrontal cortex encode abstract rules. *Nature* 411(6840):953–956.
6. Cohen JD, et al. (1997) Temporal dynamics of brain activation during a working memory task. *Nature* 386(6625):604–608.
7. Cohen JD, Servan-Schreiber D (1992) Context, cortex, and dopamine: A connectionist approach to behavior and biology in schizophrenia. *Psychol Rev* 99(1):45–77.
8. Miller EK, Cohen JD (2001) An integrative theory of prefrontal cortex function. *Annu Rev Neurosci* 24:167–202.
9. Higo T, Mars RB, Boorman ED, Buch ER, Rushworth MFS (2011) Distributed and causal influence of frontal operculum in task control. *Proc Natl Acad Sci USA* 108(10): 4230–4235.
10. Zanto TP, Rubens MT, Thangavel A, Gazzaley A (2011) Causal role of the prefrontal cortex in top-down modulation of visual processing and working memory. *Nat Neurosci* 14(5):656–661.
11. Collette F, et al. (2007) Mapping the updating process: Common and specific brain activations across different versions of the running span task. *Cortex* 43(1):146–158.
12. Crone EA, Wendelken C, Donohue SE, Bunge SA (2006) Neural evidence for dissociable components of task-switching. *Cereb Cortex* 16(4):475–486.
13. Montojo CA, Courtney SM (2008) Differential neural activation for updating rule versus stimulus information in working memory. *Neuron* 59(1):173–182.
14. Roth JK, Serences JT, Courtney SM (2006) Neural system for controlling the contents of object working memory in humans. *Cereb Cortex* 16(11):1595–1603.
15. Brass M, Ullsperger M, Knoesche TR, von Cramon DY, Phillips NA (2005) Who comes first? The role of the prefrontal and parietal cortex in cognitive control. *J Cogn Neurosci* 17(9):1367–1375.
16. Astle DE, Jackson GM, Swainson R (2008) Fractionating the cognitive control required to bring about a change in task: A dense-sensor event-related potential study. *J Cogn Neurosci* 20(2):255–267.
17. Fox MD, et al. (2005) The human brain is intrinsically organized into dynamic, anticorrelated functional networks. *Proc Natl Acad Sci USA* 102(27):9673–9678.
18. Lenartowicz A, Escobedo-Quiroz R, Cohen JD (2010) Updating of context in working memory: An event-related potential study. *Cogn Affect Behav Neurosci* 10(2): 298–315.
19. Hochreiter S, Schmidhuber J (1997) Long short-term memory. *Neural Comput* 9(8): 1735–1780.
20. Todd MT, Niv Y, Cohen JD (2009) Learning to use working memory in partially observable environments through dopaminergic reinforcement. *Advances in Neural Information Processing Systems*, eds Koller D, Schuurmans D, Bengio Y, Bottou L (Curran Associates, Inc., Red Hook, NY), pp 1689–1696.
21. Braver T, Cohen J (2000) On the control of control: The role of dopamine in regulating prefrontal function and working memory. *Attention Performance* XVIII:713–737.
22. Frank MJ, Loughry B, O'Reilly RC (2001) Interactions between frontal cortex and basal ganglia in working memory: A computational model. *Cogn Affect Behav Neurosci* 1(2):137–160.
23. Rougier NP, Noelle DC, Braver TS, Cohen JD, O'Reilly RC (2005) Prefrontal cortex and flexible cognitive control: Rules without symbols. *Proc Natl Acad Sci USA* 102(20): 7338–7343.
24. Montague PR, Dayan P, Sejnowski TJ (1996) A framework for mesencephalic dopamine systems based on predictive Hebbian learning. *J Neurosci* 16(5):1936–1947.
25. Schultz W, Dayan P, Montague PR (1997) A neural substrate of prediction and reward. *Science* 275(5306):1593–1599.
26. Roesch MR, Calu DJ, Schoenbaum G (2007) Dopamine neurons encode the better option in rats deciding between differently delayed or sized rewards. *Nat Neurosci* 10 (12):1615–1624.
27. Servan-Schreiber D, Printz H, Cohen JD (1990) A network model of catecholamine effects: Gain, signal-to-noise ratio, and behavior. *Science* 249(4971):892–895.
28. Cohen JD, Braver TS, Brown JW (2002) Computational perspectives on dopamine function in prefrontal cortex. *Curr Opin Neurobiol* 12(2):223–229.
29. Seamans JK, Yang CR (2004) The principal features and mechanisms of dopamine modulation in the prefrontal cortex. *Prog Neurobiol* 74(1):1–58.
30. Dahlin E, Neely AS, Larsson A, Bäckman L, Nyberg L (2008) Transfer of learning after updating training mediated by the striatum. *Science* 320(5882):1510–1512.
31. Bäckman L, et al. (2011) Effects of working-memory training on striatal dopamine release. *Science* 333(6043):718.
32. Boehler CN, et al. (2011) Substantia nigra activity level predicts trial-to-trial adjustments in cognitive control. *J Cogn Neurosci* 23(2):362–373.
33. Boehler CN, et al. (2011) Task-load-dependent activation of dopaminergic midbrain areas in the absence of reward. *J Neurosci* 31(13):4955–4961.
34. D'Ardenne K, McClure SM, Nystrom LE, Cohen JD (2008) BOLD responses reflecting dopaminergic signals in the human ventral tegmental area. *Science* 319(5867): 1264–1267.
35. Nuechterlein KH, Dawson ME (1984) Information processing and attentional functioning in the developmental course of schizophrenic disorders. *Schizophr Bull* 10(2): 160–203.
36. Servan-Schreiber D, Cohen JD, Steingard S (1996) Schizophrenic deficits in the processing of context. A test of a theoretical model. *Arch Gen Psychiatry* 53(12): 1105–1112.
37. Beck LH, Bransome ED, Jr., Mirsky AF, Rosvold HE, Sarason I (1956) A continuous performance test of brain damage. *J Consult Psychol* 20(5):343–350.
38. Cohen JD, Braver TS, O'Reilly RC (1996) A computational approach to prefrontal cortex, cognitive control and schizophrenia: Recent developments and current challenges. *Philos Trans R Soc Lond B Biol Sci* 351(1346):1515–1527.
39. Braver TS (2012) The variable nature of cognitive control: A dual mechanisms framework. *Trends Cogn Sci* 16(2):106–113.
40. Guynn MJ (2003) A two-process model of strategic monitoring in event-based prospective memory: Activation/retrieval mode and checkin. *Int J Psychol* 38:245–256.
41. Einstein GO, et al. (2005) Multiple processes in prospective memory retrieval: Factors determining monitoring versus spontaneous retrieval. *J Exp Psychol Gen* 134(3): 327–342.
42. Chatham CH, et al. (2009) Pupillometric and behavioral markers of a developmental shift in the temporal dynamics of cognitive control. *Proc Natl Acad Sci USA* 106(14): 5529–5533.
43. Cohen JD, O'Reilly RC (1996) A preliminary theory of the interactions between prefrontal cortex and hippocampus that contribute to planning and prospective memory. *Prospective Memory: Theory and Applications*, eds Brandimonte M, Einstein GO, McDaniel MA (Erlbaum, Hillsdale, NJ), pp 267–295.
44. Mottaghy FM, Gangitano M, Krause BJ, Pascual-Leone A (2003) Chronometry of parietal and prefrontal activations in verbal working memory revealed by transcranial magnetic stimulation. *Neuroimage* 18(3):565–575.
45. Barch DM, et al. (1997) Dissociating working memory from task difficulty in human prefrontal cortex. *Neuropsychol* 13:1371–1380.
46. D'Esposito M, Ballard D, Zarahn E, Aguirre GK (2000) The role of prefrontal cortex in sensory memory and motor preparation: An event-related fMRI study. *Neuroimage* 11(5 Pt 1):400–408.
47. Walsh V, Cowey A (2000) Transcranial magnetic stimulation and cognitive neuroscience. *Nat Rev Neurosci* 1(1):73–79.
48. Rainer G, Asaad WF, Miller EK (1998) Selective representation of relevant information by neurons in the primate prefrontal cortex. *Nature* 393(6685):577–579.
49. Tomita H, Ohbayashi M, Nakahara K, Hasegawa I, Miyashita Y (1999) Top-down signal from prefrontal cortex in executive control of memory retrieval. *Nature* 401(6754): 699–703.
50. White IM, Wise SP (1999) Rule-dependent neuronal activity in the prefrontal cortex. *Exp Brain Res* 126(3):315–335.
51. Barceló F, Suwazono S, Knight RT (2000) Prefrontal modulation of visual processing in humans. *Nat Neurosci* 3(4):399–403.
52. Yago E, Duarte A, Wong T, Barceló F, Knight RT (2004) Temporal kinetics of prefrontal modulation of the extrastriate cortex during visual attention. *Cogn Affect Behav Neurosci* 4(4):609–617.
53. Gibbs SE, D'Esposito M (2006) A functional magnetic resonance imaging study of the effects of pergolide, a dopamine receptor agonist, on component processes of working memory. *Neuroscience* 139(1):359–371.
54. Robbins TW, Arnsten AF (2009) The neuropsychopharmacology of fronto-executive function: Monoaminergic modulation. *Annu Rev Neurosci* 32:267–287.
55. Ashby FG, Casale MB (2003) A model of dopamine modulated cortical activation. *Neural Netw* 16(7):973–984.
56. Williams SM, Goldman-Rakic PS (1998) Widespread origin of the primate mesofrontal dopamine system. *Cereb Cortex* 8(4):321–345.
57. Koechlin E, Ody C, Kouneiher F (2003) The architecture of cognitive control in the human prefrontal cortex. *Science* 302(5648):1181–1185.
58. Badre D (2008) Cognitive control, hierarchy, and the rostro-caudal organization of the frontal lobes. *Trends Cogn Sci* 12(5):193–200.
59. Mull BR, Seyal M (2001) Transcranial magnetic stimulation of left prefrontal cortex impairs working memory. *Clin Neurophysiol* 112(9):1672–1675.
60. Oliveri M, et al. (2001) Parieto-frontal interactions in visual-object and visual-spatial working memory: Evidence from transcranial magnetic stimulation. *Cereb Cortex* 11 (7):606–618.
61. Koch G, et al. (2005) rTMS evidence of different delay and decision processes in a fronto-parietal neuronal network activated during spatial working memory. *Neuroimage* 24(1):34–39.
62. McConnell KA, et al. (2001) The transcranial magnetic stimulation motor threshold depends on the distance from coil to underlying cortex: A replication in healthy adults comparing two methods of assessing the distance to cortex. *Biol Psychiatry* 49 (5):454–459.
63. Stewart LM, Walsh V, Rothwell JC (2001) Motor and phosphene thresholds: A transcranial magnetic stimulation correlation study. *Neuropsychologia* 39(4):415–419.

64. O'Shea J, Muggleton NG, Cowey A, Walsh V (2004) Timing of target discrimination in human frontal eye fields. *J Cogn Neurosci* 16(6):1060–1067.
65. Pitcher D, Garrido L, Walsh V, Duchaine BC (2008) Transcranial magnetic stimulation disrupts the perception and embodiment of facial expressions. *J Neurosci* 28(36):8929–8933.
66. Cattaneo Z, Silvanto J (2008) Time course of the state-dependent effect of transcranial magnetic stimulation in the TMS-adaptation paradigm. *Neurosci Lett* 443(2):82–85.
67. Silvanto J, Lavie N, Walsh V (2005) Double dissociation of V1 and V5/MT activity in visual awareness. *Cereb Cortex* 15(11):1736–1741.
68. Nichols MJ, Newsome WT (2002) Middle temporal visual area microstimulation influences veridical judgments of motion direction. *J Neurosci* 22(21):9530–9540.
69. Hung J, Driver J, Walsh V (2005) Visual selection and posterior parietal cortex: Effects of repetitive transcranial magnetic stimulation on partial report analyzed by Bundesen's theory of visual attention. *J Neurosci* 25(42):9602–9612.
70. Oikawa H, Sasaki M, Tamakawa Y, Ehara S, Tohyama K (2002) The substantia nigra in Parkinson disease: Proton density-weighted spin-echo and fast short inversion time inversion-recovery MR findings. *AJNR Am J Neuroradiol* 23(10):1747–1756.
71. Cox RW (1996) AFNI: Software for analysis and visualization of functional magnetic resonance neuroimages. *Comput Biomed Res* 29(3):162–173.
72. Forman SD, et al. (1995) Improved assessment of significant activation in functional magnetic resonance imaging (fMRI): Use of a cluster-size threshold. *Magn Reson Med* 33(5):636–647.
73. Guimaraes AR, et al. (1998) Imaging subcortical auditory activity in humans. *Hum Brain Mapp* 6(1):33–41.
74. Napadow V, Dhond R, Kennedy D, Hui KKS, Makris N (2006) Automated brainstem co-registration (ABC) for MRI. *Neuroimage* 32(3):1113–1119.
75. Bayer HM, Glimcher PW (2005) Midbrain dopamine neurons encode a quantitative reward prediction error signal. *Neuron* 47(1):129–141.
76. Hikosaka O, Takikawa Y, Kawagoe R (2000) Role of the basal ganglia in the control of purposive saccadic eye movements. *Physiol Rev* 80(3):953–978.
77. Reynolds JN, Hyland BI, Wickens JR (2001) A cellular mechanism of reward-related learning. *Nature* 413(6851):67–70.