

Ventral Striatum and the Evaluation of Memory Retrieval Strategies

David Badre¹, Sophie Lebrecht², David Pagliaccio³, Nicole M. Long⁴,
and Jason M. Scimeca¹

Abstract

Adaptive memory retrieval requires mechanisms of cognitive control that facilitate the recovery of goal-relevant information. Frontoparietal systems are known to support control of memory retrieval. However, the mechanisms by which the brain acquires, evaluates, and adapts retrieval strategies remain unknown. Here, we provide evidence that ventral striatal activation tracks the success of a retrieval strategy and correlates with subsequent reliance on that strategy. Human participants were scanned with fMRI while performing a lexical decision

task. A rule was provided that indicated the likely semantic category of a target word given the category of a preceding prime. Reliance on the rule improved decision-making, as estimated within a drift diffusion framework. Ventral striatal activation tracked the benefit that relying on the rule had on decision-making. Moreover, activation in ventral striatum correlated with a participant's subsequent reliance on the rule. Taken together, these results support a role for ventral striatum in learning and evaluating declarative retrieval strategies. ■

INTRODUCTION

The human brain is capable of efficiently recovering information from long-term memory with high utility in our current context while also minimizing the costs associated with retrieval itself (Anderson & Milson, 1989). Although artificial memory systems are superior to human memory in their ability to store vast amounts of information in a durable, content-addressable state, even the most sophisticated current search engines lack the efficiency and success of human memory when considered within the adaptive frame of this “information retrieval problem.” Nevertheless, little is currently understood about how the brain solves this problem and retrieves information useful for our current goals.

One means of solving the information retrieval problem is to develop retrieval strategies that structure or elaborate inputs to the memory system to influence the likelihood that task-relevant information is recovered. Strategic, goal-directed influence over memory retrieval could be supported by cognitive control systems that maintain a rule or context in working memory to provide a top-down bias on on-going processing (e.g., Badre & Wagner, 2007; Miller & Cohen, 2001). For example, top-down signals from frontoparietal control systems can bias perceptual and semantic systems, thereby influencing input to the memory system (Barredo, Oztekin, & Badre, in press; Badre & Wagner, 2007; Moscovitch, 1992). Likewise, control pro-

cesses can operate on the output of memory retrieval to relate recovered information to decision criteria and response contingencies (Badre & Wagner, 2007; Benjamin, 2007; Atkinson & Shiffrin, 1971). The frontal lobes are known to be broadly necessary for effective use of retrieval strategies (Gershberg & Shimamura, 1995; Stuss et al., 1994; Moscovitch, 1992). And fMRI studies have further implicated specific regions of dorsolateral PFC (DLPFC) and ventrolateral PFC (VLPFC) in networks supporting cognitive control of memory (Badre, Poldrack, Pare-Blagoev, Insler, & Wagner, 2005; Anderson et al., 2004; Dobbins, Foley, Schacter, & Wagner, 2002; Rugg & Wilding, 2000; Gabrieli, Poldrack, & Desmond, 1998).

Importantly, however, it remains an open question how memory retrieval strategies are acquired. Moreover, not all retrieval strategies are equally effective, and a previously effective strategy can become obsolete. Thus, a cognitive control system must have not only the capability to implement a retrieval strategy but also to evaluate the efficacy of that strategy given one's goals, shifting to new strategies when necessary (Becker & Lim, 2003). Currently, it is unknown how retrieval strategies are evaluated during the course of memory retrieval.

Outside of the memory domain, evidence suggests that cognitive control strategies, implemented by the PFC, are partly learned and evaluated via mechanisms of reinforcement learning supported by the nigra-striatal dopamine system (Cools, 2011; O'Reilly & Frank, 2006; Miller & Cohen, 2001). Positive or negative deviations from the expected outcome of a course of action are referred to as positive and negative reward prediction

¹Brown University, ²Carnegie Mellon University, ³Washington University, ⁴University of Pennsylvania

error (RPE), respectively, and can reinforce or punish a governing behavioral strategy (O'Reilly & Frank, 2006; O'Doherty et al., 2004; Schultz, Dayan, & Montague, 1997). In this way, these learning signals make a given strategy more or less likely to be used in the future.

A similar process of reinforcement learning has been hypothesized to support the acquisition and adjustment of memory retrieval strategies (Scimeca & Badre, 2012). However, evidence for such a mechanism during declarative memory retrieval remains limited and indirect (Schwarze, Bingel, Badre, & Sommer, 2013; Han, Huettel, Raposo, Adcock, & Dobbins, 2010; Han & Dobbins, 2009). Indeed, to date, there has been no evidence that reinforcement learning signals in the brain lead to behaviorally evident changes in reliance on a declarative retrieval strategy. Thus, this study sought to provide an initial investigation of the neural systems supporting learning of a retrieval strategy and to test the hypothesis that whereas enacting retrieval strategies requires frontoparietal cognitive control systems, the selection and evaluation of these strategies involve nigra-striatal systems that encode positive or negative RPE associated with task outcomes.

To test this hypothesis, we scanned human participants using fMRI while they took advantage of an explicit retrieval strategy that would aid their performance on a lexical decision task (LDT). During the standard LDT, participants are presented a prime word or picture and then a target letter string, separated by a temporal offset termed the SOA. Participants are asked to decide if the target letter string is a word or not. Participants are faster to respond to a word target when it is semantically related to the prime word than when it is not. This short-term semantic priming effect is thought to derive from automatic associative retrieval and is evident at both short and long SOAs (Neely, 1977, 1991; Meyer & Schvaneveldt, 1971). Importantly, however, when the SOA is sufficiently long, participants can employ expectation-based retrieval strategies to further enhance this semantic priming effect (Gold et al., 2006; Neely, 1991).

In the present experiment, we took advantage of this strategic aspect of LDT to test (a) how the brain can use expectation to guide retrieval, without support from automatic associative mechanisms, and (b) how it evaluates a retrieval strategy in terms of its utility for an upcoming decision. Specifically, we provided participants with explicit retrieval rules that indicated what category a target word would come from given the category of the prime (Figure 1). The rule was valid for most trials (75%; "Expected" trials). So, on these trials, reliance on the rule could guide retrieval of semantic features related to the target and thereby improve decision-making performance relative to "Unexpected" trials, when the rule was violated.

We further manipulated SOA between Short (50 msec) and Long (1000 msec) durations, thereby varying the opportunity participants had to complete rule-guided

control processes. The manipulation of SOA is crucial for establishing the effects of control in behavioral measures in that the impact of control is slow to develop and so takes time to express in behavior. However, as participants do not know whether a given trial will involve a Long or Short SOA, we do not anticipate that the engagement of control will differ between these two trial types, which is of primary relevance to the brain measures. Thus, blocks on which a rule was available ("Rule" blocks) were interleaved with blocks for which there was no rule ("Neutral" blocks). Thus, the inclusion of a Neutral condition permitted us to identify regions generally involved in guiding retrieval according to the rule.

Finally, prior work has indicated that it is important to distinguish whether LDT performance is affected by factors related to the decision process itself versus other non-decision components, like cue encoding (Ratcliff, Gomez, & McKoon, 2004). Reliance on the retrieval strategy in the present task could hypothetically affect multiple components. Thus, we sought to distinguish the component process in the lexical decision using the drift diffusion model (DDM) framework (Ratcliff & McKoon, 2008; Ratcliff, 1978). The DDM is a unified means of relating RT and error rates to interpretable component processes during binary decisions, like in the LDT. The DDM explains the decision process as arising from noisy evidence accumulation or drift toward one of two response thresholds. Applied to a Word/Nonword decision in the LDT, evidence is any information recovered from memory that supports an item being a word. Given the generality of the Word/Nonword decision, this evidence could range from semantic information about the target to its lexical entry to its phonotactic structure. The amount, degree, and ease with which this evidence is recovered would affect the strength of evidence. Stronger evidence accumulation during the decision is reflected in a larger drift rate parameter. Prior modeling using the DDM framework has implicated differences drift rate in accounting for a number of major behavioral phenomena during LDT, such as effects of repetition and word frequency. However, there are a number of other parameters in the DDM that affect RT and errors, such as the variability in the drift rate itself, the starting position or bias toward one or the other boundary, the separation between the boundaries, and an initial non-decision time that is thought to reflect processes related to item encoding occurring before the decision process. Some of these other parameters would not be expected to relate to Expectation manipulations in the present experiment, as Expectation was independent of the Word/Nonword decision. However, others, such as non-decision time, have been shown previously to improve fits to behavior in the LDT (Donkin, Heathcote, Brown, & Andrews, 2009; Ratcliff, Thapar, Gomez, & McKoon, 2004) and could relate to reliance on the retrieval rule. In this study, we used the DDM to (a) parameterize the degree to which reliance on the retrieval rule affected these components of the lexical decision and (b) to assess the

relationship between individual differences in these decision components and activation in striatum.

METHODS

Participants

Seventeen (12 women) right-handed adults (age = 19–32 years, mean = 22 years) with normal or corrected-to-normal vision were recruited. All participants were without psychiatric or neurological conditions, contraindications for MRI, or medication affecting the CNS. Three additional participants were recruited, but their data were excluded before analysis because they either had their MRI session interrupted (one participant), failed to complete the task (one participant), or had excessive head movement (>3 mm; one participant). All participants gave written informed consent and were compensated for participation according to guidelines established and approved by the Institutional Review Board of the Research Protections Office at Brown University. Participants were compensated \$15/hr.

Stimuli

Stimuli were object pictures or English orthography letter strings that either spelled real English words (Word) or nonwords (Nonword). Sixteen semantic categories were used in the experiment. Each category contained 30 images and 15 words. Individual exemplars were not included as both a word and picture. For example, in the category “bird,” the exemplar “robin” was included either as an image of a robin or the word “robin,” but not both. Thus, 45 exemplars were included for each category across words and pictures.

Picture stimuli depicted single, nameable real-world objects and were 400 × 400 pixels in size subtending 3° of visual angle. Four hundred eighty pictures were used in the experiment (i.e., only during Rule blocks) drawn from the Internet and were selected such that there were 30 pictures in each of 16 object categories. Eight additional pictures taken from two semantic categories were presented in the practice section of the experiment. The category membership of each object pictured was established in a preexperimental category norming pilot, and representativeness was equated across categories. A visual noise stimulus with the same dimensions as the picture stimuli was used during the Neutral blocks.

A total of 600 lexical items (Words + Nonwords) were used in the experiment. Two hundred forty Words and an equivalent number of Nonwords were used during Rule blocks. Likewise, 60 Words and Nonwords were used during Neutral blocks. Nonwords were selected at random from the ARC Nonword Database (Rastle, Harrington, & Coltheart, 2002; www.maccs.mq.edu.au/~nwdb/). Non-

word items were balanced for length across experimental conditions. Words were selected as representative of 1 of the 16 semantic categories used for the picture sets (15 Words per category), although they named objects that differed from the exemplars featured in the picture set. The Neutral Word stimuli came from categories other than those represented in the picture set. Words were counter-balanced across experimental conditions for length, syllables, and frequency of use in the English language (Kucera & Francis, 1967).

All words were presented in black Helvetica font, 30 pt., in the center of a white screen using Matlab (Mathworks, Inc., Natick, MA) and Psychtoolbox (Brainard, 1997) displayed with an InFocus IN34 DLP projector (1024 × 768 resolution).

Logic and Design

Rule-guided retrieval was tested using a modified form of the LDT priming procedure (Figure 1; Neely, 1977, 1991; Favreau & Segalowitz, 1983). Each trial of the experiment began with presentation of a picture prime for 200 msec followed by a black fixation cross and then a target letter string that was either a Word or a Nonword with equal frequency. The period between offset of the prime and the onset of the target (SOA) was either 50 msec (Short) or 1000 msec (Long). The target letter string was on the screen for 500 msec and then was replaced by a green fixation cross for a maximum of 1000 msec. Participants were required to decide whether the target letter string was a Word or Nonword and indicated their response

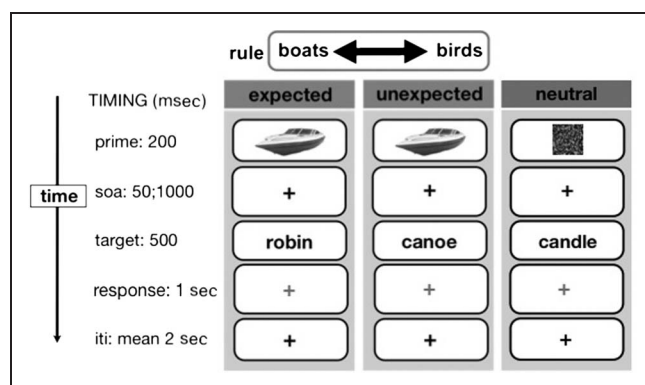


Figure 1. Schematic of task events and conditions in the LDT. Each column depicts a time line of trial events for each condition of the experiment. During rule blocks, participants were provided a rule (top) that related two arbitrary semantic categories (e.g., boats and birds). On all trials (bottom), a prime image was presented followed by an SOA and then a target word string, which the participant identified as a word or nonword. For Expected trials (left), the category of the prime image predicted the category of the target word in accord with the rule. On occasional Unexpected trials (middle; 25% of Rule trials), the target came from the same category as the prime, which violated the rule. During Neutral blocks (right), trial events followed the same structure as for Rule blocks, except that there was no rule and the prime was a visual noise image.

using a key press response with the index or middle finger of their left hand. Participants could respond at any point while either the letter string or the green fixation cross was on the screen, resulting in a 1500-msec response deadline. Once a response was made or the 1500 msec had expired, the fixation cross turned to red and no further response was recorded. Trials were separated by a variable intertrial interval (mean = 2 sec) that was determined by an algorithm (optseq2; surfer.nmr.mgh.harvard.edu/optseq/) that optimizes the efficiency of the design for event-related fMRI analysis.

Trials were grouped into blocks. For Rule blocks, participants were provided a single retrieval rule for each block that related two unrelated semantic categories. The rule indicated that when the prime depicted an object from a particular category one could expect a word from the paired category. For example, consider the rule, “Bird \longleftrightarrow Building.” In this case, if a picture of a bird was presented as the prime, the participant could expect that an upcoming Word target would name a type of building. Rules were bidirectional (i.e., bird pictures cued building names with equal frequency as building pictures cued bird names) and conveyed no information about the likelihood of an upcoming Word versus Nonword response (i.e., both responses were 50% likely).

Throughout a Rule block, Word targets conformed to the rule on 75% of trials. These were termed Expected trials. On the remaining 25% of Word trials, the target word was from the category of the picture prime. Hence, on these Unexpected trials, the target word was semantically related to the prime but violated the retrieval rule. Participants were instructed that the rule would apply most of the time. Thus, they were aware that there would be violations, but they were not told exactly how often.

Importantly, this Expectation manipulation permitted a test of reliance on the retrieval rule. By relying on the rule, participants could start retrieving information about features of the upcoming target before its presentation, thereby facilitating item-level retrieval for category members and allowing for faster and more accurate Word/Nonword discrimination. The pairing between the prime and target categories was arbitrary. Moreover, any benefit from following the retrieval rule cannot be accounted for based on an effect of priming, passive spreading activation or automatic retrieval, as rule expectancy is pitted against prior semantic association (i.e., the Unexpected targets are semantically related to the primes). Finally, a new rule relating two novel semantic categories was provided for each block, diminishing the influence of associative learning and requiring cognitive control to direct retrieval to the new cued category. A further implication of this manipulation is that any learning over blocks must be at an abstract level, as opposed to between two specific categories.

Unexpected targets should result in conflict, relative to Expected targets because of the prior retrieval of features specified by the rule. This conflict could affect retrieval and/

or postretrieval decision processes because of (a) retrieval interference from the retrieved but now-irrelevant features of the cued category, (b) the initial interpretation of unexpected category information as evidence for a Nonword response, or (c) the suppression of features related to the target category that occurred following presentation of the prime. Moreover, as control processes take longer to unfold, this conflict effect should emerge at the Long SOA, specifically. At the Short SOA when participants would have insufficient time to implement the category rule, responses to Unexpected targets may not differ from Expected targets or even may be facilitated because of automatic semantic priming.

Neutral blocks were interleaved with Rule blocks. During Neutral blocks, there was no retrieval rule. The trial structure was the same as that for Rule blocks, except that, instead of a picture prime, a static noise pattern was presented for 200 msec. A noise pattern rather than an object was used to minimize the potential for any semantic relatedness or retrieval strategy to influence performance on these trials. Thus, there were no Expected or Unexpected trials during Neutral blocks. The Word/Nonword and SOA manipulations were conducted exactly as in Rule blocks.

Over the course of the experiment, participants performed a total of eight Rule blocks (60 trials/block) and four Neutral blocks (30 trials/block). Two Rule blocks were scanned per run and alternated with runs featuring a Neutral block. A new retrieval rule was presented for 10 sec at the beginning of each Rule block (followed by a 10-sec interval before the first trial began). No rule was repeated within a given participant.

Behavioral and Computational Modeling Analysis

Behavioral data were RTs and error rates across experimental conditions. RT was measured from the onset of the target letter string until the response. Failure to respond before the 1500-msec deadline was coded as a non-response, and these trials were excluded from calculations of error rate or RT. Behavioral effects were assessed with a Rule [Rule/Neutral] \times Response [Word/Nonword] \times SOA [Long/Short] ANOVA across all experimental trials and an Expectation [Expected/Unexpected] \times SOA [Long/Short] ANOVA within the Rule blocks. Basic RT analysis was restricted to correct trials. To provide a posterior predictive check for the model, RT distributions and quantiles were derived using the empirical (i.e., Kaplan–Meier) cumulative distribution function for each condition.

The lexical decision was modeled as a drift diffusion process, following prior work (Ratcliff, Gomez, et al., 2004). In this framework, evidence accumulates noisily over time, as estimated by drift rate (ν), toward one of two decision bounds (i.e., Word or Nonword). The accumulation process itself begins after an initial non-decision interval (t), thought to reflect processes like encoding of the target. In addition, the model includes parameters for

thresholds used for display purposes are listed in the figure captions.

Whole-brain analyses were complemented by ROI analyses to test predicted effects in a priori hypothesized regions. Where possible, a priori ROIs were taken from prior studies of memory retrieval and learning that were representative of a putative function of interest. ROIs that we did not have a particular prior effect or representative study to use were functionally defined based on all significant voxels within an 8-mm radius of a chosen maximum from the unbiased contrast of all correct trials versus fixation. The maximum was chosen as the local peak *t* value within the anatomical area of interest.

We defined four ROIs based on the prior literature. First, we constructed a bilateral ROI in ventral striatum that encompassed ventral caudate and nucleus accumbens using the conjunction of 8-mm spheres around three peak foci reported in two published studies of reward processing (O'Doherty et al., 2004 [*xyz* = 14, 10, -10; *xyz* = 6, 14, 2]; Bray & O'Doherty, 2007 [*xyz* = -9, 15, -3]) along with their contralateral hemisphere homologues (ROI mask is shown in Figures 5 and 6). To ensure that our ventral striatum ROI could be considered representative of the literature on reinforcement learning, we used the Neurosynth database (neurosynth.org; Yarkoni, Poldrack, Nichols, Van Essen, & Wager, 2011) to assess whether our ROI is representative of ventral striatum as activated during studies of reward and reinforcement learning. Neurosynth is a large-scale meta-analytic database for functional neuroimaging data. This analysis yielded posterior probabilities for the peak foci of the ROIs ranging from .81 to .91, supporting the representativeness of this ROI. Second, in VLPFC, we defined ROIs in anterior [*xyz* = -47 30 -6] and mid-VLPFC [*xyz* = -50 25 14] defined from a prior study of semantic retrieval that functionally dissociated these subregions in terms of controlled retrieval and postretrieval selection (Badre et al., 2005). Finally, in parietal cortex, we defined an ROI in angular gyrus (AG) [*xyz* = -56 -54 36] based on a study of expectation violation in recognition memory (O'Connor, Han, & Dobbins, 2010). O'Connor et al. (2010) is among the few studies of expectation violation in memory, analogous to the present manipulation, and so provides a functional definition within inferior parietal cortex that we wish to test in the present manuscript. The remaining ROIs in DLPFC (-47 26 35) and IPS (-32 -56 40) were defined from all correct trials versus fixation, as discussed above.

Estimates of the shape of the HRF for each condition as a function of peristimulus time was calculated using a finite impulse response model implemented in Marsbar (marsbar.sourceforge.net/). The signal change from each condition was extracted, and the peak percent signal change was calculated as the integral of the peak time point within condition \pm one time point. The resulting integrated percent signal change estimates were subjected to repeated-measures analyses of variance, *t* tests, and linear regression as noted in the Results.

Brain-behavior correlations were considered significant at $p < .05$.

RESULTS

Behavioral Results

Behavioral performance indicated that participants used the rule to proactively guide retrieval and aid their decisions. Participants performed with a high degree of accuracy overall (mean commission error rate = 4%) while making most of their responses before the response deadline (mean omission error rate = 3%). Error rates were comparable to prior similar versions of the LDT (Ratcliff, Gomez, et al., 2004; Favreau & Segalowitz, 1983). Table 1 lists the mean correct trial RT and commission error rates in each condition of the experiment. In general, responses to Nonwords were slower (713 msec) than Words (669 msec), $F(1, 16) = 28.3, p < .0001$. Rule blocks (735 msec) were slower overall and were associated with more errors (4%) than Neutral blocks (647 msec; 1% error) $F_s(1, 16) > 13.6, p_s < .005$. This overall slowing for Rule relative to Neutral likely reflects the added working memory load associated with maintaining the rule. Finally, there was a main effect of SOA such that RT was faster and error rates lower for Long relative to Short SOAs, $F(1, 16) = 39.7, p < .0001$.

Consistent with reliance on the retrieval rule, Expected trials were facilitated relative to Unexpected trials (Figure 2A; Table 1), $F(1, 16) = 26.6, p < .0005$. Moreover, when crossed with SOA, this effect was only evident at the long SOA (SOA \times Expectation: $F(1, 16) = 39.1, p < .0001$; Unexpected-Expected: Long SOA: $F(1, 16) = 93.9, p < .0001$; Short SOA: $F = .42$), supporting the contribution of a slower cognitive control process.

Finally, as there has been substantial work in the literature on the LDT focusing on distributional effects (e.g., Ratcliff, Gomez, et al., 2004; Balota & Spieler, 1999) and because the DDM fits depend on the full RT distribution, we analyzed the effect of Expectation on RT distributions by computing RT quantiles (Figure 2B; Ratcliff, Gomez, et al., 2004; Balota & Spieler, 1999). Comparison

Table 1. Mean Correct Trial RT (msec) and Error Rates

SOA	Rule Block			Neutral Block	
	Expected	Unexpected	Nonword	Word	Nonword
Short					
RT	741	749	793	627	686
Error	4%	5%	3%	1%	2%
Long					
RT	653	742	702	604	669
Error	4%	9%	3%	2%	2%

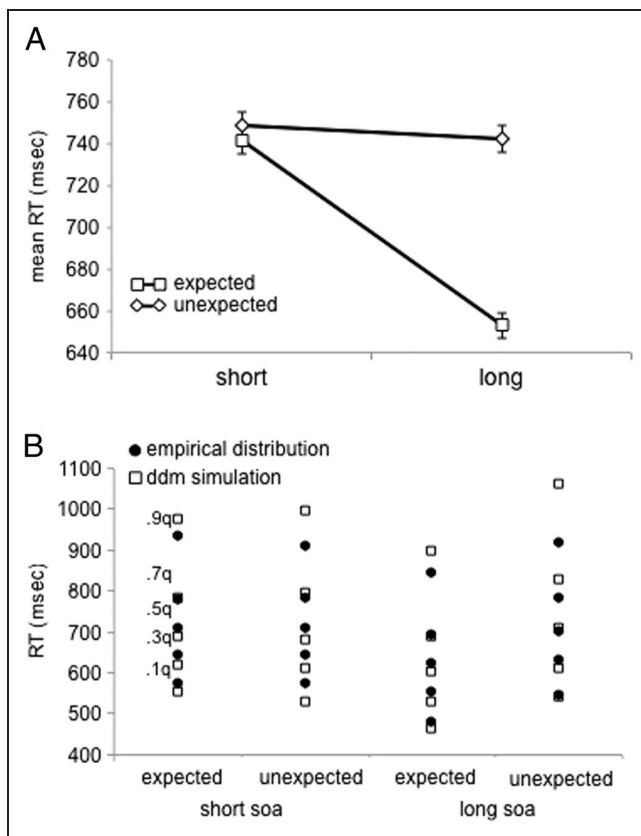


Figure 2. Behavioral and model simulation results from the LDT. (A) Expected and Unexpected conditions showed an interaction with SOA on RT such that participants were speeded for Expected conditions at the Long SOA. Error bars depict within subject standard error. (B) Five RT quantiles (.1, .3, .5, .7, .9) from the empirical RT distributions are plotted for the Expectation and SOA conditions (solid circles). Simulated RTs from the best fitting DDM model are plotted (open squares) for comparison. This model allowed drift rate and non-decision time to vary as a function of experimental condition. One thousand trials were simulated for each condition.

of the .1 and .9 quantiles of the correct trial RT cumulative distribution functions indicated that the interaction of SOA with Expectation shifted the entire RT distribution earlier for Expected trials, evident both at the lead edge, $F(1, 16) = 26.6, p < .0001$, and the long tail, $F(1, 16) = 29.7, p < .0001$, with no interaction between the two, $F = 1.7$. These behavioral effects across conditions provided constraints on the fit of the DDM model.

DDM Results

Behavior (accuracy and RT across both error and correct trials) was best fit by a DDM model allowing drift rate (ν) and non-decision time (t) to vary as a function of Expectation condition (Figure 2B; Table 2). This is consistent with a prior DDM investigation of the LDT that found inclusion of non-decision along with drift rate fit the lead of RT distribution across word frequency conditions better than including drift rate alone (Donkin et al., 2009). Monte Carlo error was less than

1% of the standard deviation of the posterior distribution for all parameters, providing evidence of convergence.

Parameter estimates from the best fitting model are shown in Table 3. Expected and Unexpected conditions expressed their effects on drift rate and non-decision time estimates. Unexpected trials showed a slowing of drift rate relative to Expected, $F(1, 16) = 10.6, p < .01$, and Neutral trials, $F(1, 16) = 104.3, p < .001$, suggesting that the violation of expectation on these trials impeded the Word/Nonword decision process. Notably, although the difference in drift rate between Expected and Unexpected conditions was quantitatively larger at the Long SOA, the Expectation by SOA interaction was only marginal, $F(1, 16) = 3.1, p = .096$.

Expectation also affected non-decision time given a sufficiently long SOA. Specifically, non-decision time for Expected trials was estimated slower than both Unexpected and Neutral at the Short SOA, but following a Long SOA, non-decision time for Expected trials was estimated to be faster than Unexpected or Neutral conditions. This cross-over was supported by reliable Expectation by SOA interactions for Expected relative to Unexpected, $F(1, 16) = 24.0, p < .001$, and Expected relative to Neutral items, $F(1, 16) = 9.6, p < .01$. Thus, the facilitative effect of Expectation on non-decision time suggests that participants relied on the retrieval rule in advance of the target, thereby speeding encoding of the target when it was consistent with the rule and sufficient time was provided. Unexpected showed a slower non-decision than Neutral, $F(1, 16) = 5.7, p < .05$. There was no interaction of this effect with SOA, $F(1, 16) = 1.6$.

To assess whether the drift rate and non-decision time estimates were related across participants, we correlated the drift rate and non-decision time parameters. Collapsing across SOA, this correlation was not significant for Unexpected ($R = .01, p = .9$) and marginal for Expected ($R = .47, p = .06$). For the critical Expected Long and Unexpected Long conditions, neither correlation was reliable ($R_s < .16, p_s > .5$). Thus, we do not find strong

Table 2. Comparison among Alternative Models

Model	DIC	Deviance	pD
Null model	7462	7414	48.2
a only	2226	2033	192.7
t only	685.9	483	202.0
z only	-4365	-4375	9.6
ν only	-7197	-7414	217.5
ν and a	-9299	-9644	344.3
ν and z	-9303	-9574	271.0
ν and t^a	-9524	-9877	352.3

^aBest fitting model.

Table 4. fMRI Activations from Major Contrasts (FWE Cluster Corrected $p < .05$)

Region	Stereotaxic Coordinates			~Brodmann's Area	Peak Z
	x	y	z		
<i>Rule Minus Neutral</i>					
Left DLPFC	-40	22	16	9, 46	4.0
	-42	20	30	9, 46	3.6
	-50	18	34	9, 46	3.5
Right lateral occipital	36	-78	8		4.5
	44	-74	-6		4.3
	32	-80	-2		4.2
Left lateral occipital	-44	-72	-12		4.5
	-32	-74	-14		4.3
	-44	-74	-20		4.2
Right cerebellum	30	-50	-24		4.3
	34	-38	-32		4.3
	38	-46	-42		3.2
Right IPS	22	-92	36		3.9
	30	-90	26		3.6
	30	-80	30		3.3
<i>Unexpected Minus Expected</i>					
Left VLPFC	-46	22	-2	47/45	4.0
	-44	28	4	47/45	3.6
Precuneus	-10	-62	38		4.6
	8	-68	36		3.7
	0	-68	38		3.6
Left SMA	-14	8	70		4.4
	-18	6	60		4.0
	-10	10	62		4.0
Left DLPFC	-38	20	32	9, 46	4.0
	-40	22	40	9, 46	3.4
	-50	12	46	9, 46	4.0
Right DLPFC	42	14	40	9, 46	3.7
	50	10	52	9, 46	4.0
Rostral medial frontal cortex	14	44	32		3.7
	20	54	34		3.6
	8	40	28		3.4
Left angular gyrus	-56	-54	36		3.6
	-48	-52	40		3.6

Table 4. (continued)

Region	Stereotaxic Coordinates			~Brodmann's Area	Peak Z
	x	y	z		
<i>Expected Minus Unexpected</i>					
Right superior temporal cortex	58	-22	8		4.3
	50	-22	0		3.9
	66	-30	16		3.7
Right insula	52	-10	6		4.2
Right posterior putamen ^a	30	-12	2		4.2
Right nucleus accumbens ^a	14	18	-8		3.6

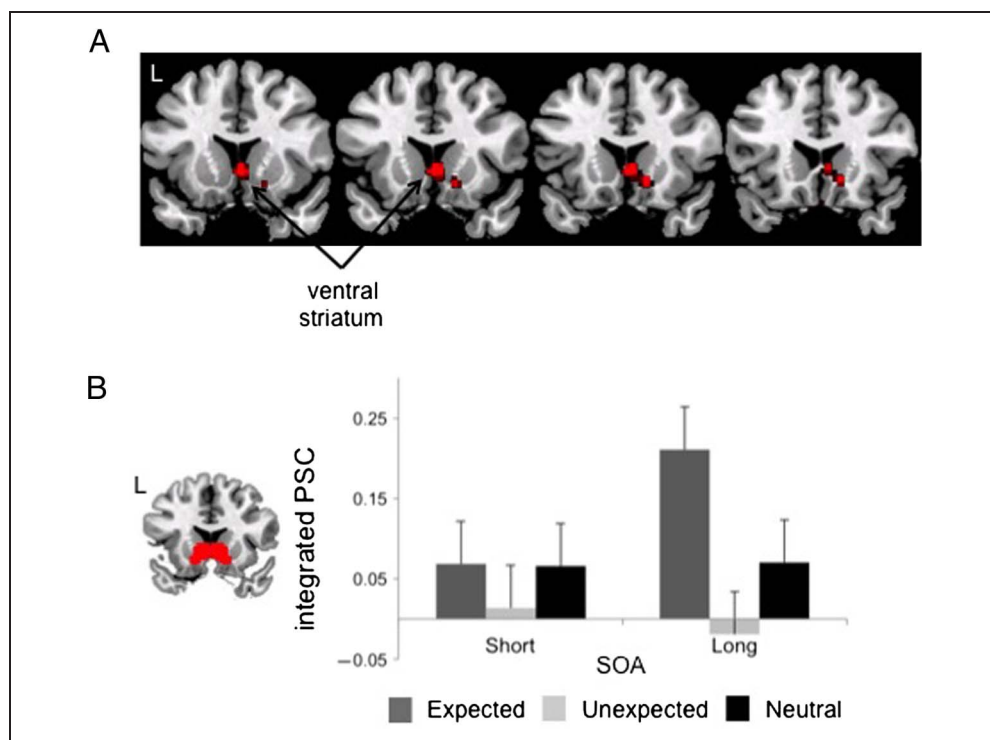
^aSmall volume FDR height corrected ($p < .05$).

with activation on Neutral Long trials ($R = .08, p = .75$) was not significant, and this difference from the correlation with Long Expected activation was supported in a reliable Expectation (Expected/Neutral) by effect interaction ($p < .05$). The correlation of Unexpected Long activation with drift rate was also unreliable ($R = .22, p = .38$). However, in this case, the Expectation (Expected/Unexpected) by effect interaction was not significant ($p = .23$), and so we could not rule out a positive correlation between drift rate and Unexpected Long activation. Although also showing a positive trend, the correlation between Expected Long trials and non-decision time was not significant ($R = .35, p = .31$). Thus, these analyses provide evidence that ventral striatal activation was related to the degree to which

the lexical decision was made easier following adherence to the retrieval rule. We next sought to test the degree to which activity in ventral striatum was related to learning, the hallmark of an RPE.

If activation in ventral striatum reflects a learning signal related to the retrieval rule, then individual differences in this response should also account for learning-related changes in retrieval strategy. To test this, we first calculated the Expectation Effect on drift rate as the Expected drift rate minus Unexpected drift rate for Long SOA trials. This Expectation Effect was computed separately for each half of the experiment. We then operationalized the learning effect as an increase in this Expectation Effect from the first to the second halves of the experiment

Figure 3. Effects of Expectation in ventral striatum. (A) The whole-brain voxel-wise contrast of Expected > Unexpected is plotted on coronal slices. To show the spread of activation, the contrast is plotted at $p < .005$ uncorrected. However, the peaks of activation in ventral striatum correct for multiple comparison over the whole striatal volume. (B) A large ROI covering ventral caudate and nucleus accumbens was constructed from previous studies of reinforcement learning and is shown at left in a coronal slice. Bar plot depicts the percent signal change integrated over a 4- to 10-sec window following presentation of the prime stimulus across Rule and Neutral conditions by SOA. Error bars show within-subject standard error.



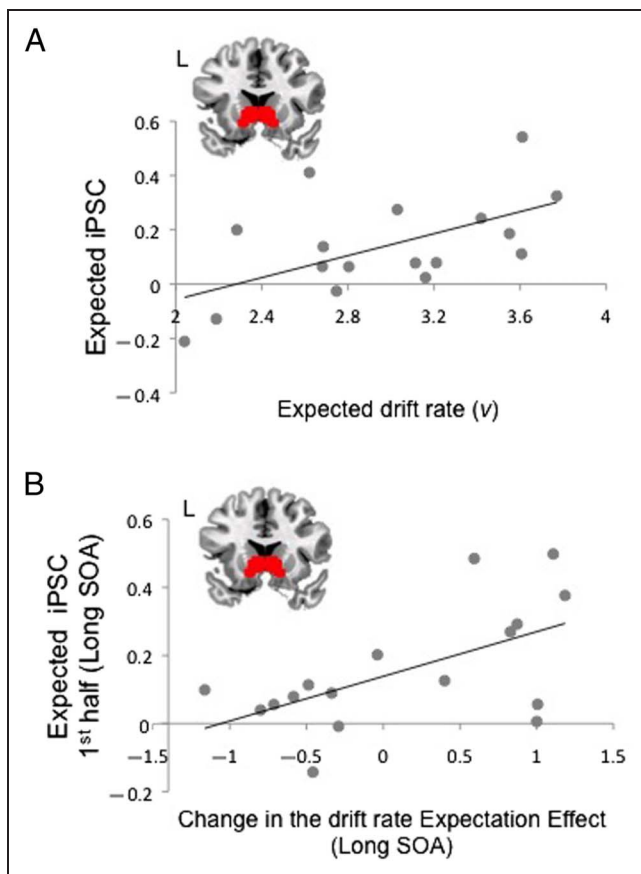


Figure 4. Correlations of ventral striatal activation with individual differences in drift rate. (A) Activation in the ventral striatum was positively correlated with individual differences in drift rate. (B) Activation in ventral striatum in the first half of the experiment was correlated with the shift in the Unexpected versus Expected difference in drift rate between the first and second halves of the experiment. The ventral striatal ROI, covering ventral caudate and nucleus accumbens, is shown on each plot.

(i.e., [Expectation Effect 2nd half] minus [Expectation Effect 1st half]). In other words, a growth in the Expectation Effect over the course of the experiment indicated greater reliance on the rule and so larger drift rates for Expected trials and smaller ones for Unexpected trials. We computed learning for non-decision time in a similar way, except that we calculated the Expectation Effect as the Unexpected non-decision time minus the Expected non-decision time because greater reliance on the rule would be reflected in slower non-decision time for Unexpected trials and faster for Expected trials.

Across participants, there was no overall trend in the change in Expectation Effects from the first to the second halves of the experiment, either in drift rate (mean v shift = .13) or in non-decision time (mean T shift = .0002), with some participants shifting toward a larger Expectation Effect and others toward a smaller one by the end of the experiment. Critically, however, activation in ventral striatum on Expected Long trials during the first half of the experiment positively accounted for individual differences

in the change in the drift rate Expectation Effect from the first to the second half of the experiment ($R = .58, p < .05$; Figure 4B). This correlation held ($R = .64, p < .01$) even when we first removed any correlation of the shift in v with activation from Expected Long trials in the second half of the experiment. As the activation in the second half of the experiment would follow any shift in the Expectation Effect, this control should remove any variance in the Expected Long activation that is not attributable to a learning effect. Thus, those participants showing greater activation in ventral striatum for Expected Long trials during the first half of the experiment showed a greater increase in the Expectation Effect in drift rate in the second half of the experiment. There were no reliable effects of ventral striatal Expected Long activation on the shift in non-decision time ($R = .1, p = .7$). Hence, as with the basic brain-behavior correlations, the learning effects related to ventral striatum were most evident in the drift rate.

To the degree that participants increase their reliance on the cue over the course of the experiment based on striatal signals, then one should predict that the difference in drift rate between Rule and Neutral trials should also correlate with the striatal signals. Specifically, to the degree that participants learn about cue reliability over the course of the experiment, then the difference in drift rate between Expected and Unexpected should relate to the difference between Neutral and Rule events. To test this, we correlated the change in Expected minus Unexpected drift rates between the first and second halves of the experiment with the change in Neutral minus Rule (Long) drift rates, we found the two to be correlated ($R = .63, p < .01$). Moreover, striatal activation on Expected Long trials during the first half of the experiment should positively correlate with the change in the difference between the Neutral and Rule drift rates. In other words, the more positive RPE that is experienced in the first half of the experiment, the more that cue reliance is reinforced as a strategy. This should translate into a larger Neutral versus Rule difference in the second half of the experiment. This correlation was also reliable ($R = .59, p < .05$). As with the correlation with the Unexpected vs. Expected difference, Expected Long activation in the second half was not correlated with the change in Neutral versus Rule ($R = .06, p = .8$), nor was the activation for Neutral Long trials themselves either in the first half ($R = .001, p = .9$) or over the entire experiment ($R = .38, p = .14$), suggesting that this effect was related to reliance on the rule as rewarded during the Rule blocks.

To summarize, activation in ventral striatum tracked the efficiency of decision-making following application of the retrieval rule. Moreover, individual differences in activation in this region during Expected trials predicted subsequent reliance on the retrieval rule, providing evidence of a learning effect. Hence, these results are consistent with the hypothesis that ventral striatum is a component of the system that selects, evaluates, and adapts retrieval strategies. In contrast to striatum, frontoparietal control

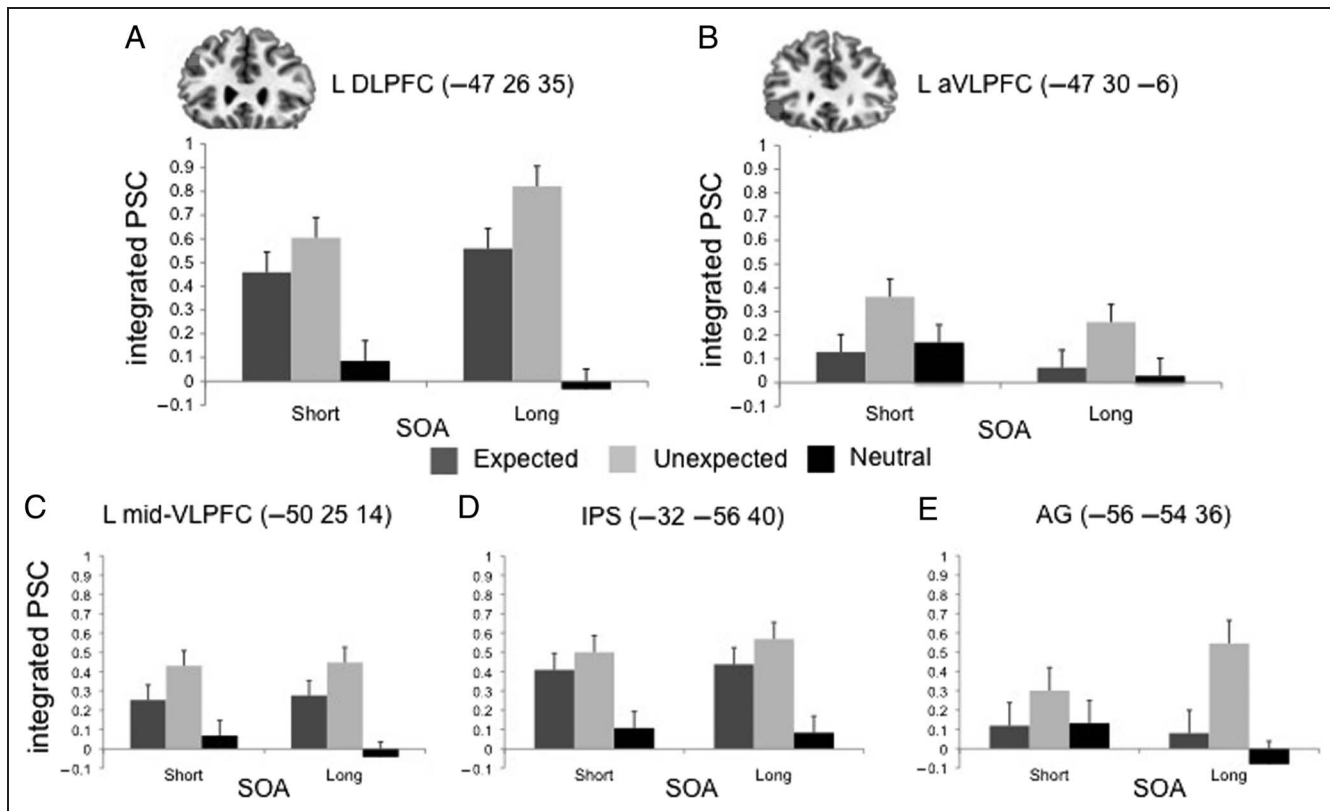


Figure 6. Plots of percent signal change in ROIs in prefrontal and parietal cortex. Bar plots depict percent signal change integrated over a 4- to 10-sec window following presentation of the prime stimulus. Plots depict Expected (dark gray), Unexpected (light gray), and Neutral (black) across SOA conditions. Results from ROIs in (A) DLPFC, (D) mid-VLPFC, and (C) IPS show similar patterns of Unexpected greater than Expected, with both Rule conditions greater than Neutral. By contrast, (B) aVLPFC and (E) AG show greater activation for Unexpected than Expected and Neutral. Error bars depict within subject standard error.

activation on Expected trials did not differ from Neutral items ($F = .06$), whereas Unexpected activation was greater than Neutral, $F(1, 16) = 4.6, p < .05$. The relative differences of the Expected and Unexpected trials from Neutral differentiated the ROIs in VLPFC and DLPFC, as was evident in a region [IVLPFC, IDLPFC] by effect [Expected, Unexpected, Neutral] interaction, $F(2, 32) = 7.3, p < .005$.

Prior work has suggested that the anterior portion of VLPFC highlighted in whole-brain and ROI analysis here is functionally dissociable from more caudal and dorsal portion of mid-VLPFC, corresponding approximately to inferior frontal gyrus pars triangularis (BA 45; Badre & Wagner, 2007). Thus, we next tested an ROI in this mid-VLPFC region based on a prior fMRI experiment (Badre et al., 2005). As shown in Figure 6C, activation in this region was unlike aVLPFC and similar to DLPFC, with both Expected and Unexpected conditions being more activated than Neutral ($F_s > 6.8, p_s < .05$) and Unexpected showing marginally greater activation than Expected, $F(1, 16) = 3.3, p = .08$. The difference in activation across Expected, Unexpected, and Neutral conditions between aVLPFC and mid-VLPFC resulted in a region by effect interaction, $F(2, 32) = 4.4, p < .05$.

Finally, we further explored the effects of expectation violation in AG (-51 -66 39) and IPS (-32 -56 40; Figure 6D, E). AG showed greater activation for Unexpected

than Expected trials, $F(1, 16) = 4.2, p < .05$. Moreover, like aVLPFC, this effect was driven by an increase in Unexpected relative to Neutral, $F(1, 16) = 6.4, p < .05$, whereas Expected and Neutral did not differ ($F = .2$). In contrast to AG, IPS showed a pattern similar to that observed for DLPFC, with greater activation for both Expected and Unexpected trials relative to Neutral ($F_s > 11.9, p_s < .01$). Although Unexpected trials also showed quantitatively greater activation than Expected, akin to that observed in DLPFC, this effect was not reliable in IPS ($F = 1.4$).

Thus, these analyses identify two networks of frontal and parietal neocortical regions during the application of strategy to memory retrieval and when compensating for violations of expectation. One network, including DLPFC, mid-VLPFC, and IPS, was generally more activated on Rule relative to Neutral blocks and for Unexpected than Expected trials. The other network, including aVLPFC and AG, was more activated for Unexpected than Expected trials but did not further differentiate Expected from Neutral conditions.

Brain-Behavior Correlations in PFC and PPC

Finally, we tested the relationship between activation patterns observed in four frontal and parietal ROIs—left

aVLPFC, left DLPFC, IPS, and AG—and individual differences in non-decision and drift rate parameters from the DDM. Left DLPFC and IPS did not correlate with any of the individual difference measures or overall RT. aVLPFC reliably correlated with drift rate on Expected Long trials ($R = .58, p < .05$). However, there was no evidence of a learning effect on drift rate in left aVLPFC ($R = .05, p = .86$). Expected Long activation in AG also showed a marginal correlation with drift rate ($R = .48, p = .053$).

Finally, to further address a potential RT confound in the present results, we tested the correlation of the RT difference between participants with the observed activation differences for the (1) Expected > Unexpected contrast in ventral striatum, (2) the Rule > Neutral contrast in DLPFC, and (3) the Unexpected > Expected contrast in DLPFC. There was no evidence of a correlation between the Expected versus Unexpected RT differences and the corresponding difference in activation in ventral striatum ($R = .006$). Likewise, there was no correlation between the Rule versus Neutral trial RT difference and the corresponding activation difference in DLPFC ($R = .02$). There was a marginal correlation between the Unexpected versus Expected RT difference and this contrast in DLPFC ($R = .44, p = .07$). However, the fact that this correlation is condition specific (i.e., marginally evident in the Unexpected vs. Expected case, but not in the Rule vs. Neutral case), despite the clear difficulty differences across these conditions argues against a global difficulty account for this DLPFC activation (see Discussion for additional elaboration of this point).

DISCUSSION

To be adaptive, it is crucial that our memory system accurately and precisely recovers information that is useful for achieving our behavioral goals. This study provides evidence that the ventral striatum may track the efficacy of retrieval strategies in terms of their impact on decision-making and goal attainment. Importantly, based on these signals, the control system may adjust its future reliance on the governing strategy. Thus, our results provide initial support for the hypothesis that memory retrieval strategies may be partly acquired and adjusted through striatal-dependent reinforcement learning mechanisms, akin to what is observed in the action and working memory domains (Scimeca & Badre, 2012).

Using a rule that related the category of a prime to the likely category of an upcoming target, participants could start retrieving information about features of the upcoming target before its presentation, thereby facilitating item-level retrieval for category members. As the relationship between the prime and target category was arbitrary and changed regularly, relying on this rule required cognitive control systems to direct retrieval to the cued category. Importantly, reliance on the rule was adaptive to the degree that successful retrieval aided a separate decision

regarding the lexical status of the target letter string. The adaptive nature of the rule was evident in facilitated behavioral measures for Expected relative to Unexpected trials. Moreover, this expectation effect emerged in behavior once sufficient time within a trial was provided (the Long SOA), consistent with the contribution of a slow, cognitive control process.

DDM estimates provided specificity regarding potential sources of these effects. In particular, the effects of Expected and Unexpected conditions on the non-decision and drift rate components of the model suggest that reliance on the rule had separable consequences on the lexical decision. Expected trials were partly faster because the non-decision component was speeded relative to Neutral and Unexpected items. This speeding could reflect faster encoding of the item, in essence a priming effect on the target because of prior rule-guided retrieval at presentation of the prime. However, as non-decision time can be difficult to distinguish from the effects of intertrial shifts in decision boundary (Ratcliff & Frank, 2012), another possibility is that participants dynamically shift their decision threshold depending on the fluency of retrieval, such as by adopting a higher boundary when fluency is lower, conflict arises, or expected conditions change (Cavanagh et al., 2011).

Unexpected trials were marked by a lower drift rate relative to Expected and Neutral trials. In other words, evidence accumulated toward the boundary more slowly when the rule was violated, impacting the decision process itself. One account of this is that recovery of item-specific information on Unexpected trials suffered interference from the prior retrieval of the cued, but now irrelevant, category and/or because of lingering suppression of the prime-related category that is unexpectedly relevant again (recall that, on Unexpected trials, the prime and target were semantically related). We note that there is no particular relationship of rule violations to word versus non-word decision frequency itself and that the unexpected item is a word. Hence, it is unlikely that the expectancy violation itself provided evidence that the item is a non-word (i.e., countervailing evidence). Rather, we interpret the slowed drift rate for Unexpected trials as reflecting (a) the reduced benefit of prior controlled retrieval present on Expected trials and/or (b) the impact of interference, as in blocking or proactive interference, arising because of the prior retrieval of irrelevant information from the expected category. Although it is important to note that “conflict” itself is not directly estimated as a parameter by the DDM. Thus, overall, the DDM indicates that participants generally relied on the rule to guide retrieval and indirectly improve their performance on the lexical decision.

In this context, we observed novel evidence that the BG and specifically ventral striatum may evaluate the success of a retrieval strategy (i.e., the retrieval rule). These signals, in turn, drive the cognitive control system to adjust reliance on this strategy. There were three primary

Thus, the ventral striatal relationship to the ease of decision-making appeared to be specific to trials on which expectations were established by the Rule. This is particularly notable given that the Rule incurred an overhead on working memory that resulted in globally slower RT and worse errors for the Rule than Neutral blocks. Hence, these observations are difficult to reconcile with general accounts of these results on the basis of simple reward (i.e., participants who are doing better at the task generally are positively reinforced) or of the learning effect on the basis of motivation (i.e., participants who are generally doing better continue to try at the task later). Rather, given the selectivity of these effects to Rule conditions, an account of these effects must make reference to the retrieval Rule specifically and its impact on decision-making.

A second key point is that these ventral striatal effects are not easily explained as learning at the motor or “response” level. As described later, prior work has provided suggestive evidence regarding a role for reinforcement learning and striatum in coding memory success and the achievement of retrieval goals (Schwarze et al., 2013; Han et al., 2010; Han & Dobbins, 2009). However, an open question from these investigations has been whether RPE is being coded at the level of the retrieval process or at the level of specific responses or categorical reports (Maddox & Bohil, 2005; Lauwereyns, Watanabe, Coe, & Hikosaka, 2002). In the present experiment, the retrieval rule has no direct bearing on the report of the participant; Word and Nonword responses were equally likely following any given category prime. Rather, the rule established expectations regarding the type of evidence that could contribute to the decision (i.e., the semantic features of the word). Thus, learning affected reliance on the retrieval rule rather than choosing a particular response.

Most of the observed effects of expectation and learning were related to positive activation on Expected trials rather than negative activation on Unexpected trials. To some degree, this is surprising to the degree that we predict negative and positive RPEs during this task. However, it is not necessarily the case that symmetric activation and deactivation for positive and negative RPEs must be observed in striatum. First, in the present experiment, we did not have a strong baseline against which to contrast positive and negative RPE, as it is not necessarily appropriate to assume that the fixation baseline is a true “zero RPE” baseline. The Neutral condition provides something of a baseline and notably does fall between the Expected and Unexpected signal change (Figure 3). However, there were other differences between this condition and Rule conditions (such as working memory demand) that can complicate its interpretation. Second, in the present experiment, we do not know to what degree individual participants expected certain outcomes, and so though it may be possible that they tend to expect positive outcomes and so have large negative

RPEs, they may tend to expect negative outcomes and so have larger positive RPEs. In the absence of a trial-to-trial fit, we must rely on the relative mean difference. Finally, although reinforcement learning models directly relating prediction error to BOLD signal in the striatum are often discussed as having symmetric effects, this is rarely tested in separate contrasts for positive and negative RPE in the human fMRI literature. When such an analysis has been done, there are often asymmetries, and indeed there is a precedent for somewhat better correlation with positive than negative RPE. For example, Badre, Doll, Long, and Frank (2012) tested positive and negative RPE separately using trial-to-trial estimates of prediction error during a more traditional reinforcement learning task. In this experiment, we only found evidence of striatum correlating with positive RPEs and not negative RPEs, despite negative RPEs contributing to a learning effect. Similar positive RPE-specific effects have been observed in striatum by others when positive and negative valences are tested separately (e.g., Pessiglione, Seymour, Flandin, Dolan, & Frith, 2006). Regardless of these uncertainties, the core prediction for RPE in the present design is a relative mean difference in activation between Expected and Unexpected in the direction observed here along with the observed relationship to behavior.

Following from this discussion, there are limitations of this study that necessitate further research in this domain to fully confirm our interpretation of these results. First, as noted above, the reported brain–behavior correlations are at the whole subject/session level. We interpret these effects as arising from trial-to-trial differences in drift rate, learning, and so forth, that then emerge in aggregate at the whole-subject level. However, we recognize that correlations observed across the group, although common in neuroimaging, may not represent the within-subject patterns of correlation (Simpson, 1951). This aggregate approach was necessary in this study as our design and the latent nature of memory retrieval itself made it difficult to fit individual participant learning rates, trial-to-trial changes in RPE, and/or belief about the utility of the retrieval rule. Nevertheless, future work in this domain should attempt to employ designs that can model individual subject learning more directly and so permit this type of trial-by-trial analysis. Second, although the learning correlation did include a temporal component, in that striatal activation early in the experiment correlated with shifts in reliance on the rule in the second half of the experiment, it is not possible to draw a conclusive causal link based on fMRI data or brain–behavior correlation. Nevertheless, the present results provide a clear motivation to further test learning of declarative retrieval strategies in populations with known disruption of striatal and BG systems, such as in Parkinson’s disease.

These caveats notwithstanding, the present results advance an emerging literature on the role of striatum in the cognitive control of declarative memory retrieval (reviewed in Scimeca & Badre, 2012). Considerable evidence

Of interest, whether DLPFC is maintaining the retrieval rule to influence striatal gating or is maintaining retrieved features, the working memory gating account ties the role of DLPFC to a process that could support what has been termed postretrieval selection (Badre & Wagner, 2007) and is discussed in more detail below.

The working memory gating account also potentially raises an alternative interpretation of the activation observed in striatum. From this perspective, striatal activation might reflect the response of go pathways that disinhibit thalamic input to PFC following encounter with the Expected item or rule-relevant semantic features. However, it is not clear why additional gating would not also be required for Unexpected trials, either to update the retrieval rule or to input/select new features. One possibility is that striatal subregions distinct from those identified here, such as in dorsal striatum, would gate DLPFC (Bornstein & Daw, 2011; O'Doherty et al., 2004). However, the present design did not permit us to separate between the prime and target phases of the trial in fMRI that could have better distinguished effects of gating/expectation from those related to outcome.

As with DLPFC, VLPFC was more activated for Unexpected than Expected trials. However, Expected trial activation in VLPFC did not differ from Neutral. Moreover, VLPFC activation was correlated with drift rate during the Expected Long condition, associating its activity with the evidence accumulation process itself. We should note that the present effects are related to the "signed drift rate" or the drift toward the Word decision specifically. This may be an important distinction, as prior work has associated DLPFC with the absolute drift or the overall ease of the decision (e.g., Heekeren, Marrett, Ruff, Bandettini, & Ungerleider, 2006). Thus, in this study, whereas DLPFC appeared crucial for maintaining the arbitrary retrieval rule or retrieved features, VLPFC may have engaged reactively to guide retrieval of evidence for a word response when the rule was violated.

To the degree that VLPFC operates reactively to support the decision process, its activation could reflect (a) the demand to activate (or reactivated) target information from memory (i.e., controlled retrieval) and/or (b) the demand to select against the now-irrelevant information in working memory (i.e., postretrieval selection). Both of these processes could theoretically increase drift rates by making diagnostic evidence more available and prior evidence has implicated VLPFC in both controlled retrieval and selection (Badre & Wagner, 2007). Moreover, some studies indicate that these functions are independently supported by aVLPFC and more caudal mid-VLPFC, respectively (Badre et al., 2005), including during lexical decision (Gold et al., 2006). In this regard, it is notable that the pattern described above was in the anterior, orbitalis portion of VLPFC that prior work has associated with controlled retrieval. Moreover, an ROI placed in mid-VLPFC revealed a pattern of response that was similar to that observed for DLPFC rather than aVLPFC. This is consistent

with previous functional dissociations between mid- and aVLPFC and also concurs with recent work suggesting that each subregion tends to correlate with separate functional networks (Barredo et al., in press). However, although it is possible that differences in the present data set could reflect functional distinctions between aVLPFC and mid-VLPFC/DLPFC networks observed previously (e.g., Barredo et al., in press; Badre & Wagner, 2007), this study was not designed to specifically test a controlled retrieval versus postretrieval account of this distinction over other possible functional differences.

Finally, the present results may inform recent debates concerning the functional role of PPC in memory retrieval. Broadly, regions of PPC have been repeatedly associated with episodic retrieval success (Wagner, Shannon, Kahn, & Buckner, 2005). However, PPC is functionally heterogeneous (Hutchinson et al., 2014; Nelson et al., 2010; Cabeza, Ciaramelli, Olson, & Moscovitch, 2008; Wagner et al., 2005) with dorsal regions, like IPS, being associated with top-down attention (Cabeza et al., 2008) or decision-making processes (Hutchinson et al., 2014) and ventral regions, like AG, with processes related to retrieval outcome (Hutchinson, Uncapher, & Wagner, 2009; Vilberg & Rugg, 2008) and bottom-up attention (O'Connor et al., 2010; Cabeza et al., 2008).

The current study involved semantic rather than episodic retrieval. Nevertheless, we observed activation in both IPS and AG subregions of PPC. And consistent with prior observations, these regions appeared to functionally dissociate, with the IPS showing a pattern of response similar to that of DLPFC and mid-VLPFC and the AG response being similar to aVLPFC.

Of particular note, AG demonstrated an effect of Expectation similar to that of aVLPFC whereby Unexpected trials were associated with more activation than either Expected or Neutral conditions. Thus, similar to aVLPFC, this implicates AG in the network supporting retrieval in response to an unexpected cue. AG has been widely associated with semantic processing and is particularly activated under conditions in which richer semantic content is available (Binder & Desai, 2011). Likewise, accounts of AG function during episodic retrieval relate it to bottom-up attention to memory (Cabeza et al., 2008) or retrieval output (Vilberg & Rugg, 2008; Wagner et al., 2005). In the present work, the demands on retrieval may similarly result in greater semantic analysis and activation of semantic representations during Unexpected relative to Expected and Neutral conditions. However, prior episodic retrieval studies have also indicated that AG may be related to the match of expectations rather than retrieval per se. In particular, O'Connor et al. (2010) found an effect of expectancy violation in AG when expected retrieval experiences (Old or New) were violated at a memory probe. The present results replicate this effect, although, as noted above, it is possible that the violation of expectation gated retrieval in the present design. Thus, the present results do not distinguish between retrieval-based and attentional accounts of AG function.

To conclude, we have provided evidence that ventral striatal activation tracks the positive outcomes that follow from an effective retrieval strategy, and this signal is related to an individual's tendency to subsequently abandon or rely on this strategy. These results suggest a closer relationship than previously demonstrated between nigrostriatal dopamine systems related to basic reinforcement learning functions and the cognitive control of memory retrieval.

Acknowledgments

This work was supported by the National Institute of Neurological Disease and Stroke (NS065046), the Alfred P. Sloan Foundation, and the James S. McDonnell Foundation.

Reprint requests should be sent to David Badre, Box 1821, Brown University, Providence, RI 02912-1978, or via e-mail: David_Badre@brown.edu.

REFERENCES

- Anderson, J. R., & Milson, R. (1989). Human memory: An adaptive perspective. *Psychological Review*, *96*, 703–719.
- Anderson, M. C., Ochsner, K. N., Kuhl, B., Cooper, J., Robertson, E., Gabrieli, S. W., et al. (2004). Neural systems underlying the suppression of unwanted memories. *Science*, *303*, 232–235.
- Atkinson, R. C., & Shiffrin, R. M. (1971). The control of short-term memory. *Scientific American*, *225*, 82–90.
- Badre, D. (2012). Opening the gate to working memory. *Proceedings of the National Academy of Sciences, U.S.A.*, *109*, 19878–19879.
- Badre, D., Doll, B. B., Long, N. M., & Frank, M. J. (2012). Rostrolateral prefrontal cortex and individual differences in uncertainty-driven exploration. *Neuron*, *73*, 595–607.
- Badre, D., & Frank, M. J. (2012). Mechanisms of hierarchical reinforcement learning in cortico-striatal circuits 2: Evidence from fMRI. *Cerebral Cortex*, *22*, 527–536.
- Badre, D., Poldrack, R. A., Pare-Blagoev, E. J., Insler, R. Z., & Wagner, A. D. (2005). Dissociable controlled retrieval and generalized selection mechanisms in ventrolateral prefrontal cortex. *Neuron*, *47*, 907–918.
- Badre, D., & Wagner, A. D. (2007). Left ventrolateral prefrontal cortex and the cognitive control of memory. *Neuropsychologia*, *45*, 2883–2901.
- Balota, D. A., & Spieler, D. H. (1999). Word frequency, repetition, and lexicality effects in word recognition tasks: Beyond measures of central tendency. *Journal of Experimental Psychology: General*, *128*, 32–55.
- Barredo, J., Oztekin, I., & Badre, D. (in press). Ventral fronto-temporal pathway supporting cognitive control of episodic memory retrieval. *Cerebral Cortex*.
- Becker, S., & Lim, J. (2003). A computational model of prefrontal control in free recall: Strategic memory use in the California Verbal Learning Task. *Journal of Cognitive Neuroscience*, *15*, 821–832.
- Benjamin, A. S. (2007). Memory is more than just remembering: Strategic control of encoding, accessing memory, and making decisions. In A. S. Benjamin & B. H. Ross (Eds.), *The psychology of learning and motivation: Skill and strategy in memory use* (pp. 175–223). London: Academic Press.
- Binder, J. R., & Desai, R. H. (2011). The neurobiology of semantic memory. *Trends in Cognitive Sciences*, *15*, 527–536.
- Bornstein, A. M., & Daw, N. D. (2011). Multiplicity of control in the basal ganglia: Computational roles of striatal subregions. *Current Opinion in Neurobiology*, *21*, 374–380.
- Brainard, D. H. (1997). The Psychophysics Toolbox. *Spatial Vision*, *10*, 433–436.
- Braver, T. S., & Cohen, J. D. (2000). On the control of control: The role of dopamine in regulating prefrontal function and working memory. In S. Monsell & J. Driver (Eds.), *Attention and performance XVIII* (pp. 713–737). Cambridge, MA: MIT Press.
- Bray, S., & O'Doherty, J. (2007). Neural coding of reward-prediction error signals during classical conditioning with attractive faces. *Journal of Neurophysiology*, *97*, 3036–3045.
- Cabeza, R., Ciaramelli, E., Olson, I. R., & Moscovitch, M. (2008). The parietal cortex and episodic memory: An attentional account. *Nature Reviews Neuroscience*, *9*, 613–625.
- Cavanagh, J. F., Wiecki, T. V., Cohen, M. X., Figueroa, C. M., Samanta, J., Sherman, S. J., et al. (2011). Subthalamic nucleus stimulation reverses mediofrontal influence over decision threshold. *Nature Neuroscience*, *14*, 1462–1467.
- Chatham, C. H., & Badre, D. (2013). Working memory management and predicted utility. *Frontiers in Behavioral Neuroscience*, *7*, 83.
- Chatham, C. H., Frank, M. J., & Badre, D. (2014). Cortico-striatal systems supporting output gating of working memory. *Neuron*, *81*, 930–942.
- Cohen, N. J., Eichenbaum, H., & Poldrack, R. A. (1997). Memory for items and memory for relations in the procedural/declarative memory framework. *Memory*, *5*, 131–178.
- Cools, R. (2011). Dopaminergic control of the striatum for high-level cognition. *Current Opinion in Neurobiology*, *21*, 402–407.
- Cools, R., Barker, R. A., Sahakian, B. J., & Robbins, T. W. (2001). Mechanisms of cognitive set flexibility in Parkinson's disease. *Brain*, *124*, 2503–2512.
- Crescentini, C., Marin, D., Del Missier, F., Biasutti, E., & Shallice, T. (2011). Interference from retrieval cues in Parkinson's disease. *Neuropsychology*, *25*, 720–733.
- Crescentini, C., Mondolo, F., Biasutti, E., & Shallice, T. (2008). Supervisory and routine processes in noun and verb generation in nondemented patients with Parkinson's disease. *Neuropsychologia*, *46*, 434–447.
- D'Ardenne, K., Eshel, N., Luka, J., Lenartowicz, A., Nystrom, L. E., & Cohen, J. D. (2012). Role of prefrontal cortex and the midbrain dopamine system in working memory updating. *Proceedings of the National Academy of Sciences, U.S.A.*, *109*, 19900–19909.
- Daw, N. D., Gershman, S. J., Seymour, B., Dayan, P., & Dolan, R. J. (2011). Model-based influences on humans' choices and striatal prediction errors. *Neuron*, *69*, 1204–1215.
- Dobbins, I. G., Foley, H., Schacter, D. L., & Wagner, A. D. (2002). Executive control during episodic retrieval: Multiple prefrontal processes subserve source memory. *Neuron*, *35*, 989–996.
- Donkin, C., Heathcote, A., Brown, S., & Andrews, S. (2009). Non-decision time effects in the lexical decision task. Paper presented at Proceedings of the 31st Annual Conference of the Cognitive Science Society. Austin, TX.
- Favreau, M., & Segalowitz, N. S. (1983). Automatic and controlled processes in the first- and second-language reading of fluent bilinguals. *Memory and Cognition*, *11*, 565–574.

- Gabrieli, J. D., Poldrack, R. A., & Desmond, J. E. (1998). The role of left prefrontal cortex in language and memory. *Proceedings of the National Academy of Sciences, U.S.A.*, *95*, 906–913.
- Gelman, A., Carlin, J. B., Stern, H. S., & Rubin, D. B. (2004). *Bayesian data analysis* (2nd ed.). Boca Raton, FL: Chapman and Hall/CRC.
- Gershberg, F. B., & Shimamura, A. P. (1995). Impaired use of organizational strategies in free recall following frontal lobe damage. *Neuropsychologia*, *33*, 1305–1333.
- Glascher, J., Daw, N., Dayan, P., & O'Doherty, J. P. (2010). States versus rewards: Dissociable neural prediction error signals underlying model-based and model-free reinforcement learning. *Neuron*, *66*, 585–595.
- Gold, B. T., Balota, D. A., Jones, S. J., Powell, D. K., Smith, C. D., & Andersen, A. H. (2006). Dissociation of automatic and strategic lexical-semantic: Functional magnetic resonance imaging evidence for differing roles of multiple frontotemporal regions. *Journal of Neuroscience*, *26*, 6523–6532.
- Han, S., & Dobbins, I. G. (2009). Regulating recognition decisions through incremental reinforcement learning. *Psychonomic Bulletin & Review*, *16*, 469–474.
- Han, S., Huettel, S. A., Raposo, A., Adcock, R. A., & Dobbins, I. G. (2010). Functional significance of striatal responses during episodic decisions: Recovery or goal attainment? *Journal of Neuroscience*, *30*, 4767–4775.
- Heekeren, H. R., Marrett, S., Ruff, D. A., Bandettini, P. A., & Ungerleider, L. G. (2006). Involvement of human left dorsolateral prefrontal cortex in perceptual decision-making is independent of response modality. *Proceedings of the National Academy of Sciences, U.S.A.*, *103*, 10023–10028.
- Hutchinson, J. B., Uncapher, M. R., & Wagner, A. D. (2009). Posterior parietal cortex and episodic retrieval: Convergent and divergent effects of attention and memory. *Learning and Memory*, *16*, 343–356.
- Hutchinson, J. B., Uncapher, M. R., Weiner, K. S., Bressler, D. W., Silver, M. A., Preston, A. R., et al. (2014). Functional heterogeneity in posterior parietal cortex across attention and episodic memory retrieval. *Cerebral Cortex*, *24*, 49–66.
- Kucera, H., & Francis, W. N. (1967). *Computational analysis of present-day English*. Providence, RI: Brown University Press.
- Lauwereyns, J., Watanabe, K., Coe, B., & Hikosaka, O. (2002). A neural correlate of response bias in monkey caudate nucleus. *Nature*, *418*, 413–417.
- Li, J., & Daw, N. D. (2011). Signals in human striatum are appropriate for policy update rather than value prediction. *Journal of Neuroscience*, *31*, 5504–5511.
- Maddox, W. T., & Bohil, C. J. (2005). Optimal classifier feedback improves cost-benefit but not base-rate decision criterion learning in perceptual categorization. *Memory & Cognition*, *33*, 303–319.
- McNab, F., & Klingberg, T. (2008). Prefrontal cortex and basal ganglia control access to working memory. *Nature Neuroscience*, *11*, 103–107.
- Meyer, D. E., & Schvaneveldt, R. W. (1971). Facilitation in recognizing pairs of words: Evidence of a dependence between retrieval operations. *Journal of Experimental Psychology*, *90*, 227–234.
- Miller, E. K., & Cohen, J. D. (2001). An integrative theory of prefrontal cortex function. *Annual Review of Neuroscience*, *24*, 167–202.
- Moscovitch, M. (1992). Memory and working-with-memory—A component process model based on modules and central systems. *Journal of Cognitive Neuroscience*, *4*, 257–267.
- Neely, J. H. (1977). Semantic priming and retrieval from lexical memory: Roles of inhibitionless spreading activation and limited capacity attention. *Journal of Experimental Psychology: General*, *106*, 226–254.
- Neely, J. H. (1991). Semantic priming effects in visual word recognition: A selective review of current findings and theories. In D. Besner & G. W. Humphreys (Eds.), *Basic processing in reading: Visual word recognition* (pp. 264–336). Hillsdale, NJ: Lawrence Erlbaum Associates.
- Nelson, S. M., Cohen, A. L., Power, J. D., Wig, G. S., Miezin, F. M., Wheeler, M. E., et al. (2010). A parcellation scheme for human left lateral parietal cortex. *Neuron*, *67*, 156–170.
- O'Connor, A. R., Han, S., & Dobbins, I. G. (2010). The inferior parietal lobule and recognition memory: Expectancy violation or successful retrieval? *Journal of Neuroscience*, *30*, 2924–2934.
- O'Doherty, J., Dayan, P., Schultz, J., Deichmann, R., Friston, K., & Dolan, R. J. (2004). Dissociable roles of ventral and dorsal striatum in instrumental conditioning. *Science*, *304*, 452–454.
- O'Reilly, R. C., & Frank, M. J. (2006). Making working memory work: A computational model of learning in the prefrontal cortex and basal ganglia. *Neural Computation*, *18*, 283–328.
- Pessiglione, M., Seymour, B., Flandin, G., Dolan, R. J., & Frith, C. D. (2006). Dopamine-dependent prediction errors underpin reward-seeking behaviour in humans. *Nature*, *442*, 1042–1045.
- Rastle, K., Harrington, J., & Coltheart, M. (2002). 358,534 nonwords: The ARC Nonword Database. *Quarterly Journal of Experimental Psychology A*, *55*, 1339–1362.
- Ratcliff, R. (1978). A theory of memory retrieval. *Psychological Review*, *85*, 59–108.
- Ratcliff, R., & Frank, M. J. (2012). Reinforcement-based decision making in corticostriatal circuits: Mutual constraints by neurocomputational and diffusion models. *Neural Computation*, *24*, 1186–1229.
- Ratcliff, R., Gomez, P., & McKoon, G. (2004). A diffusion model account of the lexical decision task. *Psychological Review*, *111*, 159–182.
- Ratcliff, R., & McKoon, G. (2008). The diffusion decision model: Theory and data for two-choice decision tasks. *Neural Computation*, *20*, 873–922.
- Ratcliff, R., Thapar, A., Gomez, P., & McKoon, G. (2004). A diffusion model analysis of the effects of aging in the lexical decision task. *Psychology and Aging*, *19*, 278–289.
- Rugg, M. D., & Wilding, E. L. (2000). Retrieval processing and episodic memory. *Trends in Cognitive Sciences*, *4*, 108–115.
- Schultz, W., Dayan, P., & Montague, P. R. (1997). A neural substrate of prediction and reward. *Science*, *275*, 1593–1599.
- Schwarze, U., Bingel, U., Badre, D., & Sommer, T. (2013). Ventral striatal activity correlates with memory confidence for old- and new-responses in a difficult recognition test. *PLoS One*, *8*, e54324.
- Scimeca, J. M., & Badre, D. (2012). Striatal contributions to declarative memory retrieval. *Neuron*, *75*, 380–392.
- Shenhav, A., Botvinick, M. M., & Cohen, J. D. (2013). The expected value of control: An integrative theory of anterior cingulate cortex function. *Neuron*, *79*, 217–240.
- Simpson, E. H. (1951). The interpretation of interaction in contingency tables. *Journal of the Royal Statistical Society, Series B*, *13*, 238–241.
- Spaniol, J., Davidson, P. S., Kim, A. S., Han, H., Moscovitch, M., & Grady, C. L. (2009). Event-related fMRI studies of episodic encoding and retrieval: Meta-analyses using activation likelihood estimation. *Neuropsychologia*, *47*, 1765–1779.
- Squire, L. R. (1992). Memory and the hippocampus: A synthesis from findings with rats, monkeys, and humans. *Psychological Review*, *99*, 195–231.

- Stuss, D. T., Alexander, M. P., Palumbo, C. L., Buckle, L., Sayer, L., & Pogue, J. (1994). Organizational strategies of patients with unilateral or bilateral frontal lobe injury in word list learning tasks. *Neuropsychologia*, *8*, 355–373.
- Sutton, R. S., & Barto, A. G. (1998). *Reinforcement learning: An introduction*. Cambridge, MA: MIT Press.
- Vilberg, K. L., & Rugg, M. D. (2008). Memory retrieval and the parietal cortex: A review of evidence from a dual-process perspective. *Neuropsychologia*, *46*, 1787–1799.
- Wagner, A. D., Shannon, B. J., Kahn, I., & Buckner, R. L. (2005). Parietal lobe contributions to episodic memory retrieval. *Trends in Cognitive Sciences*, *9*, 445–453.
- Wiecki, T. V., Sofer, I., & Frank, M. J. (2013). HDDM: hierarchical Bayesian estimation of the drift-diffusion model in Python. *Frontiers in Neuroinformatics*, *7*, 14.
- Yarkoni, T., Poldrack, R. A., Nichols, T. E., Van Essen, D. C., & Wager, T. D. (2011). Large-scale automated synthesis of human functional neuroimaging data. *Nature Methods*, *8*, 665–670.