Learning to Remember: Striatal Contributions to Declarative Memory Retrieval

by

Jason Michael Scimeca

B.A., University of Chicago, 2008

A dissertation submitted in partial fulfillment of the

requirements for the degree of Doctor of Philosophy

in Cognitive, Linguistic, and Psychological Sciences at Brown University

Providence, Rhode Island

May 2016

This dissertation by Jason M. Scimeca is accepted in its present form
by the Department of Cognitive, Linguistic, and Psychological Sciences as
as satisfying the dissertation requirement for the degree of Doctor of Philosophy.

Date _____ _____

David Badre, Advisor

Recommended to the Graduate Council

Date _____ _____

Michael Frank, Reader

Date _____ _____

William Heindel, Reader

Approved by the Graduate Council

Date _____ _____

Peter Weber, Dean of the Graduate School

Jason M. Scimeca

Birthdate: February 18, 1986
Brown University | Department of Cognitive, Linguistic, and Psychological Sciences (CLPS)
Box 1821, 190 Thayer Street, Brown University, Providence, RI, 02912
jason_scimeca@brown.edu | jscimeca@gmail.com
http://scimeca.net

EDUCATION

Brown University (2010-Present)
Progress towards Ph.D. in Cognitive Science; Advisor: David Badre

University of Chicago (2004-2008)
B.A. in Biological Sciences with Specialization in Neuroscience; Psychology (with Honors)

AWARDS AND HONORS

2013	Brown University Presidential Award for Excellence in Teaching
2012	NIMH Summer Institute in Cognitive Neuroscience Fellowship
2010	National Science Foundation Graduate Research Fellowship Honorable Mention
2007	University of Chicago College Research Opportunity Grant
2004	National Merit Scholarship

PUBLICATIONS

Journal Articles

9. **Scimeca, J. M.** & Franconeri, S. L. (2015). Selecting and tracking multiple objects. *Wiley Interdisciplinary Reviews: Cognitive Science, 6*, 109-118.

8. Badre, D., Lebrecht, S., Pagliaccio, D., Long, N. M., & **Scimeca, J. M.** (2014). Ventral striatum and the evaluation of memory retrieval strategies. *Journal of Cognitive Neuroscience, 26*, 1928-1948.

7. Cheon, B. K., Im., D., Harada, T., Kim, J., Mathur, V. A., **Scimeca, J. M.**, Parrish, T. B., Park, H., &
Chiao, J. Y. (2013). Cultural modulation of the neural correlates of emotional pain perception: The role of other-focusedness. *Neuropsychologia, 51*, 1177-1186.

6. **Scimeca, J. M.** & Badre, D. (2012). Striatal contributions to declarative memory retrieval. *Neuron, 75*, 380-392.

5. Franconeri, S. L., **Scimeca, J. M.**, Roth, J. C., Helseth, S. A., & Kahn, L. E. (2012). Flexible visual processing of spatial relationships. *Cognition, 122,* 210-227.

4. **Scimeca, J. M.**, McDonough, I. M., & Gallo, D. A. (2011). Quality trumps quantity at reducing memory errors: Implications for retrieval monitoring and mirror effects. *Journal of Memory and Language, 65*, 363-377.

3. Cheon, B. K., Im, D., Harada, T., Kim, J., Mathur, V. A., **Scimeca, J. M.**, Parrish, T. B., Park, H., & Chiao, J. Y. (2011). Cultural influences on neural basis of intergroup empathy. *NeuroImage*, *57*, 642-650.

2. Franconeri, S. L., Jonathan, S. V., & **Scimeca, J. M.** (2010). Tracking multiple objects is limited only by object spacing, not by speed, time, or capacity. *Psychological Science*, *21*, 920-925.

1. Gallo, D. A., McDonough, I. M., & **Scimeca, J.** (2010). Dissociating source memory decisions in prefrontal cortex: fMRI of diagnostic and disqualifying monitoring. *Journal of Cognitive Neuroscience, 22*, 955-969.

- **Google Scholar Citation Profile:** http://scholar.google.com/citations?user=-Dw1o2AAAAAJ

Forthcoming Manuscripts

- **Scimeca, J. M.**, Katzman, P. L., & Badre, D. (under review). Striatal prediction errors support dynamic control of declarative memory decisions.

- **Scimeca, J. M.**, Katzman, P. L., & Badre, D. (under review). Learning to remember: Prediction errors influence recognition memory decision criterion.

- Chiao, J. Y., Cheon, B. K., Harada, T., **Scimeca, J. M.**, & Mathur, V. A. (under review). Cultural influences on neural basis of environmental concern.

- Mathur, V. A., Cheon, B. K., Harada, T., **Scimeca, J. M.**, & Chiao, J. Y. (under review). Overlapping neural response to the pain or harm of people, animals, and nature.

- **Scimeca, J. M.**, Jonathan, S. V., & Franconeri, S. L. (in revision). Explaining performance in multiple object tracking: Competitive spatiotopic maps versus general resource theories.

CONFERENCE PROCEEDINGS

- **Scimeca, J. M.**, Katzman, P. L., & Badre, D. (November, 2013). *Evaluating and updating control processes in recognition memory*. Poster presented at the 43rd Annual Meeting of the Society for Neuroscience, San Diego, CA.
- **Scimeca, J. M.**, Katzman, P. L., & Badre, D. (April, 2013). *The role of prediction errors in control of recognition memory decisions*. Poster presented at the 20th Annual Meeting of the Cognitive Neuroscience Society, San Francisco, CA.
- Franconeri, S., **Scimeca, J.**, & Jonathan, S. (May, 2012). *Maintaining selection of multiple moving objects*. Poster presented at 2012 Meeting of Vision Sciences Society, Naples, FL.
- **Scimeca, J. M.**, McShane, L. M., Brew, J. A., & Badre, D. (November, 2011). *Acquisition and adaptation of rule-guided retrieval strategies in memory*. Poster presented at the 41st Annual Meeting of the Society for Neuroscience, Washington, DC.
- Franconeri, S., **Scimeca, J.**, Roth, J., & Helseth, S. (May, 2011). *Flexible visual processing of spatial relationships*. Talk presented at 2011 Meeting of Vision Sciences Society, Naples, FL.
- Levinthal, B. R., Jonathan, S., **Scimeca, J.**, & Franconeri, S. (May, 2011). *Competition limits spatial selection*. Poster presented at 2011 Meeting of Vision Sciences Society, Naples, FL.
- Franconeri, S., **Scimeca, J.**, Roth, J., Kahn, L., & Helseth, S. A. (April, 2011). *Flexible visual processing of spatial relationships*. Talk presented at the 18th Annual Meeting of the Cognitive Neuroscience Society, San Francisco, CA.
- Cheon, B., Im, D., Harada, T., Park, J., Mathur, V., **Scimeca, J.**, Park, H., Chiao, J. (April, 2011). *The role of culture on the automaticity of empathy: The modulating influence of interdependence.* Poster presented at the 18th Annual Meeting of the Cognitive Neuroscience Society, San Francisco, CA.
- Franconeri, S. L., Jonathan, S.V., & **Scimeca, J. M.** (November, 2009). *Tracking multiple objects is limited only by interobject crowding, and not object speed.* Talk presented at the 50th Annual Meeting of the Psychonomic Society, Boston, MA.

- Rotella, K., Richeson, J., **Scimeca, J.**, & Chiao, J. (October, 2009). *Neural basis of economic trust with religious ingroup and outgroup members*. Poster presented at the 39th Annual Meeting of the Society for Neuroscience, Chicago, IL.
- Cheon, B.K., Im, D., Harada, T., Kim, J., Mathur, V. A., **Scimeca, J.**, Park, H., & Chiao, J. Y. (October, 2009). *Cultural variation in neural response within temporo-parietal junction during intergroup empathy*. Talk presented at the 39th Annual Meeting of the Society for Neuroscience, Chicago, IL.
- Mathur, V.A., Harada, T., Cheon, B.K., **Scimeca, J.**, & Chiao, J.Y. (March, 2009). *Empathic neural response to living things as a function of agency and experience*. Poster presented at the 16th Annual Meeting of the Cognitive Neuroscience Society, San Francisco, CA.
- Cheon, B.K., Im, D., Harada, T., Mathur, V.A., **Scimeca, J.**, Park, H., & Chiao, J.Y. (March, 2009). *Universal and culturally-specific neural basis of ingroup bias in empathy*. Poster presented at the 16th Annual Meeting of Cognitive Neuroscience Society, San Francisco, CA.
- Gallo, D. A., McDonough, I. M., & **Scimeca, J. M.**† (April, 2008). *Prefrontal Regions Differentially Contribute to Source Monitoring Processes*. Poster presented at the 15th Annual Meeting of the Cognitive Neuroscience Society, San Francisco, CA. († Presenting author)

MISCELLANEOUS RESEARCH AND PRESENTATIONS

- **Scimeca, J.M.** (June, 2014). *The role of reinforcement learning systems in recognition memory*. Talk presented to the Brown Leadership Alliance Summer Research Program, Providence, RI.
- **Scimeca, J.M.** (April, 2014). *Dynamic control of recognition memory decisions: The important role of prediction errors*. Talk presented at Brown CLPS Cognition Seminar, Providence, RI.
- **Scimeca, J.M.** (October, 2012). *The role of feedback and prediction error signals in controlling recognition memory*. Talk presented at Brown CLPS Cognition Seminar, Providence, RI.
- **Scimeca, J. M.,** Franconeri, S. L., & Jonathan, S. V. (October, 2010). *Tracking multiple objects is limited only by object spacing, not by speed, time, or capacity*. Talk presented at the inaugural meeting of the Brown CLPS Cognition Seminar, Providence, RI.
- **Scimeca, J. M.** (2008). Quality and Quantity: Differential Contributions to Recognition Memory. University of Chicago Undergraduate Honors Thesis.

OTHER EDUCATION, TRAINING, AND CERTIFICATION

2012	**NIMH Summer Institute in Cognitive Neuroscience (Santa Barbara, CA)** Episodic Memory and Adaptive Behavior / Brain Plasticity
2011-2012	**Brown University – Sheridan Center for Teaching and Learning** Sheridan Teaching Seminar – Certificate I Program
2009	**Northwestern University** School of Continuing Studies

TEACHING

2014 Fall	**Functional Magnetic Resonance Imaging: Theory and Practice (CLPS1490)** Graduate Teaching Assistant
2013 Fall	**Functional Magnetic Resonance Imaging: Theory and Practice (CLPS1490)** Graduate Teaching Assistant
2012-2013	**Brown University Presidential Award for Excellence in Teaching** Annual prize awarded to recognize outstanding pedagogical achievement; awarded to two recipients out of 400 Brown graduate students with teaching appointments
2013 Spring	**Elementary Psychology: An Introduction to Mind and Behavior (CLPS0010)** Graduate Teaching Assistant

2012 Fall **Functional Magnetic Resonance Imaging: Theory and Practice (CLPS1490)**
Graduate Teaching Assistant

MENTORING

Brown University: Undergraduate Directed Research & Senior Honors Thesis
- Perri Katzman

Brown University: Undergraduate Directed Research
- Caryn Cobb; Celia Ford; Amanda Ruggieri; Becca Zak

Brown University: Leadership Alliance Summer Research Early Identification Program
- Ari Meilech

PROFESSIONAL SERVICE

Ad Hoc Referee: Behavioural Brain Research; Cerebral Cortex; Human Brain Mapping

PROFESSIONAL MEMBERSHIPS

- Association for Psychological Science
- Cognitive Neuroscience Society
- Society for Neuroscience

UNIVERSITY SERVICE

University of Chicago Alumni Schools Committee, Member, June 2009-Present
- Conduct interviews and information sessions with prospective undergraduate students

Graduate Student Representative, Brown CLPS Department, May 2011-September 2013
- Graduate student liaison to CLPS Department faculty and Brown Graduate Student Council

Cognition Seminar Series Coordinator, Brown CLPS Department, June 2011-June 2012
- Coordinated external and internal speakers for the CLPS Cognition Seminar Series

Brown Alzheimer's Activists, Member, September 2010-December 2012
- Provide support to local healthcare centers through volunteering and fundraising; raise awareness of Alzheimer's Disease at Brown campus and throughout RI through lectures and educational events

APO National Co-ed Service Fraternity, Member, December 2004-January 2008
- Performed community service across UChicago, Hyde Park neighborhood, and Chicago

RESEARCH POSITIONS AND EXPERIENCE

Graduate Student, **Brown University** (August 2010-Present)
- Advisor: David Badre (Cognitive Neuroscience of Cognitive Control and Memory Laboratory)

Lab Manager, **Northwestern University** (September 2008-August 2010)
- Advisor: Joan Y. Chiao (Social Affective and Cultural Neuroscience Laboratory)
- Advisor: Steven L. Franconeri (Visual Attention and Cognition Laboratory)

Research Assistant, **University of Chicago** (April 2007-August 2008)
- Advisor: David A. Gallo (Memory Research Laboratory)

Student Technician, **University of Chicago Medical Center** (October 2006-August 2008)
- UChicago Medical Center – Center for Parkinson's Disease and Movement Disorders

Research Assistant, **Wayne State University School of Medicine** (June 2006-Sept. 2006)
- Advisor: Manuel Tancer (Brain Research And Imaging Neuroscience Division, Detroit, MI)

Updated July 2015

ACKNOWLEDGEMENTS

This dissertation is dedicated to my many teachers, especially my parents, who instilled in me a curiosity about the world around and inside us.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

CHAPTER 1

Striatal Contributions to Declarative Memory Retrieval[1]

Declarative memory has long been known to depend on the medial temporal lobe memory system. Recently, there has been renewed focus on the relationship between the basal ganglia and declarative memory, including the involvement of striatum. However, the specific contributions of striatum to declarative memory retrieval remain unknown. In this chapter, I review neuroimaging and neuropsychological evidence for the involvement of the striatum in declarative memory retrieval. From this review I propose that the striatum primarily supports cognitive control of memory retrieval, in concert with the prefrontal cortex.

I then propose three non-exclusive hypotheses for the specific role of striatum and the associated value system in declarative retrieval: (1) Striatum modulates the re-encoding of retrieved items in accord with their expected utility (adaptive encoding); (2) striatum selectively admits information into working memory that is expected to increase the likelihood of successful retrieval (adaptive gating); and (3) striatum enacts adjustments in retrieval control policies based on the outcome of retrieval (reinforcement learning). These hypotheses motivate the experimental work which comprises the remainder of this dissertation. I conclude the chapter with a brief outline of these experiments.

[1] This chapter has been adapted from a published manuscript:
Scimeca, J. M., & Badre, D. (2012). Striatal contributions to declarative memory retrieval. *Neuron*, 75, 380-392.

Introduction

Declarative memory retrieval refers to the conscious recovery of previously stored experiences, facts, and concepts that are verifiable through verbal report (Tulving, 1972). It has long been known that the medial temporal lobe (MTL) system is necessary for the formation, consolidation, and retrieval of declarative memory (N. J. Cohen, Poldrack, & Eichenbaum, 1997; Squire, 1992). By contrast, other types of long-term memory, such as skill learning or classical conditioning do not appear to require the MTL memory system (N. J. Cohen et al., 1997; Corkin, 1968; Knowlton, Squire, & Gluck, 1994). Rather, these forms of "non-declarative" memory are strongly associated with the reward driven mechanisms of the basal ganglia (N. J. Cohen et al., 1997; Knowlton, Mangels, & Squire, 1996; Packard, Hirsch, & White, 1989; Shohamy et al., 2004). However, mounting evidence from both neuroimaging and neuropsychological studies of declarative memory have renewed focus on interactions between the declarative and non-declarative systems, and have highlighted the potential role of the basal ganglia, including striatum (**Figure 1.1A**), in declarative memory performance, both at encoding and retrieval (Cohn, Moscovitch, & Davidson, 2010; Han, Huettel, Raposo, Adcock, & Dobbins, 2010; Long, Oztekin, & Badre, 2010; Moustafa & Gluck, 2011; Poldrack & Foerde, 2008; Shohamy & Adcock, 2010).

Outside of the long-term memory domain, there has been growing recognition of a broader role for striatal-frontal interactions beyond basic motor control. Specifically, recent years have seen a growth in our understanding of the mechanisms by which striatum supports higher cognitive functions like working memory, decision making, categorization, and cognitive control (Badre, Doll, Long, & Frank, 2012; Badre & Frank,

2012; Cools, 2011; Doll & Frank, 2009; Graybiel & Mink, 2009; Landau, Lal, O'Neil, Baker, & Jagust, 2009; Lewis, Dove, Robbins, Barker, & Owen, 2004; Seger & Miller, 2010; Stelzel, Basten, Montag, Reuter, & Fiebach, 2010).

However, to date, we still have a limited understanding of the role of these striatal mechanisms in declarative memory retrieval. Here, we review evidence for the involvement of the striatum in declarative memory retrieval. First, based on evidence from neuroimaging and neuropsychological studies of declarative memory, we argue that, along with the prefrontal cortex (PFC), the striatum supports the cognitive control of memory retrieval. Then, leveraging models of reinforcement learning and cognitive control theory outside of the memory domain, we propose a set of novel hypotheses regarding the potential mechanistic role of the striatum in declarative memory as a basis for future research.

Striatal responses to item recognition

An adaptive function of the declarative memory system is the ability to discriminate items and contexts with which an animal has prior experience versus those that are novel. The ability to recognize previously encountered items is known to require MTL structures, including perirhinal, parahippocampal, and hippocampal cortex (Eichenbaum, Yonelinas, & Ranganath, 2007; Schacter & Wagner, 1999; Squire & Wixted, 2011; Squire, 1992).

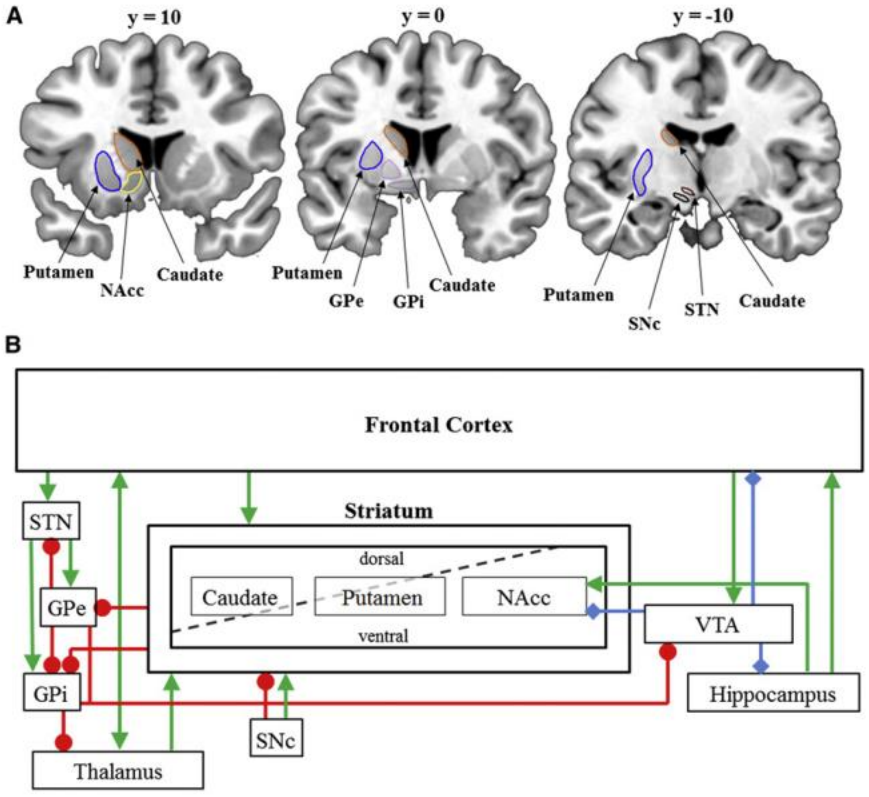

Figure 1.1 Basic Anatomy of the Basal Ganglia and Schematic Frontostriatal Circuitry. (A) Locations of basal ganglia structures are outlined on coronal slices (caudate, putamen, and nucleus accumbens [NAcc]); GPe, external segment of the globus pallidus; GPi, internal segment of the globus pallidus; SNc, substantia nigra pars compacta; and STN, subthalamic nucleus). (B) Schematic of frontal-striatal circuitry and the VTA-hippocampal loop. Projections from frontal cortex through the striatum (caudate, putamen, and nucleus accumbens), subthalamic nucleus (STN), globus pallidus, and thalamus drive adaptive gating of working memory representations. The VTA-hippocampal loop includes direct projections from ventral tegmental area (VTA) to hippocampus, and a circuit through the nucleus accumbens (NAcc) and globus pallidus from the hippocampus to VTA. The striatum is frequently characterized by a dorsal/ventral division (e.g. Bornstein and Daw, 2011), with different subregions of caudate and putamen existing on both sides of the divide. Green arrows indicate excitatory connections; red circles indicate inhibitory connections; blue diamonds indicate modulatory connections. GPe: external segment of the globus pallidus; GPi: internal segment of the globus pallidus; SNc: substantia nigra pars compacta.

In the item recognition paradigm, participants first encode a series of items, usually words or pictures, and are then shown a mix of items that they had seen

previously during encoding along with new items that have not been seen before. For each item, the participant judges whether the item has been seen previously (old) or not (new). Thus, contrasting trials on which participants correctly judged an old item as "old" (hits) against trials on which a participant correctly judged a new item as "new" (correct rejections [CR]) probes the neural correlates of "retrieval success".

Since the earliest event-related fMRI studies of the item-recognition task (e.g. Buckner, Koutstaal, Schacter, Wagner, & Rosen, 1998; Donaldson, Petersen, Buckner, & Louis, 2001; Rombouts, Barkhof, Witter, Machielsen, & Scheltens, 2001), retrieval success has yielded striatal activation. Moreover, this effect has been replicated across multiple variants of encoding tasks and stimulus materials. Spaniol and colleagues (2009) analyzed 81 fMRI studies of episodic memory, a subset of which included contrasts of encoding success (subsequent hits greater than misses) and/or retrieval success (hits greater than CR). A quantitative meta-analytic procedure indicated that retrieval success consistently activated striatum across studies, including both dorsal striatum in the left caudate and ventral striatum in regions of caudate, putamen, and nucleus accumbens (see also Kim, 2011). **Figure 1.2** shows this effect in an updated recoding and reanalysis of these data conducted for this review. Moreover, a contrast between retrieval success and encoding success showed that the ventral caudate was more reliably associated with retrieval success than encoding success across studies. Importantly, retrieval success in striatum is not dependent on an actual prior experience with an item. Rather, striatum shows greater activation for false alarms (new items incorrectly judged as old) than CR or misses (Abe et al., 2008). Thus, like most regions showing retrieval success effects (Wagner, Shannon, Kahn, & Buckner, 2005), striatal activation tracks the perception of

an item as being old during a recognition memory task, rather than it having been previously encountered on the study list. Thus, striatal retrieval success effects cannot be trivially explained based on a prior association with positive reinforcement formed at encoding.

Generally consistent with the neuroimaging data, deficits in patients with Parkinson's disease (PD) – a disease arising from degeneration of cells in the substantia nigra that are a primary source of dopaminergic input into the striatum (**Figure 1.1B**) – indicate that the basal ganglia are broadly necessary for normal levels of recognition memory performance. In particular, though not suffering from the profound amnesias accompanying MTL damage, PD patients do demonstrate deficits in recognition memory relative to controls in studies with sufficient power (Whittington, Podd, & Kan, 2000).

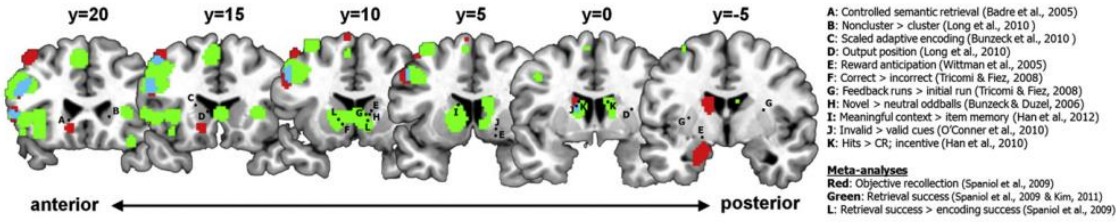

Figure 1.2. Comparisons of Activations from Meta-analyses and Neuroimaging Data Regarding Basal Ganglia Involvement in Novelty Detection, Reward Processing, and Declarative Memory. Points labeled A through K are the basal ganglia foci reported in the respective studies. Colored areas are from our recoding and reanalysis of the relevant studies from Spaniol et al. (2009) and Kim (2011) meta-analyses in GingerALE Version 2.1.1 (brainmap.org) in Talaraich space. Areas associated with objective recollection are shown in red and areas associated with retrieval success are shown in green; their overlap is marked in blue. Points labeled L are foci from the meta-analysis reported by Spaniol et al. (2009). See the main text and individual references for more information about individual contrasts. CR, correct rejections.

Accounting for these recognition deficits in PD has proven difficult and multifaceted. Across studies, deficits have been evident sometimes in recollection (Barnes, Boubert, Harris, Lee, & David, 2003; Drag, Bieliauskas, Kaszniak, Bohnen, & Glisky, 2009; Edelstyn, Mayes, Condon, Tunnicliffe, & Ellis, 2007; Edelstyn, Shepherd, Mayes, Sherman, & Ellis, 2010) and sometimes in familiarity (Davidson, Anaki, Saint-Cyr, Chow, & Moscovitch, 2006; Weiermann, Stephan, Kaelin-Lang, & Meier, 2010). Moreover, there seems ample evidence that at least a portion of memory deficits observed in PD arise from a failure to engage in effective encoding strategies (Knoke, Taylor, & Saint-Cyr, 1998; Vingerhoets, Vermeule, & Santens, 2005). However, a recent study has provided convincing evidence for a recollection deficit in PD when encoding strategy was controlled. Cohn et al. (2010) had PD patients and age-matched controls study word-pairs under shallow and deep encoding conditions, and estimated familiarity and recollection using the process-dissociation procedure (Yonelinas, Regehr, & Jacoby, 1995). Relative to shallow encoding, deep encoding normalized the patients' familiarity relative to the controls, demonstrating the efficacy of the deep encoding manipulation. However, whereas recollection improved for controls when items were deeply encoded, patients showed no improvement in recollection for deeply encoded items. As also noted by the authors, this retrieval deficit could be interpreted as a failure of the "executive" component of retrieval, such that patients did not take strategic advantage of the elaborative encoding strategy. We will return to the question of executive (i.e. cognitive control) deficits below. However, regardless of the specific source of the deficit, the evidence for a component of impaired retrieval in PD from this and prior work seems compelling.

It should be noted, however, that PD is not a selective striatal disorder, making it difficult to assign deficits to striatum specifically, as opposed to frontal disruption or dysfunctional dynamics within the broader basal ganglia system. However, recognition deficits in PD indicate that the nigra-striatal dopamine system is broadly necessary for retrieval. Moreover, declarative memory deficits have been demonstrated in a variety of disease conditions involving the nigra-striatal dopamine system such as Huntington's Disease, which is more localizable to striatum, and schizophrenia (Hodges, Salmon, & Butters, 1990; Oostrom et al., 2003; Solomon et al., 2007). Thus, when considered together with the neuroimaging data that localizes retrieval effects within the striatum, the evidence begins to converge on a necessary role for these structures during retrieval. However, as will be discussed below, this role likely relates to the way that memory retrieval is modulated by retrieval goals, as opposed to as a source of the mnemonic signal itself.

Striatal responses to item novelty

The apparent sensitivity of striatum to perceived oldness is, perhaps, surprising in light of the established association of the broader mesolimbic/nigra-striatal dopamine system with the opposite property, namely item novelty. Physiological recording studies in the rodent (Horvitz, Stewart, & Jacobs, 1997; Horvitz, 2000; Schultz, 1998) have observed activation to stimulus novelty of cells in the ventral tegmental area (VTA) and substantia nigra (SN). Importantly, novelty responses in these cells are modulated by salience and goal-relevance of the novel stimulus and are separable experimentally from the established responses of these cells to expected reward (e.g. Horvitz, 2000). Similar

effects of item novelty SN/VTA have also been observed in human fMRI studies (Bunzeck & Düzel, 2006) and are again separable from reward-related activation.

Novelty responses in the SN/VTA are hypothesized to arise via inputs from the hippocampus (Lisman & Grace, 2005), which computes the novelty of encountered items. Novelty responses in VTA neurons, in turn, are hypothesized to project back to hippocampus where they may enhance encoding of the novel item through dopaminergic modulation of hippocampal long-term potentiation (LTP). This hippocampal-VTA loop can serve the adaptive function of selectively enhancing encoding for novel items that are behaviorally relevant for the animal (Lisman & Grace, 2005; Shohamy & Adcock, 2010). Evidence from human fMRI studies have been consistent with this hypothesis, demonstrating that novelty at encoding elicits responses in VTA/SN that are associated with beneficial effects on subsequent memory (Krebs, Schott, Schütze, & Düzel, 2009; Schott et al., 2004; Wittmann, Bunzeck, Dolan, & Düzel, 2007).

Importantly, as noted above, VTA/SN cells also provide dopaminergic input into the striatum (**Figure 1.1B**) where the information they convey about expected reward and other behaviorally relevant features of an input, like novelty, can influence action selection and decision-making. For example, when harvesting reward in a stochastic environment, strategically directing exploratory behavior to novel items has the potential to glean the most new information about that environment (Badre et al., 2012; Daw, O'Doherty, Dayan, Seymour, & Dolan, 2006; Frank, Doll, Oas-Terpstra, & Moreno, 2009; Kakade & Dayan, 2002). Indeed, striatal novelty responses have been specifically associated with novelty- driven choices during economic decisions (Guitart-Masip, Bunzeck, Stephan, Dolan, & Düzel, 2010; Krebs et al., 2009; Wittmann, Daw, Seymour,

& Dolan, 2008). Moreover, many studies citing SN/VTA activation in response to novelty, also report responses to novel greater than familiar items in the striatum (Bunzeck & Düzel, 2006; Guitart-Masip et al., 2010). Notably, these activations fall in close proximity to those associated with retrieval success (**Figure 1.2**).

Thus, considered together with retrieval success effects, the evidence for novelty responses in the striatum argues against obligatory coding of item oldness in striatum as a consequence of episodic retrieval. Rather, striatal responses to episodic memory signals are likely modulated depending on the adaptive significance of "oldness" or "newness" to the animal's current actions and desired outcomes.

Two recent findings provide support for this hypothesis. Bunzeck and colleagues (2010) showed that responses in the striatum are scaled adaptively based on expectations of the relative novelty and oldness of items in the environment. Han and colleagues (2010) more directly manipulated the goal-relevance of item novelty versus oldness and revealed a similar dynamic flexibility in striatal responses. Specifically, retrieval goals were manipulated by associating either "old" or "new" responses with potential monetary reward. When "old" responses were incentivized, participants earned money for correct old responses (hits) and lost money for incorrect old responses (false alarm) and neither gained nor lost money for "new" responses (and vice versa when "new" was incentivized). Activity in the caudate tracked the incentivized response independent of whether the item was studied or novel (**Figure 1.3**). It should be noted that this pattern was seen regardless of whether or not participants received explicit feedback after their response, suggesting that striatal activity can be driven by satisfaction of internal goals and incentives.

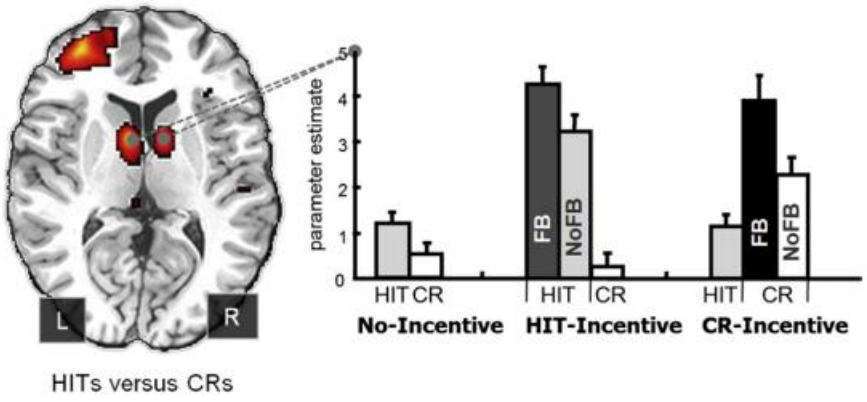

Figure 3. Striatal Activation Modulated by Incentive in Han et al. (2010). Activation in caudate tracking the incentivized response – either hits or correct rejections (CR) – regardless of whether or not explicit feedback (FB) was present, suggesting that striatal activity can be driven by incentives and goal-relevant responses (adapted with permission from Han et al., 2010).

Of course, a key question is whether these results can be reconciled with retrieval success effects, when there is no overt incentive to locate old versus new items. First, as is evident in **Figure 1.2**, the subregion of caudate that demonstrated these dynamic effects matched closely that observed across studies of retrieval success. Second, in a condition where neither response was incentivized, Han and colleagues found greater activity for hits compared to correct rejections, consistent with previous work. Similarly, striatal activity was seen for hits even when new responses were incentivized. Thus, all else being equal, participants subjectively valued "old" responses more than "new" responses when performing recognition memory tasks.

In summary, the evidence from studies of retrieval success and novelty detection indicates that striatum plays a role in the basic ability to behave according to the oldness or novelty of an item. Though in light of the qualitative differences in the severity of

memory deficits accompanying striatal versus MTL dysfunction, it is unlikely that striatum is the source of memory signals conveying oldness versus novelty, which is the province of the MTL memory system. Accordingly, as with perceptual and other inputs to the striatal system, MTL signals coding item novelty or oldness will elicit striatal responses dependent on the value of this information for current behavioral goals.

Importantly, however, goals need not be restricted to outcomes achieved through overt behavior. Rather, the process of retrieval itself can be conducted with the expectation of a particular information retrieval outcome. For example, when trying to remember a recent conversation with a good friend, we might try thinking of our friend's face as a cue. We adopt this strategy with the implicit expectation that it will yield an outcome that meets our goal, namely remembering our previous conversation. To distinguish this type of outcome from an exogenous reward or behavioral goal, we will refer to this type of desired information retrieval outcome as a *retrieval goal*. In what follows, we will argue that the striatum is particularly important for declarative memory when cognitive control is required to achieve a retrieval goal.

Striatum and the cognitive control of memory

The ability to internally modulate ongoing processing based on goals, expectations, and strategies is generally referred to as cognitive control. As introduced above, in the context of memory, cognitive control mechanisms are important for guiding and monitoring retrieval in order to achieve a particular retrieval goal. Cognitive control of memory has an established dependence on frontal lobe function, evident in the unique memory impairments of frontal lobe patients. In contrast to the profound amnesia seen in

patients with damage to the hippocampus and associated MTL areas (Scoville & Milner, 1957; Squire, 1986, 1992), patients with damage to PFC demonstrate deficits in memory tasks that involve strategic control of memory, goal- directed manipulation of mnemonic information, or overcoming interference (Aly, Yonelinas, Kishiyama, & Knight, 2011; Moscovitch, 1992; Stuss et al., 1994; Thompson-Schill et al., 1998; Wheeler, Stuss, & Tulving, 1995). Neuroimaging studies have contributed specificity, highlighting different frontal systems in support of separate control processes that contribute to these demanding retrieval tasks (J. R. Anderson, Fincham, Qin, & Stocco, 2008; M. C. Anderson et al., 2004; Badre, Poldrack, Paré-Blagoev, Insler, & Wagner, 2005; Badre & Wagner, 2007; Buckner et al., 1998; Buckner, 1996; Gallo, McDonough, & Scimeca, 2010; Kuhl, Dudukovic, Kahn, & Wagner, 2007; Long et al., 2010; Poldrack et al., 1999; Yonelinas, Otten, Shaw, & Rugg, 2005). Importantly, similar lines of neuroimaging and neuropsychological evidence also implicate the striatum in the cognitive control of declarative memory retrieval.

Cognitive control of episodic retrieval

Within the episodic retrieval domain, source memory tasks place explicit demands on cognitive control. In a source memory experiment, participants are required to verify a specific detail from a prior encoding event, such as indicating what type of task was performed with the item. In these tasks, the retrieval goal is explicit and highly specific, and so retrieval must be directed to successful recovery of only the task-relevant "source" detail to exclusion of other competing details. Thus, source memory decisions involve greater demands on cognitive control mechanisms than do simple item recognition decisions.

Contrasts between source and item recognition memory consistently locate activation in a network of frontal and parietal regions that include the striatum. In their meta-analysis, Spaniol et al. (2009) reported consistent source memory effects (i.e., "objective recollection") in left dorsal caudate, overlapping with the left dorsal striatal focus observed for retrieval success (**Figure 1.2**). In our reanalysis and recoding of these data, we found that the effects in caudate were evident both for studies contrasting correct source versus correct item decisions and those contrasting correct versus incorrect source decisions. Thus, the preferential effects of source memory in caudate were neither simply due to performing the more difficult source task nor merely to successful retrieval, irrespective of whether it was goal directed or not.

Importantly, the association of striatum with source memory relative to item decisions is not necessarily reflective of the contribution of recollection versus familiarity in these two types of tasks. Studies that have distinguished between spontaneous recollection versus familiarity during item recognition (such as is assessed by using the remember/know procedure) have not consistently located activation in the striatum when participants merely experienced recollection relative to familiarity. Direct contrast of source retrieval versus recollection during item recognition indicated that left caudate was more consistently observed across studies of source memory (Spaniol et al., 2009). Thus, the need for cognitive control, as opposed to the mere occurrence of recollection, modulates striatal activation during episodic memory retrieval.

PD patients also demonstrate consistent deficits in cognitive control of memory. In general, PD patients show greater deficits in less structured retrieval contexts, such as free recall paradigms, relative to recognition memory paradigms (Taylor, Saint-Cyr, &

Lang, 1990; Zgaljardic, Borod, & Foldi, 2003). Though likely partially arising from ineffective encoding (Knoke et al., 1998; Vingerhoets et al., 2005), their deficits on these tasks could also be traced to a failure to employ effective retrieval strategies. For example, studies using the California Verbal Learning Test (Delis, Massman, Butters, & Salmon, 1987) have shown that PD patients show decreased semantic clustering at retrieval relative to controls (Brønnick, Alves, Aarsland, Tysnes, & Larsen, 2011; Oostrom et al., 2003). Thus, deficits in recall among PD patients may partially be traced to a failure to effectively employ strategic control processes at retrieval.

Cognitive control during memory retrieval is also important to overcome interference, such as that arising through automatic retrieval of irrelevant information. PD patients show difficulty in overcoming such memory interference (Helkala, Laulumaa, Soininen, & Riekkinen, 1989; Rouleau, Imbault, Laframboise, & Bédard, 2001), though some studies have failed to find differences (Sagar, Sullivan, Cooper, & Jordan, 1991). Again, though likely partly due to encoding, these effects may also be attributable to retrieval deficits. For example, Crescentini and colleagues (2011) employed a part-list cuing paradigm designed to induce interference via external retrieval cues. PD patients and healthy age-matched controls studied separate word lists under shallow and deep encoding. Following shallow encoding, both groups showed decreased retrieval in the interference condition relative to a non-interference control. In contrast, following deep encoding, control participants showed equivalent performance in the interference and control condition, while the patients still showed impaired retrieval in the interference condition. Thus, akin to the result from Cohn and colleagues (2010) in recognition memory, this part-list cueing effect could be interpreted as a failure to effectively take

advantage of a good encoding strategy at retrieval; in this case, in order to overcome interference.

Cognitive Control of Semantic Retrieval

Striatal involvement in the cognitive control of declarative memory retrieval generalizes beyond MTL-dependent episodic memory to include semantic memory retrieval. Semantic memory refers to knowledge of facts, concepts, and word meanings that are independent of a specific encoding context and that may be stored in a distributed neocortical representation outside of the medial temporal lobe (McClelland & Rogers, 2003; Tulving, 1972). As with episodic retrieval, access to semantic knowledge can be bottom up and cue driven or it can be goal directed, requiring cognitive control (Badre & Wagner, 2007). Unlike episodic memory, however, it is easier to isolate observed effects, such as in patients, as arising during retrieval rather than encoding. Here again, the evidence generally suggests that the striatum is important for control of semantic memory retrieval.

Badre and colleagues (2005) investigated the neural systems supporting the cognitive control of semantic memory retrieval. This study focused on the contribution of left ventrolateral PFC (VLPFC) to different forms of cognitive control of memory retrieval. In a reanalysis conducted for this review, a manipulation of controlled semantic retrieval located activation in the left dorsal caudate (**Figure 1.2**). Perhaps consistent with this finding, a recent study from Han et al. (2012) found that VLPFC was preferentially engaged during a demanding retrieval task (source memory vs. item memory), but only for semantically meaningful items, suggesting that VLPFC was engaged in semantic

elaboration to enhance retrieval. The caudate showed a qualitatively identical pattern of activation. Thus, as with the Badre et al. (2005) result noted above, activation in caudate is observed under the same conditions requiring cognitive control of semantic memory that engaged VLPFC.

Consistent with the imaging data, at least one study has located interference-induced deficits in semantic retrieval in PD patients. Compared to age-matched controls, PD patients showed an impaired ability to produce a semantically related verb when presented with a noun (Crescentini, Mondolo, Biasutti, & Shallice, 2008). The deficit was greatest in a condition where there was no strongly associated response for the presented stimulus, and instead many weakly associated target verbs.

Hence, as with episodic retrieval, the striatum likely interacts with the PFC to play a causal role in the goal-directed retrieval and selection of semantic information from memory. Importantly, this suggests fronto-striatal cricuits may play a similar role in the cognitive control of both episodic and semantic retrieval. However, future research will need to test whether this common function in semantic versus episodic memory is instantiated the same or separable fronto-striatal circuits.

Hypotheses for the role of striatum in declarative memory retrieval

From the preceding review, it seems evident that the striatum plays a necessary role in optimal declarative retrieval performance, particularly under conditions requiring the cognitive control of memory. In this way, the contribution of striatum appears to mirror that of the frontal cortex during declarative memory tasks (Aly et al., 2011; Stuss et al., 1994; Thompson-Schill et al., 1998; Wheeler et al., 1995). However, research on

the neural mechanisms of cognitive control and reinforcement learning, outside of the context of memory, has suggested that striatum and frontal cortex have distinct but complementary roles (Badre & Frank, 2012; Braver & Cohen, 2000; Cools, Clark, & Robbins, 2004; Cools, 2011; McNab & Klingberg, 2008; O'Reilly & Frank, 2006). In particular, whereas lateral PFC supports cognitive control by sustaining task- relevant information in working memory (e.g. Miller & Cohen, 2001), the striatum plays a key role in flexibly updating and selecting from among candidate frontal motor or cognitive representations based on their utility or adaptive value for current goals (J. W. Brown, Bullock, & Grossberg, 2004; Cools, 2011; Frank & Badre, 2012; Gurney, Prescott, & Redgrave, 2001; Mink, 1996). In this way, fronto-striatal circuits allow for separable maintenance and updating (Hochreiter & Schmidhuber, 1997), with striatum playing a key role in mapping acquired value/utility to action selection.

Drawing on this basic cognitive control and reinforcement learning literature, we propose three hypotheses for striatal mechanisms during declarative memory retrieval: (1) Striatum modulates the re-encoding of retrieved items in accord with their expected utility (i.e., adaptive encoding), (2) striatum selectively admits information into working memory that is expected to increase the likelihood of successful retrieval (i.e., adaptive gating), and (3) striatum enacts adjustments in cognitive control based on the outcome of retrieval (i.e., reinforcement learning). These hypotheses are not intended as an exhaustive list nor are they mutually exclusive. However, each accounts for a portion of the extant data on striatal involvement in declarative memory (see **Table 1.1**) and has some limited evidence in its support.

Table 1. Evidence Regarding Striatum in Declarative Memory Retrieval

Effect	Example Reference
Neuroimaging	
Retrieval success (hits > CR) during item recognition	Spaniol et al., 2009
Perceived retrieval success (FA > CR)	Abe et al., 2008
Retrieval success > encoding success	Spaniol et al., 2009
Novelty responses during novelty detection	Bunzeck and Düzel, 2006
Adaptive scaling based on behavioral significance/incentive value of familiarity/novelty	Bunzeck et al., 2010
Incentive value of old or new memory response	Han et al., 2010
Objective recollection (source > item recognition)	Spaniol et al., 2009
Retrieval of weak semantic associates	Badre et al., 2005
Clustered > nonclustered items during free recall	Long et al., 2010
Increases with output position during free recall	Long et al., 2010
Violation of expectations during retrieval	O'Connor et al., 2010
Neuropsychology	
Recognition memory deficits	Whittington et al., 2000
Recollection deficits	Edelstyn et al., 2010
Familiarity deficits	Weiermann et al., 2010
Impaired recollection in deep LoP encoding condition	Cohn et al., 2010
Impaired familiarity in shallow LoP encoding condition	Cohn et al., 2010
Decreased semantic clustering	van Oostrom et al., 2003
Difficulty overcoming retrieval interference	Crescentini et al., 2011
Deficits in verb production based on semantic-relatedness	Crescentini et al., 2008

Hypothesis 1: Retrieval as adaptive encoding

Striatal activation during declarative memory retrieval may serve to modulate re-encoding of previously encoded items as a function of their behavioral relevance and their likelihood of future utility. The goal of memory retrieval may be expressed as the recovery of items that have an expected utility for an agent exceeding the costs associated

with retrieval itself (J. R. Anderson & Milson, 1989). From this perspective, it is important for the availability of items in memory to be prioritized by their expected utility, particularly in a given task context. Among the various cues to utility for a given item is its history of prior use: items that were retrieved in a particular context previously are more likely than others to be useful in that context again. So, retrieval itself is an important cue to the utility of an item. Indeed, analyses of retrieval that leverage prior use statistics – both in human declarative memory and in other analogous information retrieval contexts like library book withdrawals or Google search queries – account for a wide range of retrieval phenomena (J. R. Anderson & Milson, 1989; Burrell, 1980; S. Brin and L. Page, 1998, Seventh International World-Wide Web Conference (WWW 1998), conf.; Griffiths, Steyvers, & Firl, 2007). Thus, it is adaptive to have a means of prioritizing memories based on context-dependent prior utility.

Striatal dopamine signals elicited by retrieval could provide one mechanism by which memories are strengthened in accord with their context-dependent utility. It is well established that classical midbrain dopamine structures, such as the SN and VTA, along with medial prefrontal cortex, and ventral and dorsal striatum respond as a function of expected utility (e.g Knutson, Adams, Fong, & Hommer, 2001; Knutson, Fong, Adams, Varner, & Hommer, 2001; Knutson, Taylor, Kaufman, Peterson, & Glover, 2005). Projections from the midbrain to hippocampus can support modulation of hippocampal encoding by cells in these regions. For instance, dopamine can modulate synaptic change via long-term potentiation (LTP) within hippocampus, such as by decreasing LTP thresholds within CA1 fields (Jay, 2003; Lemon & Manahan-Vaughan, 2006; Li, Cullen, Anwyl, & Rowan, 2003). Thus, the nigra-striatal dopamine system is generally well

suited for coordinating dopaminergic modulation of hippocampal encoding while processing items associated with high expected utility (Shohamy & Adcock, 2010).

Recent evidence directly supports the hypothesis that the nigra-striatal dopamine system modulates hippocampal encoding as a function of the expected utility of an item (reviewed in Shohamy & Adcock, 2010), albeit not during retrieval itself. As already discussed, the hippocampal-VTA loop is thought to enhance memory for novel items in an adaptive fashion (Krebs et al., 2009; Schott et al., 2004; Wittmann et al., 2007). Moreover, two recent studies have provided evidence for dopaminergic modulation at encoding in accord with anticipated reward statistics. Wittmann and colleagues (2005) demonstrated that cues predicting subsequent reward lead to greater activation in ventral striatum and midbrain relative to pictures that did not predict reward. Moreover, activation in these striatal and midbrain regions was predictive of subsequent memory at the longer test delay for the rewarded but not the neutral pictures. By contrast, hippocampus showed subsequent memory effects for both the rewarded and neutral items and did not differentiate the two. Adcock and colleagues (2006) more directly incentivized retrieval itself, by providing participants a cue indicating that remembering an upcoming picture during a later recognition test would be worth either high or low reward. Again, regions of VTA and ventral striatum (nucleus accumbens) showed greater activation to high-reward cues. Moreover, correlation between these regions and hippocampus was positively correlated with enhanced subsequent memory. Thus, these results demonstrate that the basal ganglia can modulate hippocampal encoding to enhance memory based on an estimate of future, as opposed to immediate, utility.

Taken together, these data suggest that the basal ganglia support adaptive encoding (Shohamy & Adcock, 2010). Though theorizing has primarily focused on initial encoding, a similar adaptive encoding account could be extended to nigra-striatal involvement during retrieval, as well. As noted above, the successful retrieval of an item from memory is itself evidence that this item holds some utility in the current context. Thus, it is generally adaptive to increase the likelihood of future retrieval of that item, given an analogous context (also see Roediger & Butler, 2011). Hence, to the extent that cells in SN/VTA fire at retrieval – whether as an obligatory marker of retrieval success or reflective of the expected utility of the retrieved information for the current context – dopamine projections to hippocampus could enhance re- encoding and so prioritize items in memory as a function of their retrieval history. Importantly, descending inputs from ventral striatum to VTA (**Figure 1.1B**; Lisman & Grace, 2005) could provide modulatory input related to the adaptive value of retrieved information for current goals, providing greater contextual specificity to re-encoding.

Adaptive encoding can provide a reasonable account of a portion, though not all, of the evidence regarding striatal involvement at retrieval. Certainly, retrieval success and novelty effects in striatum, as observed with fMRI, could reflect encoding modulation in accord with the current utility of an item. Indeed, even differences between source retrieval and item familiarity/general recollection could relate to the degree of match between retrieved information and a maintained goal. However, the evidence of retrieval deficits in patient groups with basal ganglia dysfunction (e.g. Cohn et al., 2010) argues that striatum also plays a role in retrieval itself, rather than exclusively influencing future

retrieval attempts. With this in mind, we will now consider two hypotheses that propose a role for striatum in on-going retrieval.

Hypothesis 2: Adaptive gating of working memory to control retrieval

Striatum may modulate retrieval itself in accord with the expected utility of retrieval success in the current context. As noted above, if one takes the goal of memory retrieval to be recovering those items with high expected utility given the context, then cognitive control of memory is a means by which the priority of items in memory can be modified on-line to increase the likelihood of retrieving relevant information and minimize the influence of irrelevant information. In its specifics, this objective can be reached by a number of means, and indeed, there are likely several control mechanisms that operate complementarily at different stages of retrieval. For example, attention might be directed to cues in the environment that increase the likelihood of successful retrieval. Likewise, a cue might be maintained in working memory or semantically elaborated to allow it to influence retrieval. Following retrieval, monitoring of retrieved information and selection of information that matches decision criteria or behavioral goals can ensure accuracy and precision. Cognitive neuroscience research on the cognitive control of retrieval has provided a share of evidence regarding how PFC is organized to support these functions (Badre & Wagner, 2007; Dobbins & Wagner, 2005; Gallo et al., 2010; Gershberg & Shimamura, 1995; Oztekin & Badre, 2011). Importantly, however, all of these cognitive control mechanisms share a common demand to maintain a goal or relevant contextual feature in working memory in order to provide a top-down bias on current processing (Desimone & Duncan, 1995; Miller & Cohen, 2001). PFC is widely thought to support this working memory maintenance function. Also critical for working

memory is a "gate" that will let goal relevant contextual information into working memory and will keep irrelevant information out (Braver & Cohen, 2000; Hazy, Frank, & O'Reilly, 2006). The striatum may act as this working memory "gate" (O'Reilly & Frank, 2006).

As one example of how gating of working memory could influence retrieval, consider that certain cues are more likely to yield retrieval of goal-relevant information than others. Hence, maintaining those particular cues (and not others) in working memory – such as by sustaining a distributed pattern of neural activity in the PFC – provides a top-down input to the MTL system that will bias retrieval toward associates of that particular cue. At least two capacities are critical for this mechanism to operate: (1) Cues must be identified that are of potentially high expected value in the retrieval context, where here expected value is directly related to the likelihood of retrieving task relevant information, and (2) high value cues should be selectively allowed into working memory while inhibiting irrelevant or misleading cues.

As noted above, the striatum has been implicated in this type of adaptive gating of PFC to support working memory and cognitive control over action (Cools, 2011; Landau et al., 2009; McNab & Klingberg, 2008). In computational models of working memory (e.g. O'Reilly & Frank, 2006), neural networks simulate parallel cortico-striatal loops that are responsible for working memory gating, determining which representations are maintained in recurrent "PFC" layers. Based on dopaminergic learning signals, striatum learns to gate representations into PFC that lead to better outcomes (i.e., have high utility given the context) and suppress those leading to less rewarding outcomes. Once learned, gating proceeds upon encounter with a contextual input associated with high utility.

Gating itself can be accomplished through fronto-striatal- thalamic loops (Alexander, DeLong, & Strick, 1986) that modulate maintenance activity in PFC. Relative to the learning or evaluative component of this system that may be more associated with ventral striatum, this gating function may be differentially carried out by the dorsal striatum (M. X. Cohen & Frank, 2009; O'Doherty et al., 2004; Tricomi, Delgado, & Fiez, 2004).

This network architecture is generally supported by various lines of behavioral, pharmacological, and patient work (Badre & Frank, 2012; Cools, Ivry, & D'Esposito, 2006; Dahlin, Neely, Larsson, Bäckman, & Nyberg, 2008; Frank & Fossella, 2010), and computational models using this fronto-striatal "gating" network architecture have been applied to tasks involving working memory, task-switching, and contingent action selection (e.g. Frank & Badre, 2012; Moustafa, Sherman, & Frank, 2008; O'Reilly & Frank, 2006). Thus, extended to memory retrieval, cues or retrieval strategies that previous experience has associated with high expected value for retrieval could be gated into or excluded from working memory by these same frontro-striatal circuits.

The adaptive gating hypothesis is generally consistent with the evidence regarding the contribution of striatum to declarative retrieval, though limited direct evidence exists. Certainly, retrieval success effects and novelty responses could reflect an encounter with a cue or deploying a strategy that is relevant to a decision about oldness/novelty. Moreover, gating demands will increase in any retrieval context requiring more cognitive control; as contextual elements, goals, retrieval strategies, and interim products of retrieval are updated and maintained in working memory. Hence, evidence of greater striatal activation that accompanies PFC activation for source relative to item retrieval, during controlled semantic retrieval, or with increased output position during free recall

(Long et al., 2010) is broadly consistent with the gating hypothesis. Also, potentially consistent with this interpretation, one multi-modal imaging study using fMRI and SPECT reported a correlation of increased D2 receptor binding in striatum with greater left VLPFC activation during proactive interference resolution (Nyberg et al., 2009).

Retrieval deficits in patients under conditions requiring greater control could likewise be traced to ineffective working memory gating. For example, as already discussed, the recollective deficit observed in PD patients following deep encoding (Cohn et al., 2010) could reflect a failure to take advantage of an effective encoding strategy, perhaps because of a failure to gate adaptive cues or retrieval strategies into working memory that were afforded by the deep encoding task.

Thus, across neuroimaging and neuropsychological studies, the gating hypothesis is broadly consistent with striatal involvement in cognitive control of memory retrieval. However, none of the studies cited provide direct evidence for this interpretation over others. Directed future research will be required to test this hypothesis directly and to dissociate striatal updating/selection from PFC maintenance during memory retrieval.

Hypothesis 3: Reinforcement learning and adaptation of cognitive control at retrieval

Just as striatum may mark the expected value associated with anticipated retrieval in a particular context, it may also be important for adapting cognitive control based on deviations from expectations about retrieval outcome. As introduced in the preceding discussion, striatum must acquire expectations about the value of particular retrieval strategies and control representations in order to support a gating function. Likewise,

when these strategies prove to be ineffective or become obsolete, the system must revise its expectations or even shift to new strategies. In the reinforcement learning literature, the deviation of an outcome from an expectation is referred to as a reward prediction error (RPE; O'Doherty et al., 2004; Schultz, Dayan, & Montague, 1997; Sutton & Barto, 1998). In order to learn the relationship between a context, a course of action, and a particular outcome, a positive RPE reinforces a particular behavior and makes it more likely to be chosen in an analogous context in the future.

Conversely, a negative RPE punishes a particular course of action and makes it less likely to be chosen. Thus, over the course of learning, behavior incrementally converges on statistically optimal behavioral strategies given the context. In the striatum, which representations to gate into working memory and which to suppress may be learned through modulation of synaptic plasticity by dopaminergic RPE signals computed in the midbrain. For example, these signals may modulate the activity of separate populations of "Go" and "NoGo" neurons that express D1 and D2 dopamine receptors respectively (O'Reilly & Frank, 2006; Shen, Flajolet, Greengard, & Surmeier, 2008).

Applied to the cognitive control of memory, RPE could hypothetically operate in a similar manner, reinforcing or punishing selection/maintenance of a particular retrieval strategy given the context. Becker and Lim (2003) proposed a model of semantic clustering in free recall that provides an example of how RPE might drive adjustments in control of memory (see also Gorski & Laird, 2011). This model sought to simulate semantic clustering strategies during recall. Clustering was implemented by maintaining a semantic context in "PFC" working memory units where it influenced serial retrieval by the MTL/hippocampus. After each item was retrieved, it was assessed for its familiarity.

Items associated with too much or too little familiarity were judged as errors (i.e., repetitions or intrusions, respectively). Either of these errors produced a negative RPE that punished the maintenance of a particular semantic context (i.e., retrieval strategy) in PFC. When enough such errors accumulated, the category maintained in PFC shifted.

This model simulates classical semantic clustering, as well as reductions in recall due to a "frontal" challenge, namely dividing attention (Moscovitch, 1994). Importantly, the model highlights that recall itself can be affected not only by demands on maintaining a strategy but also detecting when a strategy has become obsolete and a shift is in order. Consistent with this insight, frontal patients have been shown to use fewer numbers of semantic categories for clustering than controls, even when controlling for deficits in the degree to which they retrieve semantically related items consecutively (Hildebrandt, Brand, & Sachsenheimer, 1998; Jetter, Poser, Freeman, & Markowitsch, 1986). Hence, this model illustrates that RPE could be an important signal used by the brain to adjust memory retrieval strategies.

Within the declarative memory domain, there is some behavioral evidence that participants adjust their retrieval strategies based on feedback about outcomes. Han & Dobbins (2009) manipulated explicit feedback to differentially reinforce "old" responses in a recognition memory task and found that participants become more or less likely to endorse memory probes as "old". This shift in behavior occurred gradually over the course of learning and persisted even in blocks after the feedback was removed. This suggests that participants had adaptively adjusted a latent criterion threshold that they use to evaluate retrieved mnemonic evidence. Theoretically, these adjustments could arise from an RPE, as in a mismatch between the expectations of participants regarding the

outcome of their report (old/new) and the feedback they received. Such an RPE could be computed in the striatum.

Considerable evidence has already linked the basal ganglia in general and striatum in particular to incremental adjustments in behavior in accord with RPE (though see Berridge, 2007). Classically, patients with basal ganglia disorders, like PD patients, show deficits in tasks, like the weather prediction task, in which links between a state, action, and outcome must be learned based on reinforcement (Gluck, Shohamy, & Myers, 2002; Knowlton et al., 1994; Poldrack et al., 2001). Similarly, evidence from reinforcement learning tasks that estimate learning rates in individual participants and model RPE based on a participant's specific sequence of responses and rewards has repeatedly shown that activation in ventral striatum tracks trial-to-trial changes in RPE (Badre & Frank, 2012; Daw, Gershman, Seymour, Dayan, & Dolan, 2011; Gläscher, Daw, Dayan, & O'Doherty, 2010; O'Doherty et al., 2004; O'Doherty, Hampton, & Kim, 2007). There is also some evidence that this type of reinforcement learning may influence learning of working memory gating functions by dorsal striatum (Badre & Frank, 2012; Frank & O'Reilly, 2006; Moustafa et al., 2008). Thus, RPE may play a similar role in memory control and either reinforce memory control strategies or drive changes in them in accord with the deviation from expected retrieval outcomes.

As with the gating hypothesis, the reinforcement learning hypothesis is broadly consistent with evidence linking striatum to cognitive control. Retrieval success effects could reflect the positive RPE associated with the success of a retrieval strategy (i.e., achieving a goal; e.g., Han et al., 2010). Likewise, evidence linking striatum to retrieval

tasks that place greater demands on cognitive control could reflect adjustments in control as retrieval unfolds.

More directly, there is also some limited evidence that striatal activation can vary as a function of deviations from expectation during memory retrieval. Tricomi and Fiez (2008) reported a paired-associate learning task, in which participants first learned the associations by randomly choosing between two answer choices and then receiving feedback on their accuracy. On subsequent memory trials, participants made their decisions based on their memory of the correct response from earlier trials, again receiving feedback on their performance. Caudate activation was evident on the memory trials but not the initial learning trials, suggesting that the caudate was selectively engaged when participants are expecting the feedback to provide information about the accuracy of their memory decisions. O'Connor and colleagues (2010) examined the interaction between expectations and retrieval success effects by manipulating participants' expectations of upcoming oldness in a recognition memory test. Participants saw a valid or invalid anticipatory cue ("likely old" or "likely new") before each recognition judgment. The caudate was active not only in the "retrieval success" contrast, but also in a contrast comparing invalid cue trials vs. valid cue trials, suggesting that the caudate activity may serve as a marker of the violation of memory expectations.

To summarize, there is evidence that people can take advantage of feedback to adjust their memory retrieval strategies; a process that could reasonably be assumed to rely on some form of RPE. And, there is evidence that striatal activation tracks deviations from expectation during retrieval tasks and so is potentially modulated by RPE. These observations motivate the hypothesis that RPE signals in striatum support experience-

driven adjustments in cognitive control strategies during retrieval. However, it remains to be demonstrated that these RPE signals are the source of behavioral adjustments in memory control.

Summary

There has been a growing recognition of the role of striatum across cognition, extending beyond basic motor control and being implicated in domains such as action selection, working memory, reinforcement learning, and cognitive control. Indeed, the results reviewed here suggest that striatum interacts with other brain regions, such as prefrontal cortex and hippocampus, in declarative memory retrieval. In particular, the extant neuroimaging and neuropsychological literature implicate striatum in oldness and novelty detection, goal-relevant decision processes in recognition memory, and the cognitive control of episodic and semantic memory (**Table 1.1**).

Considering these observations, it is evident that striatum plays a critical role in optimal memory retrieval, but the specific mechanistic contributions of striatum are underspecified. Drawing on existing theories of striatal function, we have proposed three possible ways in which striatum might contribute to declarative memory retrieval, namely through (1) adaptive encoding at retrieval, (2) adaptive gating of working memory to control retrieval, and (3) reinforcement learning of retrieval strategies. These hypotheses are likely not mutually exclusive. Indeed, it seems likely that all three may characterize a component of what striatum, and/or the broader basal ganglia system is supporting during retrieval. Moreover, there may be differences within striatum regarding how these hypothesized functions are supported. For example, the difference between adaptive

gating and reinforcement learning/evaluation of memory control strategies – a kind of actor-critic architecture for memory control (e.g. Bornstein & Daw, 2011; Botvinick, Niv, & Barto, 2009; Holroyd & Yeung, 2012) – could be supported by dorsal versus ventral striatum respectively.

It is also important to clarify that the hypothetical contributions of striatum to declarative retrieval performance proposed here need not exclusively support those cognitive control processes that lead to improved retrieval itself. Cognitive control can affect the accuracy and precision of retrieval, as illustrated by the examples provided above. However, cognitive control may also adjust decision criteria and response selection policy during memory tasks in order to gain positive outcomes, independently of the underlying retrieval outcomes. In other words, cognitive control may also bias reports during memory tasks as opposed to affecting discrimination, per se (Lauwereyns, Watanabe, Coe, & Hikosaka, 2002; Maddox & Bohil, 2005). And indeed, certain manipulations such as those that incentivize particular reports (old versus new, for example; Han & Dobbins, 2009; Han et al., 2010) are likely examples of adaptation occurring at this decision stage, as opposed to affecting retrieval or discrimination directly. Nevertheless, whether cognitive control mechanisms are directed toward achieving a particular retrieval goal, such as recovering a particular type of information from memory, or maximizing positive outcomes by biasing reports, striatum may play a similar role in utility-driven updating and selection of working memory representations.

Finally, it is important to note that though we have drawn an analogy between striatal function during declarative memory tasks and existing models of striatum developed outside of the memory domain, mapping value to memory signals and

processes – which is at the base of all three hypotheses – is different in important ways from typical reinforcement learning tasks that map value to a stimulus. In particular, declarative memory representations are abstract and multidimensional and are shaped by the retrieval process itself. Thus, items or contexts with very different features may elicit similar memory signals and conversely items with highly overlapping features may be treated differently depending on the nature of the memory signal being computed. Thus, in the context of memory, it is not helpful to conceptualize striatal function as mapping value to stimuli. Rather, one must consider the problem of assigning value to levels and types of mnemonic representations. Similarly, valuation itself within the memory domain is somewhat different than in traditional contexts. For example, value could be based on the match of a latent memory state to expectations, the degree of effort minimization that follows from successful retrieval, and/or the variability in retrieval outcome (akin to outcome variance in reinforcement learning; e.g., Niv, Edlund, Dayan, & O'Doherty, 2012). Hence, moving forward, it is crucial to study the contribution of striatum to declarative memory in the context of memory retrieval itself, rather than by analogy with other domains. Future directed investigations will be required to provide a more concrete view of the mechanistic role of striatum in declarative memory retrieval.

Outline of Dissertation

The remainder of this dissertation seeks to specify the role of striatum and the associated value system during declarative memory retrieval. I first address the third hypothesis outlined above, which posits that striatum enacts reinforcement learning of control policies at retrieval. To test this hypothesis, I employ a behavioral paradigm

inspired by reinforcement learning models, as well as functional neuroimaging, computational modeling, and a test of transfer.

Chapter 2 presents experiments that test the role of PE-based learning on recognition memory decisions. Recognition memory provides a straightforward paradigm in which to demonstrate the role of PEs and RL processes during retrieval. Chapter 2 first lays out a theoretical framework that describes how feedback-based PE signals could regulate retrieval control policies in the context of recognition memory. Using a behavioral false-feedback recognition task, I then show that the magnitude of PEs during recognition influences the degree to which individuals make adjustments to recognition decision policies (that is, adjustments to an individual's decision criterion). I find that this PE signal is scaled relative to a value signal derived from the memory strength underlying the recognition decision, as opposed to a value signal tied to the confidence of the decision.

In Chapter 3, I combine this false-feedback task with fMRI to identify the neural circuits involved in this PE-based adjustment of recognition decisions. As predicted by the theoretical background described above, I find that striatum tracks PEs during recognition memory and that this striatal PE signal predicts the magnitude of behavioral adjustment. Critically, the striatum tracks the PE signal when scaled relative to memory strength but not when scaled relative to response confidence, consistent with the behavioral results of Chapter 2.

To address the psychological locus of learning involved in these feedback-based adjustments, I use a test of transfer and computational modeling. Specifically, I aimed to

distinguish between a form of learning in which individuals make adjustments in a process involved in evaluating memory evidence (memory-level learning); and a form of learning in which individuals simply learn to bias their response tendencies towards the response most associated with positive outcomes (response-level learning). To test for transfer, a second experiment reported in Chapter 3 provided participants with a free recall false memory test (Kantner & Lindsay, 2012; Roediger & McDermott, 1995) after they completed the false-feedback recognition task. I find that the PE-based learning which occurs during the recognition test influences an individual's free recall performance, suggesting a task-general mnemonic-evaluation process as the locus of learning. Furthermore, across both Chapters 2 and 3, I apply drift diffusion modeling to the false-feedback recognition task and find that a model which allows for adjustments in an evidence evaluation parameter better fits the behavioral data compared to a model which allows for adjustments in a response bias parameter. Together, these transfer and modeling results implicate a memory evaluation process as the psychological locus of learning during PE-based adjustments of recognition memory decisions.

These studies provide the first evidence that a striatally-mediated reinforcement learning system serves to regulate declarative memory retrieval policies. Although memory theorists have long known that healthy individuals can flexibly adapt and adjust retrieval strategies based on internal goals and the external environment, the underlying mechanisms that support these adjustments have remained largely unspecified. PE-based learning provides a concrete mechanism by which these strategies can be evaluated and updated. This further extends RL principles beyond standard applications such as action selection or decision-making and into the domain of declarative memory. In particular,

the finding that PEs during recognition memory are scaled relative to a value signal derived from memory strength serves to highlight the importance of testing RL predictions in different cognitive domains.

Informed by the findings of Chapters 2 and 3, in Chapter 4 I turn to the first hypothesis outlined above, which posits that an important role of striatum during memory retrieval is to index valued retrieval outcomes and drive re-encoding of retrieved information in accord with this value. The results from the false-feedback recognition paradigm suggest that the nigra-striatal value system assigns value to perceived evidence of oldness, which is consistent with classic retrieval success effects (Spaniol et al., 2009; Wagner et al., 2005). However, the value of perceived oldness may be an obligatory, task-independent aspect of memory retrieval; alternatively, the value signal could reflect the goal-attainment associated with retrieving memory evidence in the context of a memory task (Han et al., 2010). The experiments described in Chapters 2 and 3 cannot dissociate these alternatives because the recognition paradigm was framed as a memory test with the tacit goal of retrieving information from memory. Furthermore, these alternatives are not exclusive: both perceived oldness and goal-attainment may result in positive valuation by the nigra-striatal system (see above; Elward, Vilberg, & Rugg, 2014; Han et al., 2010) and in turn drive re-encoding.

In Chapter 4, I use a variant of the testing effect procedure (Roediger & Butler, 2011) to look for behavioral signatures of retrieval-based re-encoding. In short, the testing effect refers to the phenomenon that a memory test results in enhanced encoding relative to repeated study, as measured by a subsequent memory test. To uncouple perceived oldness and task goals, I framed the initial test as either a memory task or a

novelty detection task. In a memory task, perceived oldness and task goals are aligned; whereas these factors are crossed in a novelty detection task in which the task goal is to detect newness. I find evidence that both perceived oldness and goal-attainment modulate re-encoding, when the novelty detection task is made sufficiently distinct from the memory task. This is consistent with previous studies implicating multiple co-occurring striatal value signals during retrieval (Elward et al., 2014) and demonstrates that both of these value signals are tied to retrieval-based re-encoding. More broadly, this behavioral finding supports the hypothesis that one function of striatal responses observed during retrieval is to index valuable retrieval outcomes and facilitate re-encoding of retrieved information.

CHAPTER 2

Learning to Remember: Prediction Errors Influence the Recognition Memory Decision Criterion[2]

Recognition memory requires flexible control over how retrieved evidence is evaluated, in order to accommodate different task goals, expectations, and environments. However, the mechanisms that support adjustments to memory decisions remain unclear. In this chapter, I test the third hypothesis introduced in Chapter 1: that striatum enacts reinforcement learning of cognitive control policies at retrieval. In the context of recognition memory, this hypothesis proposes that prediction errors (PE) – the deviation between an actual and expected outcome – drive adjustments to recognition memory decisions.

In recognition memory, the decision criterion marks the level of evidence above which an individual will decide that an item has been previously encountered. I provided false positive feedback to "old" or "new" responses in a recognition memory test, which elicited shifts in criterion toward the reinforced response. I sought to demonstrate the causal influence of PEs by systematically delivering false feedback to elicit different levels of PE. The results show that larger PEs resulted in larger shifts in decision criterion, but only when PE was scaled relative to a value signal derived from memory strength rather than the expected response outcome. Finally, to identify the psychological locus of learning, I used drift diffusion modeling. The data were better explained by a

[2] This chapter has been adapted from a manuscript submitted for publication:
Scimeca, J. M., Katzman, P. L., & Badre, D. (under review). Learning to remember: Prediction errors influence the recognition memory decision criterion.

model in which adjustments occurred at the level of evaluating memory evidence, rather than a change in simple reporting tendencies.

These results provide empirical evidence in support of the hypothesis that a reinforcement learning system evaluates and dynamically updates recognition memory decision policies at retrieval. In this framework, evidence of oldness yields an expectation of a positive outcome, and deviations from this expectation elicit adjustments to recognition decisions.

Introduction

The ability to rapidly identify previously encountered items is of high adaptive value. However, it can be equally important to avoid memory errors. Failing to recognize a previously encountered item as old (a "miss") or misidentifying a new item as old (a "false alarm") can lead to maladaptive responses (Roediger & McDermott, 1995; Schacter, Curran, Galluccio, Milberg, & Bates, 1996; Schacter, 1999). In recognition memory, the tradeoff between these two types of errors is reflected in the decision criterion. The decision criterion marks the level of retrieved memory evidence above which an individual will decide that an input has been previously encountered and will act accordingly.

The recognition decision criterion must be flexible in order to accommodate different task goals or expectations about the environment in which a person is behaving. To illustrate this point, consider the following example. You are in line at a coffee shop and notice a familiar looking face across the room. Deciding whether or not to say hello depends on whether the memory signal elicited by the person's face is sufficient to surpass your internal recognition decision threshold. Crucially, however, your decision threshold might differ depending on your goals or current context. For example, you might require less evidence of oldness if you are in a coffee shop in your hometown versus one in a foreign city. The placement of your criterion is important because the memory decision has real stakes: there may be a positive outcome (e.g. a pleasant conversation with an acquaintance) or a negative outcome (e.g. an embarrassing interaction with a stranger) depending on where you set the threshold. And, indeed, the

outcome of your decision on one occasion might make you more or less likely to approach a familiar-looking face in that context in the future.

In the laboratory, decisions of the type illustrated above are commonly tested using the old/new recognition memory task (Rugg & Yonelinas, 2003; Wixted, 2007; Yonelinas, 2002). In the conventional recognition memory paradigm, participants first study a series of items. Then, during a subsequent test phase, participants view a mix of items that were seen during encoding (old items) and new items, and must judge whether each test item is old or new. To make their response, participants must first retrieve information from long-term memory associated with the test item, such as a generalized sense of familiarity and/or richer recollective details from the prior encoding episode. Participants must then evaluate this mnemonic evidence and decide whether the evidence warrants responding that the test item is old (Stretch & Wixted, 1998a; Verde & Rotello, 2007). The focus of the present work is on this latter decision process and the means by which this process is adjusted based on the outcomes of recognition decisions.

A common framework for conceptualizing recognition memory decisions is in terms of a signal detection process (for reviews, see Banks, 1970; Macmillan & Creelman, 2005; Wixted, 2007; Yonelinas, 2002). From this perspective, all items encountered during retrieval yield some degree of memory strength. Test items that evoke a signal greater than a criterion threshold are endorsed as "old"; items that fail to elicit sufficient memory signal are rejected as "new." Optimal placement of criterion, then, depends on the signal distributions of old versus new items (**Figure 2.1A**). For a given distribution of mnemonic evidence, a criterion placed at the intersection of the "old" and "new" distributions maximizes correct responses while minimizing error

responses. Shifts in criterion away from this neutral placement will differentially increase certain types of errors. Specifically, a conservative criterion, requiring higher strength than neutral to say "old," results in more "new" responses and so a greater vulnerability to misses. A liberal criterion, requiring less strength than neutral to say "old," results in more "old" responses and thus drives a greater vulnerability to false alarms.

It is important to note that SDT is not a model of memory processes, but instead provides a straightforward framework for describing memory decision making. There are two broad classifications of models of recognition memory decisions, the first of which is typically referred to as "global matching models" or "strength-based models" (e.g. Clark & Gronlund, 1996; Dennis & Humphreys, 2001). In these models, an individual maintains a single threshold value as a reference and makes decisions by comparing retrieved memory evidence to this reference. In the example of the coffee shop, seeing an acquaintance elicits a high level of memory strength, and you say hello because this strength is higher than the reference (the criterion).

An alternative class of models are likelihood-ratio models of recognition memory (e.g. Murray Glanzer, Hilford, & Maloney, 2009; Wixted & Gaitan, 2002). In these models, a recognition decision is based on the relative probability that a given level of memory evidence would be produced by an old versus a new item. In the example of the coffee shop, seeing an acquaintance elicits a given level of memory strength, and you say hello because this level of strength has previously been elicited more frequently by old items than by new items. In these models, a criterion is not a maintained reference value but instead corresponds to the point along the evidence axis at which "old" items become

more likely than “new” items. This point occurs at the intersection of the evidence distributions and is akin to a neutral criterion in strength-based models.

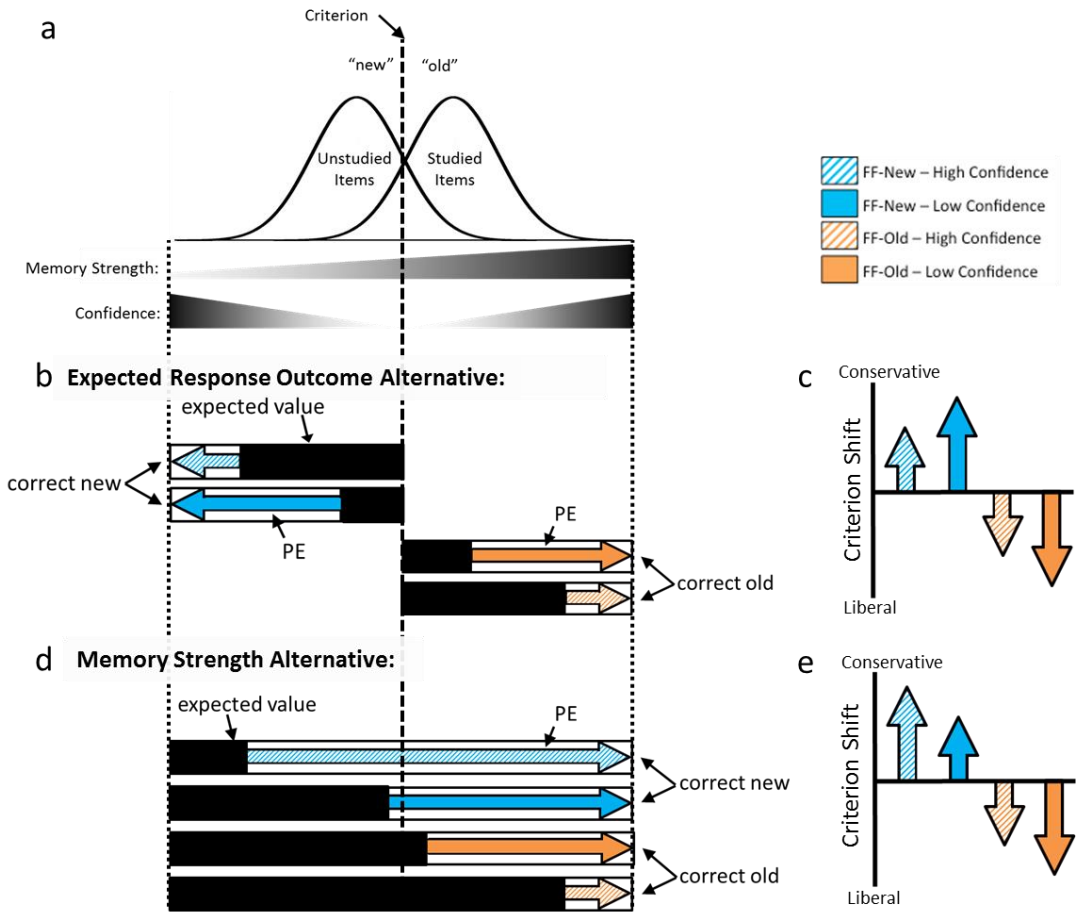

Figure 2.1. Signal Detection Memory Strength Model and Criterion Shift Predictions.
(a) Standard signal detection memory strength model of recognition memory with a neutral recognition criterion. Memory strength increases along the x-axis from left to right; confidence increases for responses further from the criterion. The legend indicates four different types of responses, which map directly to the four false feedback groups. (b) The expected values and PEs associated with positive feedback under the Expected Response Outcome alternative. Scale of bars is directly mapped to the curves depicted in panel a. (c) The predicted criterion shifts associated with the Expected Response Outcome alternative. (d) The expected values and PEs associated with positive feedback under the Memory Strength alternative. Scale of bars is directly mapped to the curves depicted in panel a. (e) The predicted criterion shifts associated with the Memory Strength alternative.

Both strength-based models and likelihood-ratio models make similar (and often identical) predictions with regard to recognition decisions. Likewise, both classes of models are well-described by signal-detection theory analyses of response rates. Therefore, we will use the terms *criterion placement* and *criterion shift* throughout the chapter to describe a change in old/new responding that results in a change in the SDT estimate of criterion. However, this terminology is not meant to only imply that the underlying mechanism is a shift in a maintained reference value, as assumed by the global matching model. In the General Discussion, we will describe how the present results relate to these various models of recognition memory.

From the perspective of an adaptive memory system, flexibly adjusting criterion allows one to maintain high levels of memory performance in the face of a changing environment. Although most theoretical models assume that people have control over their criterion location and can update it to optimize performance, empirical results are mixed. Furthermore, it remains an open question as to what signals are used to dynamically adjust criterion.

Adjustments of Recognition Criterion

Criterion is not fixed, and people are capable of exerting direct control over their recognition memory criterion when they are explicitly instructed to do so. Participants instructed to adopt a conservative or liberal criterion will comply with the instruction in their recognition decisions (Postma, 1999). Likewise, participants explicitly informed of differences in the proportions of old and new items or provided explicit reward incentives for particular responses will adjust their criteria accordingly (Curran, DeBuse, & Leynes,

2007; Ratcliff, Sheu, & Gronlund, 1992; Van Zandt, 2000). However, in our everyday lives we are often not given explicit information about how to set our criterion. Rather, we must learn to set our criterion based on internal mnemonic signals, environmental regularities, or experienced outcomes that follow from our decisions. Nevertheless, in the laboratory there is mixed evidence regarding the degree to which participants are able to shift criterion based on these types of cues.

The expected mnemonic strength of target items can serve as a signal for setting criterion. Many studies have observed a strength-based mirror effect (Bruno, Higham, & Perfect, 2009; Cary, 2003; Glanzer & Adams, 1990; Stretch & Wixted, 1998b). Specifically, if the average strength of test items is higher (effectively shifting the mean of the old distribution rightward in **Figure 2.1A**), participants will adopt a more conservative decision criterion overall, reducing their false alarm rate while still maintaining a reasonable hit rate. They do so even when they are not provided explicit instructions regarding the strength of the old items (Verde & Rotello, 2007). Hence, participants do attend to the overall strength of test items and regulate their criteria accordingly. However, the criterion appears to be resistant or "stubborn" within a memory test: for example, participants do not typically show criterion shifts when the underlying distribution of memory strength for old items is changed partway through an experiment (Hicks & Starns, 2014; Verde & Rotello, 2007; but see Brown, Steyvers, & Hemmer, 2007).

Providing overt feedback about the accuracy of recognition decisions can sometimes help participants take advantage of changes in memory strength distributions to adjust their criteria. For example, Verde and Rotello (2007) found that participants

only adaptively shifted their criteria with changes in the distribution of memory strength when external feedback was provided. In contrast, Hicks and Starns (2014) manipulated distributions of memory strength and found minimal differences between conditions with or without external feedback. We return to these two studies in more detail in the General Discussion to discuss how the ostensibly conflicting results may be reconciled in light of our findings.

Similarly, the provision of feedback can sometimes help participants take advantage of differences in the base rates of test items. Participants are relatively insensitive to differences in the base rate of old or new items (which would be reflected as a difference in heights in the two distributions depicted in **Figure 2.1A**) in the absence of explicit instruction (Cox & Dobbins, 2011; Koop, Criss, & Malmberg, 2014). However, participants do adjust criterion in accord with base rate differences when they are provided with explicit feedback about their responses (Kantner & Lindsay, 2010; Koop et al., 2014).

The above examples suggest that people can use feedback to regulate their recognition criterion, particularly in contexts where base rates or strength distributions are highly divergent. However, the extant literature suggests that the above examples are more the exception than the rule, and feedback is not universally effective for eliciting optimal shifts in criterion. Logically, a participant should adopt a criterion that optimally balances the trade-off between misses and false alarms (all else being equal, this corresponds to a neutral criterion). Unless individuals adopt a perfectly optimal criterion in the absence of feedback (which is rarely the case), feedback should provide signals that allow the individual to shift the criterion towards a more beneficial position. Instead,

explicit feedback about the accuracy of responses often has little effect on criterion placement under standard recognition memory circumstances, when there are no drastic differences in strength distributions or base rates (Han & Dobbins, 2008, 2009; Kantner & Lindsay, 2010; Verde & Rotello, 2007). Thus, the underlying mechanisms that determine when and to what degree feedback impacts recognition decisions remain unclear.

The variable impact of feedback during recognition memory is perhaps surprising given the importance of outcomes in other domains of decision making. Outcomes are highly informative learning signals and reinforcement is a central component of many theories of learning and decision making outside of the memory domain (O'Doherty, Dayan, Friston, Critchley, & Dolan, 2003; O'Doherty et al., 2004; Schultz, Dayan, & Montague, 1997). Moreover, corrective feedback is known to be important in other types of memory beyond recognition, such as categorization (Ashby & Maddox, 2005; Stewart, Brown, & Chater, 2005), free recall (e.g. Pashler, Cepeda, Wixted, & Rohrer, 2005), and real-world education settings (e.g. Kang, McDermott, & Roediger, 2007).

One experimental manipulation that consistently produces dynamic shifts in criterion in standard recognition tasks (i.e., without manipulated changes in the underlying strength distributions or base rates) is the provision of false feedback to participants. Han and Dobbins (2008, 2009) repeatedly elicited significant criterion shifts by providing false positive feedback on trials where participants made an error. For example, one group of participants received false positive feedback on all misses, and veridical feedback on all other types of responses. Another group of participants received false positive feedback on all false alarms, and veridical feedback on all other types of

responses. The criterion adopted by these groups differed significantly, with the former group adopting a more conservative criterion and the latter group adopting a more liberal criterion. A group receiving only veridical feedback failed to significantly shift criterion across the experiment. These results indicate that participants do consider feedback about outcomes when setting their criterion, even when there are no changes in memory strength distributions and base rates throughout the experimental session.

Thus, the mixed effect of feedback on recognition decisions calls for a reexamination of the underlying model and the nature of the feedback. Here, we will propose that the diverse prior results can be accommodated within a general reinforcement-learning framework. We will provide evidence that an individual's recognition criterion is influenced by Prediction Errors – that is, the impact of feedback is scaled as a function of an individual's expectations.

Reinforcement Learning and Regulating Recognition Memory Decision

The effect of false feedback on recognition memory is initially consistent with a reinforcement learning (RL) process (Han & Dobbins, 2008, 2009; Scimeca & Badre, 2012). Reinforcement learning models were developed in computer science, leveraging learning principles from Behaviorist psychology (Sutton, 1988; Sutton & Barto, 1998) to describe normative solutions by which artificial agents can acquire actions that maximize outcomes in simplified environments. Recently, reinforcement learning concepts have received renewed attention in psychology and neuroscience, and have been successful in explaining a wide variety of behavioral and neuroscience data (Frank, Moustafa,

Haughey, Curran, & Hutchison, 2007; Frank, Seeberger, & O'Reilly, 2004; Gläscher et al., 2010; O'Doherty et al., 2004, 2003; Schultz et al., 1997; Schultz & Dickinson, 2000).

The key signal that drives learning in RL is the Prediction Error (PE) signal. The PE signal reflects the degree to which the actual outcome of an action deviates from the expected outcome, and larger PEs result in greater learning (O'Doherty et al., 2004; O'Reilly & Frank, 2006; Schultz et al., 1997). In order to learn the relationship between a context, a course of action, and a particular outcome, a positive PE (i.e., receiving a better outcome than expected) reinforces the chosen action in a given context and makes it more likely to be chosen in an analogous context in the future. Conversely, a negative PE (i.e., receiving a worse outcome than expected) makes a course of action less likely to be chosen in the future. In simplified environments, RL agents will incrementally converge on statistically optimal behavioral strategies for a given context.

Several previous studies have suggested a role for reinforcement learning in declarative memory decisions (Cox & Dobbins, 2011; Koop et al., 2014; Mickes, Hwe, Wais, & Wixted, 2011; Wixted & Gaitan, 2002). Specifically, Han and Dobbins (2008, 2009) noted several reasons that the effect of false feedback on shifts in the decision criterion are consistent with typical characteristics of an RL system. (1) The shifts in criterion were demonstrated through behavior and not conscious knowledge: the majority of participants did not report being aware of the feedback manipulation or their shifting criterion in post-test questionnaires. (2) The shifts in criterion were incremental, gradually increasing in magnitude as more trials accumulated. (3) The shifted criterion persisted even when feedback was removed from the experiment. However, a central untested hypothesis is that PEs are the primary learning signal that drive dynamic shifts

in criterion. The primary purpose of the present study is to test whether recognition criteria are evaluated and adjusted in a manner consistent with PE-based learning in an RL framework.

Participants in the present study completed a recognition memory test wherein they made old/new recognition memory decisions and provided a fine-grained continuous rating of their confidence. Our approach is intended to frame recognition memory decisions in the same framework as value-based and perceptual decision-making models common to the RL literature, and we sought to adopt the conventions of these models. In a typical RL model, an individual makes a decision with some expectation of a positive outcome, which we will term the Expected Value (EV) of the decision. A critical assumption of our framework is that we can derive this EV for a given decision from an individual's response and confidence rating. If this is the case, then the magnitude of the resulting PE – the difference between the outcome and the EV – should impact the magnitude of criterion shifts.

Feedback in a typical recognition task can take two possible forms: positive feedback corresponds to an outcome value of 1 (the presence of a rewarding outcome) and negative feedback corresponds to a value of 0 (the absence of a rewarding outcome). Similarly, the Expected Value associated with a decision is scaled from 0 (low expectation of a positive outcome) to 1 (high expectation of a positive outcome). From this it follows that a Prediction Error ranges from -1 (a large negative PE, punishing the decision that resulted in the PE) to +1 (a large positive PE, reinforcing the decision that resulted in the PE).

Critically, the correspondence between recognition decisions, confidence, and EV is unknown. Therefore, we considered two alternatives for how the EV, and consequently the PE, is scaled during recognition memory decisions. The first alternative is that the EV derives from the "Expected Response Outcome". Under this alternative, the EV of a decision is directly linked to the confidence rating: low confidence ratings correspond to low EVs, and high confidence ratings correspond to high EVs. In other words, the participants are rating how confident they are that their response, whether old or new, will receive positive feedback. This view is most directly analogous to the way EV is typically conceptualized in conventional RL. **Figure 2.1B** shows the different EVs (black bars) and PEs (colored arrows) associated with positive feedback following four example responses under this alternative. **Figure 2.1C** shows the direction and magnitude of the resulting criterion shifts predicted by this alternative.

Under this alternative, the EV is essentially the degree to which the participant expects to be correct (and so receive a positive outcome) on a given trial. It is quite plausible that people could generate such expectations given that individuals are accurate in metacognitively predicting the accuracy of their recognition decisions (correct responses are made with higher confidence than incorrect responses, e.g. Stretch & Wixted, 1998). The magnitude of a positive PE (and so the degree of criterion shift), then, would be negatively correlated with their response confidence for both "old" and "new" response types, being largest when confidence is low and smallest when confidence is high.

A second alternative that is perhaps unique to the case of a memory task is that the EV may be linked to the retrieved Memory Strength on a given trial. Under this

alternative, both Memory Strength and EV increase from left-to-right along the horizontal axis of the standard SDT figure, as shown in **Figure 2.1A**. **Figure 2.1D** shows the different EVs (black bars) and PEs (colored arrows) associated with positive feedback following four example responses under this alternative. **Figure 2.1E** shows the direction and magnitude of the resulting criterion shifts predicted by this alternative. Under this alternative, memory strength or evidence of oldness is inherently valued, akin to a Pavlovian or response-independent predictor of future positive outcomes, with higher memory strength predicting better future outcomes.

Given this pattern of EVs, PEs would follow a pattern such that positive feedback on a low memory strength trial would be associated with higher positive PE than on a high memory strength trial. Hence, the lowest memory strength is associated with high confidence "new" responses, and the highest memory strength is associated with high confidence "old" responses. Put another way, the Expected Value of a recognition decision is negatively correlated with confidence ratings for "new" responses and positively correlated with confidence ratings for "old" responses. Note that this prediction for "new" responses is the inverse of what is predicted by the Expected Response Outcome alternative.

The Memory Strength alternative is consistent with theories positing that successful memory retrieval can be inherently rewarding and predictive of positive goal-relevant outcomes (Han et al., 2010; Johnson, Hashtroudi, & Lindsay, 1993). In the neuroimaging literature, successful memory retrieval is consistently shown to recruit value-related brain regions like the striatum (Elward et al., 2014; Han et al., 2010; Scimeca & Badre, 2012; Wagner et al., 2005). Similarly, classic observations from social

psychology, such as the "mere exposure effect," suggest that people exhibit approach responses and preferences for familiar items, again consistent with an inherently valued "oldness bonus" (Whittlesea & Price, 2001; Zajonc, 1968, 2001). We further discuss how the Memory Strength alternative relates to signal-detection theory and existing models of recognition memory in the General Discussion.

Summary and Logic of the Present Study

In summary, the present study aimed to test the role of learning based on PEs in the support of dynamic shifts in recognition decision criterion. We hypothesized that a reinforcement learning system that uses feedback-based PEs as a learning signal is responsible for regulating recognition decision criterion. A critical assumption of our approach is that the Expected Value of a recognition decision can be derived from an individual's recognition decision and confidence rating. We considered two alternative mappings between response confidence and Expected Values, which we termed the Expected Response Outcome (ERO) alternative and the Memory Strength (MS) alternative. This allowed us to test for a causal role of PEs in dynamic criterion shifts, and to directly pit the Memory Strength and Expected Response Outcome alternatives. As will be shown, the results are consistent with the Memory Strength alternative.

We used a modified version of the false feedback recognition procedure (Han & Dobbins, 2008, 2009) that allowed participants to provide their old/new recognition response while also providing a concurrent, continuous confidence rating. This procedure allowed us to estimate trial-by-trial EVs and PEs. Feedback was manipulated across three experimental groups. One group received completely veridical feedback (VF group). One

group, termed the False Feedback – New (FF-New) group, received false positive feedback on a subset of their "new" response errors (misses) and veridical feedback on all other trials. The third group, termed the False Feedback – Old (FF-Old) group, received false positive feedback on a subset of their "old" response errors (false alarms) and veridical feedback on all other trials. Following Han & Dobbins (2008; 2009), we predicted that false feedback would result in an increased propensity to make the reinforced decision: the FF-New group should respond "new" more frequently (a conservative criterion shift), and the FF-Old group should respond "old" more frequently (a liberal criterion shift).

In order to experimentally manipulate the magnitude of PEs, we further targeted false feedback to either low or high confidence ratings using an adaptive algorithm that accounted for individual differences in confidence and response rates. That is, a subgroup of the FF-New group received false feedback targeted to Low Confidence "new" responses (FF-New-LC) and a separate subgroup of participants received false feedback targeted to High Confidence "new" responses (FF-New-HC). Likewise, the FF-Old group was divided into subgroups that received false feedback targeted to either Low Confidence (FF-Old-LC) or High Confidence (FF-Old-HC) "old" responses.

If recognition decisions are regulated through PE-based reinforcement, we predicted that the magnitude of the criterion shift would be dependent on whether participants received false feedback targeted to Low Confidence or High Confidence responses. This would suggest that criterion shifts are driven by the magnitude of feedback-based PEs. If regulating recognition decisions does not depend on PE-based

reinforcement, we would predict no difference between the two FF-New subgroups and the two FF-Old subgroups.

Furthermore, the specific pattern of criterion shifts allows us to differentiate between the Expected Response Outcome alternative and the Memory Strength alternative. Both alternatives predict that the FF-Old – Low Confidence subgroup will show a larger criterion shift than the FF-Old – High Confidence subgroup (orange arrows in **Figure 2.1)**. However, the two alternatives make opposite predictions for the FF-New participants (blue arrows in **Figure 2.1**). Specifically, the Expected Response Outcome hypothesis predicts that false positive feedback following low confidence "new" responses will result in higher PEs than high confidence "new" responses (**Figure 2.1B**), and so the FF-New-LC subgroup will show the largest conservative criterion shift (**Figure 2.1C**). In contrast, the Memory Strength alternative predicts that false positive feedback following high confidence "new" responses will result in high PEs (**Figure 2.1D**). Thus, this hypothesis predicts that the FF-New-HC subgroup will show the largest conservative criterion shift (**Figure 2.1E**).

In a final set of analyses, we sought to verify the locus of learning in the false feedback recognition paradigm. By fitting drift diffusion models of decision making across multiple independent data sets, we specify that the PEs elicited by the false feedback procedure shift the memory evidence criterion rather than a response-level reporting bias.

Method

Participants**.** Eighty participants (45 females; age, 18-27 years; mean, 20 years) participated in exchange for course credit or payment. All participants had normal or corrected-to-normal vision and were native English speakers. Participants gave written informed consent according to guidelines established and approved by the Human Research Protections Office of Brown University. After completing the experimental session, participants were thoroughly debriefed with regard to the false feedback manipulation. Participants were randomly assigned to experimental conditions and experimenters were blinded to this assignment.

Stimuli. The stimulus set consisted of 772 nouns naming concrete objects. For each participant, a random subset of 640 of the words was selected to be used as study and lure items. The words were randomized across study and lure lists for each participant. About half of the words referred to organic objects and the rest referred to inorganic objects. About half of the words referred to objects smaller than a typical-sized shoebox and the rest referred to objects larger than a typical-sized shoebox.

Procedure. During an initial study phase, participants made semantic judgments for 240 words. On each study trial, a word was presented on the screen in black Courier New font for 300ms. Participants were next presented with one of two semantic decision task cues, indicating whether they should make an organic judgment (organic/inorganic) or a size judgment (small/large, relative to a typically-sized shoebox) about the preceding item. Participants made their response by pressing either the "1" or "2" key on a keyboard. Participants were instructed to make their semantic judgment as quickly and as accurately

as possible, and were given one second to respond to each semantic decision cue. If participants did not respond within one second, a beep indicated that they should try to respond faster. Trials were separated by a 1-8 second inter-trial interval. For reasons unrelated to the goals of the present experiment, on each trial, the background was one of three colors (yellow, red, or green). The order of study words were randomized across all participants and across the two semantic decisions.

The study phase was divided into three blocks of 80 items each. Between study blocks, participants were allowed to rest and press a key when they were ready to continue. Before the experiment, the participants practiced the semantic decisions on a separate set of example object words with the experimenter. Practice words were never tested during the recognition phase. At no point during the study phase were participants informed that their memory would be later tested for the words they saw during the semantic judgment task. The study phase was designed as part of an unrelated study on semantic decisions and we do not discuss it further in the present chapter. After the study phase, the participants completed a probabilistic selection task (Frank et al., 2004). The results from this intermediate task are not discussed here.

Participants returned approximately 24 hours later to complete the recognition test. In the test phase, participants were told that they would complete a memory test for the words they had seen in the first phase of the experiment. The test phase consisted of 288 recognition trials and was divided into six blocks consisting of 48 trials, each consisting of equal old and new items. Between blocks, participants were allowed a self-paced break. Participants were instructed to respond "old" if they recognized the word from the first phase of the experiment, and "new" if they did not recognize the word.

Participants indicated their response using the left or right arrow key: the left arrow key always corresponded to "old" responses and the right arrow key always corresponded to "new" responses. After the test item appeared on the screen, participants had 1.5 seconds to initiate their response by pressing a key. RT was coded based on the time from item presentation to the onset of the button response.

Participants were further instructed to indicate how confident they were in their response by holding down the response key to fill in an open rectangle displayed underneath the chosen response (**Figure 2.2**). The rectangle started filling as soon as the button was pressed and stopped as soon as the button was released. The proportion of the rectangle filled indicated subjective confidence: low confidence was indicated by filling a low proportion of the rectangle and high confidence was indicated by filling a high proportion of the rectangle. The highest subjective confidence rating corresponded to completely filling the rectangle and took 1.6 seconds from the time of the initial button press. The rectangle was 300 pixels high and confidence was recorded as the proportion of the total height that had been filled. Regardless of how long it took the individual to initiate their response, they were always provided the full 1.6 seconds to indicate their confidence.

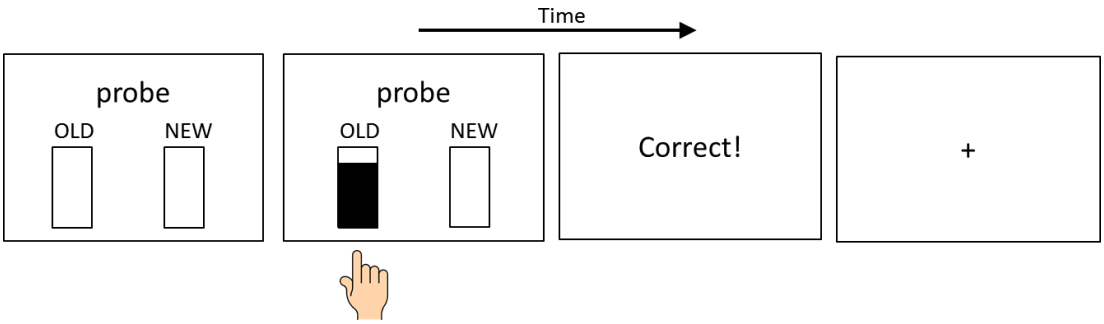

Figure 2.2. Stimulus display of recognition trial. Participants provided a continuous confidence rating by filling one of two confidence boxes concurrent with their old/new decision and received feedback immediately following the response.

We should note that this approach has two advantages over prior procedures in which the recognition confidence ratings are made as a separate response after the old/new response. First, we can obtain a more fine-grained measure of confidence on each trial, and use this confidence to test for the impact of PEs in this task. Second, when multiple actions (e.g., old/new response and confidence rating) precede a reinforcing outcome (feedback), there is an issue of temporal credit assignment: how to assign credit across multiple actions. A basic solution is to assign credit in a temporally discounted fashion, such that actions nearest in time to the reinforcing reward are assigned the most credit, and actions further in time are assigned less credit (Montague, 1996; Sutton, 1988). By collapsing the old/new decision and confidence rating into a single response, we help mitigate this temporal credit assignment problem to some degree.

Immediately after participants released the response key, feedback about the accuracy of the response was displayed on the screen for one second (“Correct!” or “Incorrect!”). This was followed by a fixation cross of a variable duration so that the total trial length (starting from the initial stimulus presentation) was 5 seconds. If a participant

did not make her/his initial button press response within the 1.5 seconds deadline, the feedback for the trial displayed "No Response!"

False feedback manipulation. Participants were assigned to one of five experimental conditions. The Veridical Feedback (VF; n = 16) group received completely veridical feedback. The other four conditions received false positive feedback targeted towards either low or high confidence errors. The False Feedback – New group received false positive feedback on a subset of "new" errors and veridical feedback on all other trials. This group was divided into Low and High Confidence subgroups (FF-New-LC and FF-New-HC; n = 16 for each subgroup). The False Feedback – Old group received false positive feedback on a subset of "old" errors and veridical feedback on all other trials. This group was likewise divided into Low and High Confidence subgroups (FF-Old-LC and FF-Old-HC; n = 16 for each subgroup).

The method for targeting false positive feedback employed an adaptive algorithm that took into account individual differences in confidence ratings and base rates of errors. This information was used to provide approximately equal instances false feedback to each participant, provide the false feedback equally across the course of the entire recognition test phase, and maximize the difference in the targeted confidence between subgroups. The script was designed to provide each participant approximately 30 instances of false feedback. The script tracked online the participant's error rate for each response type (false alarms and misses.) For each trial in the recognition test, the experimental script estimated how many more errors the participant would make and determined the proportion of these errors that would need to receive false feedback for the participant to receive approximately 30 total false feedback instances. This proportion

was updated on a trial-by-trial basis for each participant and was used to probabilistically provide false positive feedback.

For example, consider a FF-Old participant who the script estimated would make 36 instances incorrect "old" responses. Approximately 83% of her "old" errors would receive false positive feedback, in order to provide 30 total instances of false feedback. For the first four error trials, the script assumed an error rate of 36 errors (approximately equal to the error rate seen in pilot data). For the fifth and subsequent errors, the script used the error rate determined online specific to the participant.

The experimental script also tracked the distribution of confidence used for each response type ("old" or "new"), on a participant-by-participant basis. This distribution was used to determine if a given trial would receive false feedback. If the script was aiming to provide false feedback on *X* percent of the remaining error responses, the script would determine whether the participant's confidence on a given trial fell in either the lowest or highest *X* percent of the participant's confidence distribution, depending on the condition. This distribution and cut-off point *X* was updated on a trial-by-trial basis for each participant.

Continuing the previous example, the script might determine that 83% of the remaining "old" response errors should receive false feedback. For a participant in the FF-Old-HC group, an "old" response error given confidence rating above the 17th percentile mark in her specific confidence distribution would receive false positive feedback. An "old" error given a confidence rating below the 17th percentile mark in her confidence distribution would receive veridical negative feedback. This would ensure

that false feedback was provided with the appropriate frequency (as estimated by the script) and preferentially targeted to high but not low confidence "old" errors for this participant. Conversely, for a participant in the FF-Old-LC group, an "old" error given a confidence rating below the 83rd percentile received false positive feedback; an "old" error given a confidence rating above the 83rd percentile mark received veridical negative feedback. See Section I of the *Supplemental Material* for manipulation checks and discussion regarding the distribution of false feedback in this experiment.

Signal Detection Theory (SDT) and Statistical Analysis Methods. In all experiments, d-prime (*d'*) was calculated as the difference between z-scored hit rates and z-scored false alarm rates (Green & Swets, 1966; Macmillan & Creelman, 2005). Decision criterion (*c*) was calculated as negative one-half of the sum of z-scored hit rates and z-scored false alarm rates (Green & Swets, 1966; Macmillan & Creelman, 2005). Following from Han and Dobbins (2009), we divided our analysis of the test phase into blocks in order to investigate the pattern of criterion over time. We calculated criterion for an initial criterion, middle criterion, and terminal criterion; thus each calculation was based on the hit rate and false alarm rate across 96 trials.

Because all analyses of criterion were *a priori* planned test, we considered statistical results as significant at an alpha level of $p = 0.05$. Because we did not have strong *a priori* predictions about additional analyses regarding accuracy and confidence, these analyses were Bonferroni corrected for multiple comparisons.

Results

SDT analysis: Accuracy

Participants performed well at the task overall (d-prime $M = 0.94$, $SD = 0.40$). Hit rate, false alarm rate, and d-prime scores as a function of group and time are provided in **Table 2.1**. Accuracy estimates from the four feedback conditions were submitted to a three-way ANOVA with factors of feedback group (FF-New / FF-Old), targeted confidence (Low Confidence / High Confidence), and time (Initial / Middle / Terminal). D-prime decreased over the course of the experiment, supported by a main effect of time ($F(2,120) = 6.135$, $p = .003$, $\eta_p^2 = .093$). This effect may be due to fatigue, and/or criterion shifts in the false feedback groups. No other main effects or interactions reached significance.

		Time		
		Initial	**Middle**	**Final**
Criterion	**Veridical**	0.15 (0.08)	0.04 (0.09)	0.00 (0.11)
	FF-New - High Confidence	0.13 (0.09)	0.26 (0.09)	0.39 (0.09)
	FF-New - Low Confidence	0.16 (0.06)	0.18 (0.08)	0.10 (0.07)
	FF-Old - High Confidence	0.03 (0.07)	-0.08 (0.08)	-0.26 (0.08)
	FF-Old - Low Confidence	-0.12 (0.09)	-0.33 (0.10)	-0.39 (0.09)
D-Prime	**Veridical**	0.99 (0.18)	0.93 (0.15)	0.82 (0.13)
	FF-New - High Confidence	1.14 (0.11)	1.06 (0.16)	0.93 (0.10)
	FF-New - Low Confidence	1.02 (0.12)	1.01 (0.1)	0.92 (0.12)
	FF-Old - High Confidence	1.15 (0.09)	1.04 (0.11)	0.92 (0.09)
	FF-Old - Low Confidence	0.84 (0.13)	0.67 (0.1)	0.61 (0.11)
Hit Rate	**Veridical**	0.63 (0.04)	0.65 (0.04)	0.64 (0.05)
	FF-New - High Confidence	0.66 (0.03)	0.60 (0.04)	0.53 (0.03)
	FF-New - Low Confidence	0.63 (0.03)	0.62 (0.03)	0.63 (0.03)
	FF-Old - High Confidence	0.70 (0.03)	0.72 (0.03)	0.75 (0.03)
	FF-Old - Low Confidence	0.69 (0.03)	0.74 (0.03)	0.75 (0.02)
FA Rate	**Veridical**	0.28 (0.04)	0.33 (0.04)	0.35 (0.04)
	FF-New - High Confidence	0.27 (0.04)	0.24 (0.04)	0.22 (0.03)
	FF-New - Low Confidence	0.26 (0.03)	0.26 (0.04)	0.30 (0.03)
	FF-Old - High Confidence	0.29 (0.03)	0.34 (0.04)	0.42 (0.04)
	FF-Old - Low Confidence	0.39 (0.04)	0.50 (0.05)	0.53 (0.05)

Table 2.1. Criterion, D-Prime, Hit Rate, and False Alarm Rates for Main Experiment.
Values in parentheses represent standard error of the mean. FA: false alarms.

SDT analysis: Decision criterion.

Criterion estimates from all conditions are shown in **Table 2.1** and **Figure 2.3**. To assess the role of PEs in criterion shifts, decision criterion estimates from the four false feedback conditions were submitted to a three-way ANOVA with factors of feedback group (FF-New / FF-Old), targeted confidence (High Confidence / Low Confidence) and time (Initial / Middle / Terminal). We replicated the main effect of false feedback: the FF-New groups adopted a more conservative criterion and the FF-Old groups adopted a

more liberal criterion ($F(1,60) = 30.551$, $p < .001$, $\eta_p^2 = .337$; see Section II of the *Supplemental Material* for further replication in two additional data sets).

In addition to targeting false feedback to either low or high confidence responses, we titrated false feedback so that it was provided at an approximately equal rate across the entire experimental session and in approximately equal numbers to all participants. Having controlled the provision of feedback in this way, we found evidence of incremental learning: the difference between the FF-New groups and FF-Old groups increased over the course of the experiment (group by time interaction, $F(2,120) = 14.415$, $p < .001$, $\eta_p^2 = .194$).

To specifically assess the role of PEs, we next analyzed the effect of the targeted confidence manipulation. Note that any difference between the Low Confidence and High Confidence subgroups supports a role for PE-based reinforcement in regulating recognition decisions; further, the specific pattern of results allows us to discriminate between the Expected Response Outcome and Memory Strength alternatives. As shown in **Figure 2.3**, our results support the Memory Strength alternative: the largest criterion shifts were observed in the FF-New – High Confidence and FF-Old – Low Confidence groups. This alternative is supported statistically by a main effect of targeted confidence: the effect of PEs under the Memory Strength alternative manifests as a main effect because the High Confidence groups show a more positive criterion than the Low Confidence groups. This holds for both the FF-New (High Confidence shows larger positive shift) and FF-Old groups (High Confidence shows smaller negative shift). Across the whole experiment, we found a main effect of targeted confidence ($F(1,60) =$

4.042, $p < .05$, $\eta_p^2 = .063$). Consistent with the Memory Strength alternative, there was no significant group by confidence interaction ($F(1,60) = 0.200$, $p = .656$, $\eta_p^2 = .003$).

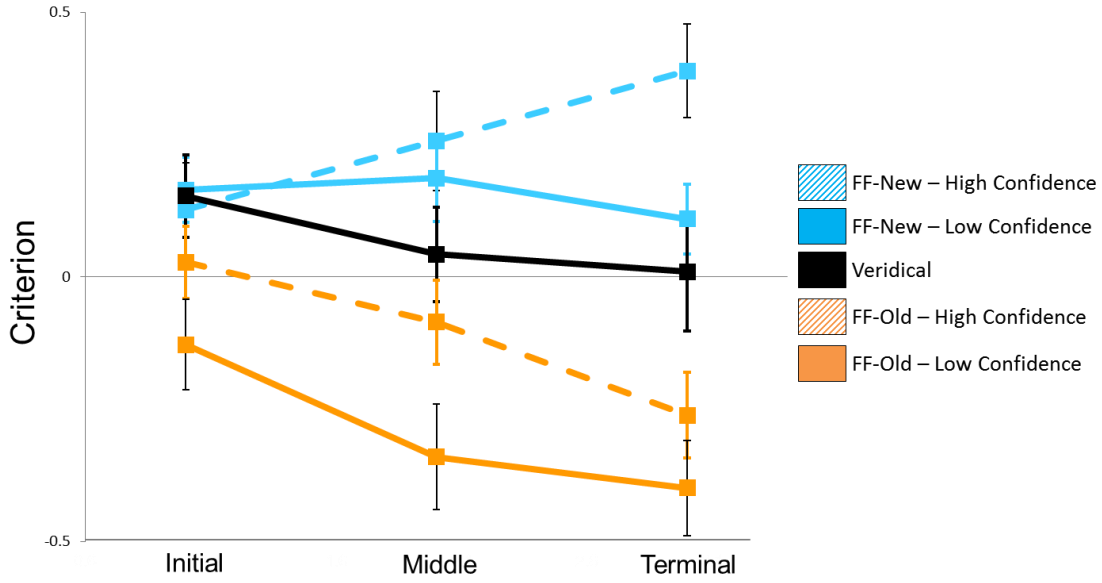

Figure 2.3. Criterion Shifts. Signal detection estimates of criterion across experimental groups over the initial, middle, and terminal thirds of the of the recognition test session. Error bars indicate standard error of the mean.

We note that we did not observe a reliable interaction between targeted confidence and time ($F(2,120) = 2.406$, $p = 0.095$). Because the experiment was double blind and employed random assignment, we have no reason to assume that there were group differences in default criterion before participants began the experiment. Thus, we interpret the main effect of targeted confidence as evidence that the false feedback manipulation caused shifts in criterion from individuals' pre-experimental default criterion. However, we cannot rule out the possibility that our effect is due to a failure of

random assignment (a Type I error with probability equal to our alpha level of 0.05). We did observe an interaction between targeted confidence and time for the FF-New group ($F(2, 60) = 5.020$, $p < 0.01$) but not the FF-Old group ($F < 1$). This null result may be due to rapid learning in the FF-Old-LC condition that resulted in a large criterion shift in the initial phase (**Figure 2.3**).

Given an incremental learning effect, we might expect the impact of targeted confidence to be strongest at the end of the experiment. That is, the effects of small versus large PEs will be largest after participants have had the most opportunities to learn from feedback. The pattern of terminal criteria confirmed this expectation: there was a main effect of targeted confidence ($F(1,60) = 6.612$, $p = .014$, $\eta_p^2 = .099$) and no group by confidence interaction ($F(1,60) = 0.901$, $p = .346$, $\eta_p^2 = .015$). The main effect of targeted confidence indicates that PEs support dynamic shifts of criterion, and that PEs are scaled relative to retrieved memory strength.

Finally, recall that the two alternatives differ primarily in their opposite predictions about the relative positions of the two FF-New groups: the Memory Strength alternative predicts that the FF-New-HC will show a more conservative criterion, while the Expected Response Outcome alternative predicts that the FF-New-LC will show a more conservative criterion. A *t*-test between the terminal criterion for these two groups confirms that the FF-New-HC group adopted a more conservative criterion ($t(30) = 2.622$, $p = .014$, Cohen's $d = 0.937$), supporting the Memory Strength alternative.

Confidence analyses

To fully characterize the pattern of confidence responses across our various manipulations, we analyzed confidence ratings as a function of response type, feedback group, targeted confidence, and time. We restricted our analysis to the four groups which received false positive feedback and so did not include Veridical. Confidence ratings were submitted to a five-way ANOVA with factors of recognition response (Old / New), accuracy (Correct / Incorrect), feedback group (FF-New / FF-Old), targeted confidence (Low Confidence / High Confidence), and time (Initial / Middle / Terminal).

A test of the Expected Response Outcome alternative requires that individuals have above-chance metacognitive accuracy; that is, correct responses are made with higher confidence than incorrect responses. This was confirmed by a main effect of accuracy ($F(1,60) = 309.343$, $p < .001$). **Table 2.2** shows the means and standard error for confidence ratings of each group as a function of correct and incorrect responses. Likewise, "old" responses were made more confidently than "new" responses (main effect of Response, $F(1,60) = 35.958$, $p < 0.001$). Both of these effects are consistent with previous research.

	Response Type	
	Correct	**Incorrect**
Veridical	0.50 (0.06)	0.39 (0.05)
FF-New - High Confidence	0.62 (0.05)	0.45 (0.04)
FF-New - Low Confidence	0.54 (0.03)	0.38 (0.03)
FF-Old - High Confidence	0.64 (0.04)	0.49 (0.04)
FF-Old - Low Confidence	0.46 (0.03)	0.35 (0.03)

Table 2.2. Mean Confidence Ratings for Correct and Incorrect Responses for Each Group.
Values in parentheses represent standard error of the mean.

There was also an accuracy by response interaction ($F(1,60) = 97.877$, $p < 0.001$). The difference between correct and incorrect responses was larger for "old" responses than for "new" responses. This effect was driven by the fact that correct "old" responses were made with higher confidence than correct "new" responses. This may reflect the qualitative difference between "old" responses (in which there is the presence of positive evidence for a previous encounter) and "new" responses (in which these is an absence of positive evidence for a previous encounter). It may also reflect contributions of recollection-based processes to "old" responses.

Participants in the Low Confidence groups made more low confidence responses than participants in the High Confidence groups (main effect of targeted confidence, $F(1,60) = 12.163$, $p < .001$), regardless of response type. Note that this main effect cannot explain the criterion shift results, because the FF-Old – Low Confidence and FF-New – High Confidence groups showed the largest shifts in criterion. Interestingly, this effect was not specific to the type of response that received false positive feedback (targeted confidence by feedback group interaction: $F(1,60) = 1.657$, $p = .204$); this effect was seen across both "old" and "new" responses for both the FF-New and FF-Old groups.

Finally, there was a main effect of time such that confidence ratings became lower over the course of the experiment ($F(2,120) = 22.049$, $p < 0.001$); this may reflect participants' tendency to reduce their confidence ratings as they garnered more experience with the recognition task and received feedback on their performance. No

other main effects of interactions reach significance after correction for multiple comparisons.

Individual differences analysis: correlating net prediction errors with criterion

To further assess the role of PEs in dynamic shifts of criterion, we computed the net PE that each participant encountered across all trials of the recognition test, and correlated this net PE with each participant's criterion. The criterion analyses reported above demonstrate a causal effect of targeting false positive feedback to low or high confidence errors. However, individuals do not only experience PEs following false positive feedback; instead, our framework suggests that each feedback instance provides a PE signal. Thus, this analysis allows us to leverage the Net PE signal that an individual participant encounters across all trials, including both veridical and false feedback trials.

Using the same logic described in the introduction and in **Figure 2.1**, we derived an Expected Value for each trial and calculated the corresponding PE for that trial. Again, a key assumption of our framework is that the Expected Value of a recognition decision can be derived from the participant's responses and confidence ratings. A reliable relationship between PEs calculated in this way and an individual's criterion would provide evidence consistent with PE-based learning of recognition criterion. We considered both alternatives for deriving an EV and computing the PE: one in which PE is scaled relative to Memory Strength, and another in which PE is scaled relative to the Expected Response Outcome. Consistent with the decision-making and RL literature, we derived EV on a scale of 0 to 1 for all decisions. Also consistent with the existing

literature, we assigned positive feedback an outcome value of 1 (both veridical and false) and negative feedback a value of 0.

To derive PEs scaled according to Memory Strength, we first computed an EV for each trial on a scale where the highest confidence "new" responses (*i.e.*, low memory strength) corresponded to 0 and the highest confidence "old" responses (*i.e.*, high memory strength) corresponded to 1 (e.g. black bars in **Figure 2.1D**). We then computed PE by taking the difference between the actual outcome and this derived EV (e.g. arrows in **Figure 2.1D** show PEs for positive feedback). To match the sign of criterion estimates (more liberal criterion is more negative), the sign of PEs was adjusted such that all PEs favoring "old" responses (positive outcomes after an "old" response; negative outcomes after a "new" response) were given a negative sign, while all PEs favoring "new" responses (positive outcomes after a "new" response; negative outcomes after an "old" response) were given a positive sign. These PEs were averaged across all trials (including trials that received false feedback and trials that received veridical feedback) to assign each participant a Memory Strength Net PE (MSNPE) score.

To derive PEs scaled relative to Expected Response Outcome, we first computed an EV for each trial on a scale where the lowest confidence responses (*i.e.*, low expected outcome) corresponded to 0 and the highest confidence responses (*i.e.*, high expected outcome) corresponded to 1, regardless of whether or not it was an "old" or "new" response (e.g. black bars in **Figure 2.1B**). As above, the PE was computed by taking the difference between the actual outcome and this EV (e.g. arrows in **Figure 2.1B** show PEs for positive feedback), and the sign of the PE was adjusted so that PEs favoring "old" responses were given a negative sign and PEs favoring "new" responses were given a

positive sign. These PEs were averaged across all trials to assign each participant an Expected Response Outcome Net PE (ERONPE).

We correlated these Net PE scores with each individual's criterion. If criterion is regulated through PE-based learning, more negative Net PE scores should be associated with more liberal criteria and more positive Net PE scores should be associated with more conservative criteria. Note that we do not use the change in criterion (the difference between terminal and initial criterion) as our dependent measure because the initial criterion is not an appropriate baseline. Because false feedback is provided immediately from the onset of the test phase, even the initial criterion estimate reflects some degree of criterion shift away from an individual's "default" criterion. In fact, in many decision-making and reinforcement learning tasks, learning rate is often highest during the initial portion of a learning experience.

We found a reliable correlation between Net PE scores and criterion, which is consistent with hypothesis that PE-based learning is engaged to regulate recognition decisions. Both the MSNPE ($R = 0.52$, $p < .001$) and the RONPE ($R = 0.51$, $p < .001$) were significantly and positively correlated with criterion and did not reliably differ in their R^2. **Figure 2.4** shows this correlation.

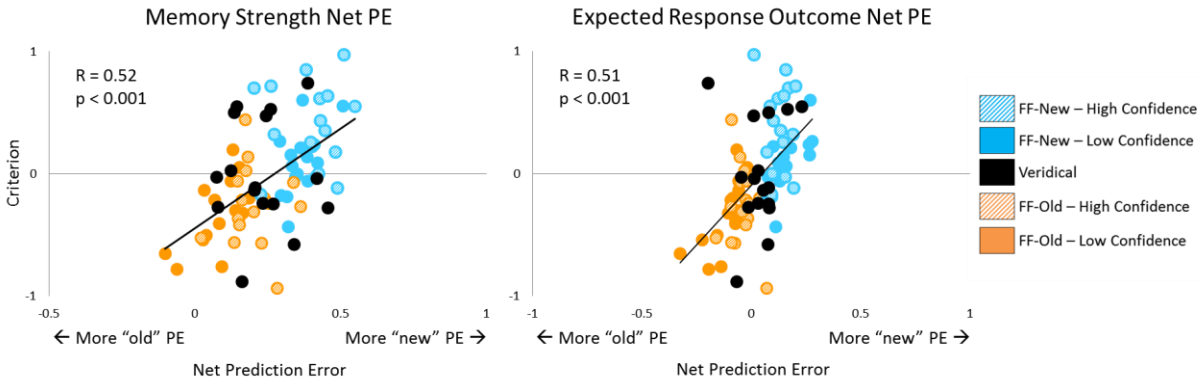

Figure 2.4. Net PE Correlational Analysis. Scatterplot and best fit line showing the correlation between Memory Strength Net PE and criterion. Experimental group is denoted by color.

Importantly, this correlation could be driven by the main effect of feedback group. That is, it may simply be the case that the FF-New group received more positive feedback on "new" responses and adopted a conservative criterion (and vice versa for the FF-Old group) and this simple group effect is driving the correlation between Net PE and criterion (*cf*. Simpson, 1951). In order to demonstrate an effect of Memory Strength PEs over and above the effect of group, we performed a multiple regression analyses that included a categorical predictor for Group (FF-Old, FF-New, Veridical) and a predictor for MSNPE. In this model, the beta weights for the PE predictors represent partial regression coefficients, controlling for the effect of group. This model provided a reliable fit to the data ($R = 0.552$; $p < 0.001$) and both the Group and MSNPE were significant predictors of criterion (both p's < 0.05). This indicates that MSNPE predicts criterion even after controlling for the effect of Group, consistent with a role for PE-based learning that occurs for both veridical and false feedback trials.

Finally, we performed a multiple regression analysis that included predictors for Group, MSNPE, and ERONPE. This model tests for three partial regression coefficients:

(1) the effect of Group, after controlling for MSNPE and ERONPE; (2) the effect of MSNPE, after controlling for Group and ERONPE; and (3) the effect of ERONPE, after controlling for Group and MSNPE. Including all of these in the same model allows us to test whether the MSNPE or ERONPE provides a better fit to the data. In this model, however, we find that none of these predictors are significant (Group: $p = 0.47$; MSNPE: $p = 0.06$; ERONPE: $p = 0.16$). That is, none of the individual predictors provides a reliable increase in predictive power after controlling for the other two predictors in the model. Therefore, this approach does not allow us to differentiate between the MS and ERO alternatives. However, the reliable multiple correlation coefficient for the model ($R = 0.568$; $p < 0.001$) indicates the model, as a whole, is a strong predictor of criterion.

We note that we replicated this general pattern of Net PE results across an independent data set collected using the false feedback recognition paradigm, as described Section II of in the *Supplemental Material*.

Discussion

We used an experimental manipulation of false feedback targeted to high or low confidence responses to show that PEs cause shifts in the recognition criterion. Moreover, the results indicated that PEs are scaled as a function of Memory Strength rather than Expected Response Outcome. Specifically, we found that the criterion shift was more liberal for the FF-Old-LC relative to the FF-Old-HC group and was more conservative for the FF-New-HC group than for the FF-New-LC group. This pattern supports the Memory Strength alternative, particularly the difference for the New groups, which are opposite of the prediction from the Expected Response Outcome alternative. These

results suggest that the EV used to compute PEs is linked to the underlying retrieved memory strength associated with the recognition decision.

We obtained evidence of change in the criterion over the course of the experimental session. As the participants encountered additional feedback, they enacted further adjustments to their recognition decision criterion. This is consistent with an incremental reinforcement learning process and with previous observations (Han & Dobbins, 2009). An important caveat is that this incremental effect was observed only for the difference between groups (FF-Old versus FF-New) but not for targeted confidence. The main effect of targeted confidence supports a PE-based learning account, but this effect did not increase over time. Therefore, we cannot rule out the chance of a false positive (a Type I error) whereby the main effect of targeted confidence is due to a failure of random assignment.

In addition to criterion shifts, we found that confidence responses themselves changed over the course of the session, with participants being more likely to make the reinforced confidence decision. This pattern may reflect some response-level learning occurring in parallel with learning of memory criterion. Specifically, based on the false feedback, participants may learn that low (or high) confidence responses are more likely to receive a positive outcome. This learning is not specific to either "old" or "new" responses, and instead the participant simply makes more low (or high) confidence responses, in general. Although our data suggest that PEs scaled relative to Memory Strength are more important for driving shifts in criterion, they do not rule out a separate or complementary role for PEs scaled relative to the Expected Response Outcome.

It is important to note that this change in confidence responses does not confound the observed change in criterion: we observed a larger shift in criterion in one Low Confidence group (FF-Old-LC) and one High Confidence group (FF-New-HC). In the same way, our observed pattern of criterion shifts is not easily explained by differences in salience or frequency of the targeted confidence levels. Because we used a continuous measure of confidence (as opposed to e.g. separate buttons for Low or High Confidence) we did not use a categorical cut-off to separate "low confidence" from "high confidence". Since we were not constrained to a few categorical responses, the types of trials targeted in the FF-New-HC condition did not differ in rarity from the types of trials targeted in the FF-New-LC (and likewise for the FF-Old groups).

Finally, we also obtained evidence that the Net PE encountered across all trials (including both veridical and false feedback) is predictive of an individual's criterion. This individual differences effect was observed even after for controlling for the main effect of group. However, this approach was not able to show that the Memory Strength alternative provides a better fit to the data than the Expected Response Outcome alternative. This may be due to the fact that the PEs predicted by the MS and ERO alternative are correlated (particularly for "old" responses) and this observation provides a cautionary note regarding using a subject-level individual differences analysis to test PE-driven learning in these contexts. However, the results of our causal manipulation (in particular, the pattern seen in the FF-New groups) suggests that PEs scaled relative to Memory Strength are more important in driving shifts in criterion.

Following from this discussion, an important open question for the criterion shift concerns the locus of learning. Concretely, within the signal detection theory frame,

criterion is computed on the proportion of "old" and "new" responses. Hence, it is not possible to distinguish whether the shift in criterion reflects learning how to categorize evidence from memory or rather reflects a change in response tendencies (i.e. a drive to respond "old" or "new" regardless of the evidence from memory). The concluding analyses use drift diffusion modeling based on accuracy and response time to address this final question.

Drift Diffusion Modeling: Identifying the Locus of Learning

A fundamental open question concerns the psychological locus of learning that underlies shifts in decision criterion during recognition memory tasks. As criterion within the SDT framework is estimated based solely on the responses of participants (their proportion of hits and false alarms), we cannot assess from this whether the shift in criterion is due to a change in the way that evidence from memory is treated when categorizing an item as old or new or simply reflects a change in the tendency to report old or new. More specifically, according to the response-level learning account, criterion shifts may occur simply because the participant learns that a specific response (e.g. "old" or "new") is more likely to result in a positive outcome. In this case, criterion shifts are simply a specific case of straightforward feedback-based response learning. As such, an apparent shift in criterion could occur even in the absence of a change in the evaluation of evidence from memory (i.e., consider that the participant simply presses "old" more often because this is rewarding, regardless of what they remember). By contrast, according to the memory-level learning account, a criterion shift occurs if the participant learns to better evaluate ambiguous retrieval outcomes with regard to the category of a cue (old or new). For example, individuals may shift their criterion as they learn to associate a

specific level of memory strength with old items. In this case, criterion shifts reflect a true dynamic shift in the underlying memory decision processes.

Both accounts are plausible given the evidence of criterion shifts based only on rates of hits and false alarms. Indeed, at least one previous study that used explicit pay-offs to reward "old" or "new" responses, interpreted their effects in terms of response level learning rather than as a change in the evaluation of evidence from memory (Curran et al., 2007). The present design holds a similar ambiguity. Thus, in order to characterize criterion placement during recognition memory as an RL process, it is crucial to determine whether the evidence to date is actually evidence of memory training or is simply standard response or motor level reinforcement in the context of a recognition decision. Though estimates of criterion based only on responses are of limited help in adjudicating among these alternatives, a change in response bias versus a change in evidence evaluation can be identifiable when one considers the response time (RT) distributions associated with error and correct responses under each model.

The widely applied drift diffusion model (DDM) is one decision making framework that incorporates both RT and accuracy in order to identify response bias from the criterion placed on evidence evaluation (e.g. Ratcliff & Starns, 2009; Ratcliff, 1978; Starns, Ratcliff, & McKoon, 2012). Thus, we complemented SDT modeling with DDM to better distinguish the locus of learning in the present work. In the DDM, the decision process arises from noisy evidence accumulation (or drift) towards one of two response bounds (**Figure 2.5**). For recognition decisions, one bound corresponds to an "old" response and the other bound corresponds to a "new" response. Thus, over time, evidence accumulation drifts towards one of these response bounds at a rate (i.e., the drift rate; v)

that is determined by the strength of evidence recovered from memory. Evidence that is evaluated as signaling "old" leads to a positive drift rate, towards the "old" boundary. Evidence that is evaluated as signaling "new" leads to a negative drift rate, towards the "new" boundary. The accumulation process is terminated when the drift reaches either boundary, and the corresponding response is made. Stronger evidence would lead to a larger drift rate, and thus, a faster response.

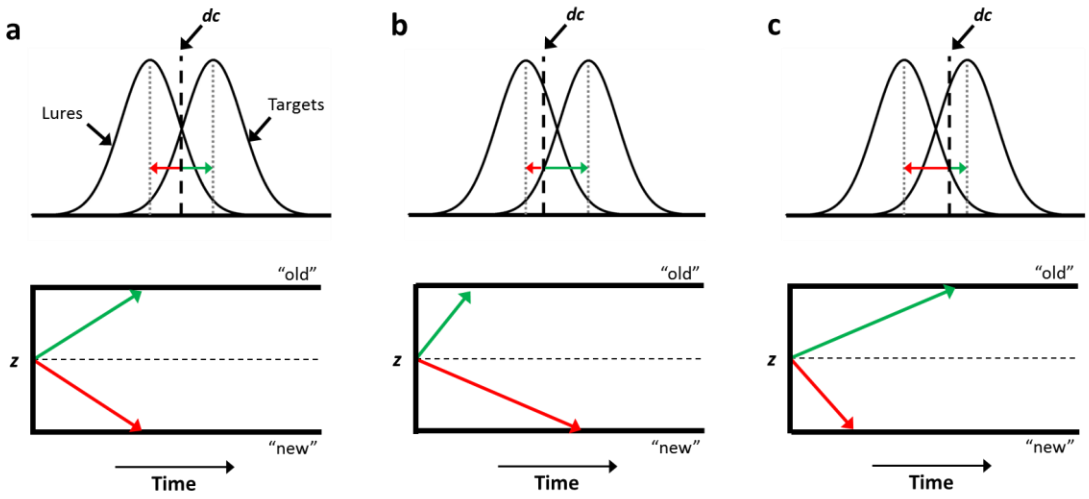

Figure 2.5. Schematic of the effects of shifts in the drift criterion on drift rates in the DDM. Plots in top row illustrate neutral (a), liberal (b), or conservative (c) placement of a drift criterion that is used for micro SDT processes at each time step of a drift process and that increment drift up or down. Plots in bottom row illustrate the impact of this drift criterion placement on the relative positive (toward "old") and negative (toward "new") drift. The response bias (z) is the starting point between the two response bounds and is depicted as a dashed line. Response bias is in the neutral position in all plots, equally between old and new. Shifting start point positively or negatively would also increase the bias to make old or new responses respectively, but would do so independent of the level of evidence on any trial.

Within this framework, the response-level versus memory-level learning alternatives could be formalized as a change in either *response bias* or *drift criterion*, respectively (Criss, 2010; Ratcliff & McKoon, 2008; Starns, Ratcliff, & McKoon, 2012;

Voss, Rothermund, & Brandtstädter, 2008). The response bias, *z*, determines the starting point of the diffusion process relative to either of the response boundaries. When bias is closer to the "old" response boundary, participants are overall more likely (requiring less evidence) to respond "old" independently of the strength of the evidence (drift rate). Shifts in response bias result in a larger effect on the leading edge of the RT distribution relative to shifts in the drift criterion (Ratcliff & McKoon, 2008). Response bias shifts are most closely analogous to the response-level learning described above: a response bias shift means that one of the responses is more likely to occur, in general, but this bias is not influenced by the actual evidence accumulation or dynamic decision process.

The drift criterion parameter, *dc*, in the DDM determines whether observation of a signal of a given strength is more likely to elicit drift towards the "old" or "new" boundary. This is not thought to change how a memory strength or match signal is retrieved from long term memory (which is impacted by factors such as strength of studied items) but instead how this match signal is evaluated and mapped onto drift rate (Ratcliff & McKoon, 2008; Ratcliff & Starns, 2009; Starns, Ratcliff, & McKoon, 2012). From this perspective, the drift process can be thought of as a series of small signal detection processes, where at each time step, the evidence for an item is sampled. If that sample is greater than the old/new criterion then the drift is incremented toward "old," and is otherwise incremented toward "new." The criterion of those signal detection processes conducted at each time step is the drift criterion.

At the aggregate decision level, shifting the drift criterion in a liberal direction will increase the average positive (i.e., toward "old") drift rate along with a corresponding decrease in the average negative drift rate (**Figure 2.5B**). Conversely, a

conservative shift in drift criterion will result in a decrease in the average positive drift rate and a corresponding increase in the average negative drift rate (**Figure 2.5C**). This change in drift criterion is most closely analogous to the memory-level learning described above: a criterion shift means that a certain level of evidence or memory strength is now more likely to be classified as "old" as opposed to "new." A change in the RT distribution based on either a change in bias or drift criterion can be determined by the effects on the tail and leading edge of the RT distribution: changes in drift criterion have a larger impact on the tail than do changes in bias (Ratcliff & McKoon, 2008).

A change in either of these parameters in the model would lead to an essentially identical change in response proportions as measured via the standard SDT estimates, but they are identifiable within the DDM. Notably, the few studies investigating these two accounts of criterion during recognition memory have found that manipulating base rates and target strength resulted primarily in changes in response bias parameters, with subtle effects on drift criterion as well (Criss, 2010; Starns, Ratcliff, & McKoon, 2012).

To summarize, we applied the drift diffusion model to address the question: when an individual shifts her/his criterion in the context of false feedback, what is the psychological locus of this adjustment? We consider two alternatives: (1) criterion shifts are the result of response-level; (2) criterion shifts are the result of changes in the evaluation of retrieved memory evidence as either "old" or "new." Thus, we contrasted various drift diffusion models to see which parameter was better able to capture the changes across our empirical data sets.

Method

Model Parameterization and Estimation Procedures

The recognition decision was modeled using the DDM, following prior work (e.g. Ratcliff, Thapar, & McKoon, 2004). For all model variants, parameters were fit to the accuracy and RT data (from both correct and error trials). We estimated bias (*z*) and drift criterion (*dc*); along with parameters for an initial non-decision interval (*t*), thought to reflect processes like encoding of the target; and the distance between the boundaries (*a*) (see **Figure 2.5**). Inter-trial variability parameters were not included in the model. Model fitting was conducted using the Hierarchical Drift Diffusion Modeling (HDDM) module (Wiecki, Sofer, & Frank, 2013). This module, implemented in Python, uses a hierarchical Bayesian estimation procedure that fits the DDM parameters based on all participant data simultaneously. This approach is analogous to random effects estimation in that it treats between subject variance as a random variable while fitting within-subject parameters simultaneously, and so has advantages over other approaches like pooling all participant data or fitting each individual participant separately. The HDDM software uses informative priors to constrain parameter estimates within a range of plausible values (Wiecki et al., 2013).

We independently fit the data from the experiment reported above, as well as from two independent data sets described in the Section II of the *Supplemental Material*. For each data set, we focused on two potential DDM models. In particular, we fit a model, termed the "drift criterion model," in which drift criterion (*dc*) was allowed to vary across false feedback group (e.g., Veridical, FF-Old, FF-New) and time (First Half

versus Second Half of the experimental session); and a model, termed the "response bias model," in which response bias (*z*) was allowed to vary across false feedback group and time. No other parameters were allowed to vary as a function of feedback group or time. We also fit a third model that allowed both *drift criterion* and *response bias* to vary as a function of feedback group and time. The deviance information criterion (DIC) and posterior parameter estimates from these models are reported in the results. Simulated RT distributions were produced using the best fitting model parameters as a posterior predictive check, as implemented in the HDDM software. Quantile plots of these simulated distributions are shown in **Supplemental Figure 2.7** (see *Supplemental* Material) and the simulated data were considered credible if the observed data was captured in the 95% credible interval of simulated data.

Response time data were examined for outliers: any response faster than 300ms was considered a guess or an error and excluded from the analyses (Criss, 2010; Ratcliff et al., 2004; Starns, Ratcliff, & White, 2012). Non-response errors (when participants failed the response deadline) were excluded from analysis. These comprised fewer than 3% of trials on average across participants.

A Markov-chain Monte-Carlo (MCMC) procedure estimated the DDM parameters' posterior distributions. 10,000 samples from the distributions were estimated. The first 3,000 samples were discarded (burn in), and of the remaining samples, every tenth sample was retained (thinning). Model convergence was assessed based on Monte-Carlo error (MC error) and visual assessments of chain convergence (Gelman, Carlin, Stern, & Rubin, 2004). Model selection was based on minimization of the deviance

information criterion (DIC), which is more readily compatible with MCMC estimation than Akaike information criterion or Bayesian information criterion.

Results

Both the drift criterion and response bias models provided good fits of the data. In each model, MC error was less than 1% of the standard deviation of the posterior distribution for all parameters, providing evidence of convergence. Critically, however, behavior was better fit by the drift criterion model than the response bias model, as indicated by lower DIC scores. **Table 2.3** shows DIC scores for each model. **Figure 2.6** depicts the plots of the parameter estimates for each experiment and shows that the posterior plots from the drift criterion models are qualitatively more similar to the pattern seen in the signal detection theory analyses of criterion than are posterior plots from the response bias models. Furthermore, we replicated this pattern of results across two independent data sets, as described in Section II of the *Supplemental Material*.

Finally, a third model that allowed both drift criterion and response bias to vary provided a better fit than either of the restricted models **(Table 2.3).** One caveat to this result is that this third model contains an extra parameter than the restricted models. The DIC procedure is designed to account for this difference in parameters; however, there remains the chance that it does not adequately penalize for the additional parameter in this model. This is not an issue for the comparison between the two restricted models since they each contain the same number of parameters.

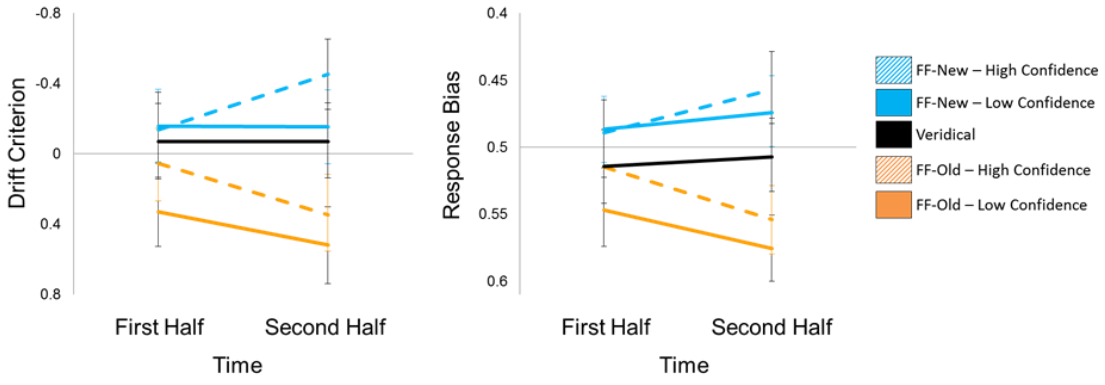

Figure 2.6. DDM Parameter Estimates. Interaction plots show the DDM estimates of the drift criterion (left column) and the response bias (right column) for each experimental group across the first and second half of the recognition test session. On each graph, more conservative values are plotted up and more liberal values are plotted down. Drift criterion and response bias were estimated in separate models. Model comparison indicated that a DDM model in which drift criterion shifted over time fit the data better than a model in which response bias shifted over time. Error bars represent 95% confidence intervals.

Model	DIC	deviance	pD
Drift Criterion	22186	21820	366
Response Bias	23042	22692	350
Full Model	21935	21469	465

Table 2.3. Comparison among alternative drift diffusion models. Lower DIC values indicate better fitting model. DIC: deviance information criterion.

Discussion

The data from three independent experiments were consistently better described by the drift criterion model than the response bias model. These outcomes suggest that the pattern of response time and accuracy is more consistent with a change in the way that evidence is evaluated and categorized as "old" or "new," and not an evidence-

independent bias to report "old" or "new." As such, criterion shifts are not simply a special case of conventional response-based RL in which participants learn that a given action (an "old" or "new" response) is more associated with a positive outcome. Instead, participants adjust the process by which they evaluate retrieved evidence. That is, they learn to adjust the 'cognitive action' that maps a given level of memory evidence to either an "old" or "new" decision.

Together with the SDT results linking the Expected Value of decisions to the underlying retrieved memory strength, these results suggest that feedback-based shifts in recognition criterion involve latent memory states and memory evaluation processes distinct from typical perceptual decision making and RL paradigms. However, we note that a third model that included both shifts in drift criterion and in response bias provided the best fit to the data, notwithstanding potential caveats regarding model complexity and DIC scores. Therefore, it is possible that both shifts in drift criterion and in response bias contribute to feedback-based criterion shifts.

The DDM is one of many competing models of recognition memory decisions (Clark & Gronlund, 1996; Dennis & Humphreys, 2001; Glanzer et al., 2009; Ratcliff & Starns, 2009; Ratcliff, 1978; Shiffrin & Steyvers, 1997; Van Zandt, 2000; Wixted & Gaitan, 2002). Importantly, the role of PEs in driving criterion shifts is not dependent on the assumptions of the DDM; we chose to apply the DDM to our data because this framework has components that allow us to clearly differentiate between the response-level and memory-level alternative explanations of criterion shifts, which are indistinguishable in the conventional SDT framework. This is a novel application of the

drift diffusion model, as to our knowledge the DDM has not previously been used to investigate feedback-based criterion shifts in recognition memory.

General Discussion

We tested the hypothesis that PE signals drive dynamic control of recognition decisions. In an old/new recognition memory task, we manipulated the magnitude of PEs by providing false positive feedback following either low confidence or high confidence errors. This approach yielded two key results. First, the magnitude of PEs causally impacted the magnitude of dynamic criterion shifts such that larger PEs led to greater shifts in criterion. Second, this pattern of results was best explained by a model wherein PE is scaled relative to the underlying retrieved memory strength associated with a recognition decision. This suggests that the Expected Value of a recognition memory decision, which is in turn used to compute PEs following outcomes, is linked to the retrieved memory strength. Finally, converging results from drift diffusion modeling show that these criterion shifts involve a change in the way that evidence retrieved from memory is evaluated, more so than a simple change in response bias. Together, these results provide empirical evidence consistent with recent theories positing an important role of PEs and reinforcement learning circuits in the cognitive control of declarative memory retrieval (Han & Dobbins, 2009; Scimeca & Badre, 2012; Wixted & Gaitan, 2002).

Cognitive control strategies are critical to the efficient use of declarative memory. Previous work has suggested that declarative memory control strategies are learned and

updated in a manner akin to skill-learning (Benjamin, 2007; Scimeca & Badre, 2012). Although we hypothesize that this is broadly applicable across declarative memory, here we used recognition memory as an experimental task to demonstrate one of the primary predictions of this hypothesis: the magnitude of PEs should impact adjustments in memory control strategies. Earlier work showing that feedback could influence shifts in criterion was consistent with this hypothesis, but did not directly manipulate the magnitude of PEs (Han & Dobbins, 2008, 2009). By targeting false feedback to either high or low confidence responses, we were able to show that the degree to which this feedback deviates from a participant's expectations impacts the magnitude of shifts, directly implicating PEs in the control of recognition decisions.

An important assumption underlying our experimental logic is that self-reported confidence is in some way linked to a participant's expectation of a successful outcome. We considered two alternatives regarding how confidence may be linked to this Expected Value. Under the Expected Response Outcome alternative, this EV is positively correlated with the level of reported confidence for both "old" and "new" responses: responses made with low confidence are associated with a low EV; responses made with high confidence are associated with a high EV. Under the Memory Strength alternative, this EV is linked to the underlying retrieved memory strength associated with any test item: items that elicit low memory strength are associated with a low EV; items that elicit high memory strength are associated with a high EV. In this case, reported confidence is positively correlated with EV for "old" responses and negatively correlated with EV for "new" responses. These two alternatives are primarily distinguished by their predictions for PEs following "new" responses, and our data support the Memory Strength

alternative: targeting false positive feedback to high confidence "new" responses resulted in a larger criterion shift than did targeting false positive feedback to low confidence "new" responses.

Interestingly, this indicates that false positive feedback made following a high confidence "new" response results in a large PE and large increase in the propensity to respond "new" in the future. Conversely, false positive feedback following a high confidence "old" response results in a small PE, and only a small increase in the propensity to respond "old" in the future. This idea may seem inconsistent with previous work in the recognition memory literature: making recognition decisions about items that elicit memory signals near the decision criterion is more difficult and more prone to erroneous decisions, while making recognition decisions about items eliciting memory signals further from the decision criterion is easier and more accurate (Stretch & Wixted, 1998a).

From this perspective, it may seem counterintuitive that participants use successful memory retrieval as a predictor of generally positive outcomes. However, in many real-life situations and indeed in many laboratory tasks, retrieving sufficient signal from long-term memory is critical to successful performance. For example, in a source memory task, retrieval of the specific source detail is critical to success (e.g. Johnson et al., 1993; Mitchell & Johnson, 2009). Old/new item recognition is a particular case in which "absence of evidence" can be used as "evidence of absence." That is, a retrieval attempt that elicits a very low level of memory evidence results in an easy, high-confidence "new" response in a recognition task. But in any situation in which an

individual is searching memory for specific information, a retrieval attempt that elicits a very low level of memory evidence is unrewarding.

This pattern of Memory Strength EVs may be task-dependent. Because the participant is in a "memory test", they may assign higher value to states with high memory evidence, since high memory evidence is more likely to be associated with the explicit (or assumed) goals a memory task (i.e. find "old" items). Alternatively, the pattern may be task-independent. Individuals may have learned, in a Pavlovian manner, that successful retrieval generally followed by positive outcomes. Intriguingly, this positive association with retrieval evidence may relate to mere exposure effects, whereby simple, implicit exposure to a stimulus results in increased preference for that stimulus (Whittlesea & Price, 2001; Zajonc, 1968, 2001).

A final possibility is that there are separate "oldness" and "novelty" inputs into the memory-decision system (Kafkas & Montaldi, 2014, 2015; Rutishauser et al., 2015). The oldness signal increases from left-to-right across the horizontal evidence axis, and the EV signal may be derived from the "oldness" signal. Similar models exist in the context of perceptual decisions (e.g. Law & Gold, 2009). The formation and generality of this hypothesized "oldness bonus" remains an open question and should be tested in future work (see *Chapter 4*).

Learning to evaluate declarative memories through a nondeclarative RL system

The present study is part of a growing literature on the interaction between memory systems that are often dichotomized as "implicit" and "explicit" (Scimeca & Badre, 2012; Shohamy & Adcock, 2010). The pattern of results we obtain highlights why

it is critical to test the predictions of nondeclarative PE-based learning in the context of declarative memory. We argue that the reinforcement learning mechanisms widely studied in decision making and working memory perform similar functions in the service of the cognitive control of declarative memory. However, declarative memory differs in qualitative ways from the conventional settings where RL is studied. For example, if the learning that occurs in our paradigm were simply a special case of RL, in which participants learned that a certain type of response was simply more predictive of a positive outcome in the context of this task, then we would expect to find a symmetrical pattern of results between "old" and "new" responses (consistent with the Expected Response Outcome alternative.) The asymmetry that we find in our results is not easily explained by this sort of simple RL process. Instead, our results indicate that EVs are related directly to the success or failure of the memory process.

This interpretation is further supported by the results from the DDM. Recognition memory decisions involve evaluating latent memory signals. We found that changes in the evaluation of memory evidence (the *drift criterion* parameter) better described our data than did changes in the propensity to make a specific response (the *response bias* parameter). As such, the learning we observe is not simply straightforward stimulus-response learning but instead points to a deeper interaction between the reinforcement-learning and declarative memory systems, whereby feedback drives changes in evaluating latent memory evidence. This suggests that these changes should broadly impact any memory task that taps into this evaluation process. However, the generality of this effect remains an open question.

The notion that criterion shifts involve changes in evidence evaluation is also relevant to recent debates about the differentiation model of criterion shifts in recognition memory (Criss, 2010; Starns, Ratcliff, & White, 2012). In short, this model posits that strength-based criterion effects are the result of differentiation between the target and lure distributions of memory strength, and that better differentiation shifts the lure distribution to the left on a strength axis of a signal detection theory model. Although our results do not preclude a role of differentiation in other instances of criterion shifts, it is difficult to explain the present results via a differentiation account. Because all groups encountered identical items during the study phase (as well as during the test phase), there should be no difference in the strength distributions between groups. Instead, we posit that our false feedback manipulation resulted in dynamic adjustments in evaluative processes used to assess memory signals retrieved from existing strength distributions.

A neural substrate for learning from PEs in declarative memory

We interpret the present results as evidence for a role of PEs and reinforcement learning processes in the control of recognition memory decisions. This makes an intriguing link between learning within the context of memory decisions and well-established circuits that compute Prediction Error in the brain. Various brain regions have been implicated in computing PEs, and recent theories have posited a potential role for striatal PEs, in particular, in the cognitive control of declarative memory (Scimeca & Badre, 2012). Within the domains of probabilistic learning and decision making, there is considerable evidence linking neural PEs in striatum to nondeclarative learning (Liljeholm & O'Doherty, 2012; O'Doherty et al., 2004, 2003; Schultz et al., 1997).

However, only recently has there been growing interest in a potential role for PE-based learning in declarative memory retrieval.

Though direct examples of the types of PE studied here have yet to be tested in the brain, there exists fMRI evidence tying striatal activity to feedback processing in declarative memory more generally. Tricomi and Fiez (2008) used a paired-associate learning task in which participants saw pairs three times and received feedback after each trial. On the first encounter with a cue, participants could only guess the correct target; on subsequent encounters, participants' decisions could be informed by their memory for the correct association. Neuroimaging revealed striatal activity to feedback on the memory trials, but not the initial learning trial, suggesting that striatal activity was only present when the feedback was informative to memory performance. O'Connor and colleagues (2010) showed that striatal activity indexes deviations of memory expectations by combining a recognition paradigm with cues regarding the likely status (old or new) of upcoming test probes. A contrast between invalid and valid cue trials included activity in striatum, consistent with the idea that this region is sensitive to deviations from expectations in a declarative memory context.

As already noted, the criterion results further suggest that retrieved memory strength is linked to an Expected Value used to compute PEs; that is, retrieved memory strength is predictive of positive outcomes in the future and should be therefore be valued as rewarding. This interpretation is consistent with neuroimaging data on the "retrieval success" effect – contrasts between hits and correct rejections consistently show activation of striatum (Spaniol et al., 2009; Wagner et al., 2005), and this striatal activity is evident even when comparing false alarms versus correct rejections (Abe et al., 2008).

This activity may be tracking the expectation that successful retrieval is predictive of a positive outcome. What constitutes a positive outcome can also depend on the relevant task goals. When hits are incentivized with monetary rewards, this striatal effect is greatly enhanced (Han et al., 2010). By contrast, when correct rejections are incentivized, the striatal response to hits persists at baseline levels even as the striatal response to correct rejections is substantially enhanced. Thus, in a typical recognition task, it appears that hits may be interpreted as rewarding both due to goal-attainment and due to pre-existing associations with successful outcomes. The present results suggest that future work using human neuroimaging techniques (see *Chapter 3*) and/or studying patient populations with pathology of the nigra-striatal dopamine system will lead to further insights into the basic mechanisms by which decision making in recognition memory is adapted to changes in task and circumstance.

Accounting for prior results within an RL framework

The PE-based learning described here provides a concrete mechanism for dynamic control of recognition memory decision criterion. There are likely many processes that can impact criterion, including both declarative and nondeclarative influences, and multiple processes probably contribute depending on task structure and task goals. For example, experimenter instructions and overt incentive manipulations likely engage more explicit top-down control to set an appropriate decision criterion. PE-based learning instead represents a nondeclarative process, and ostensibly acts dynamically over the course of the recognition test.

The RL model of controlling dynamic shifts in recognition criterion provides a framework for understanding the positive and negative results in the extant literature. One pattern of apparently conflicting results is the interaction between strength manipulations and external feedback. Two studies that manipulated both of these factors provide an instructive example. As discussed earlier, Verde and Rotello (2007) reported no shifts in criterion between test blocks containing strong and weak target items in the absence of external feedback, even when the blocks were marked with semantic cues. However, they found that providing external feedback did result in a criterion shift between strong and weak blocks. In contrast, Hicks and Starns (2014) reported no impact of external feedback on criterion shifts between strong and weak memory strength test blocks.

One potential explanation for this discrepancy is the amount of feedback that was provided in each study. The experiments conducted by Hicks and Starns (2014) used block sizes of 10-40 items and only provided explicit feedback following incorrect responses (although positive feedback was implied by the absence of negative feedback following correct responses). Verde and Rotello (2007) used a block size of 80 items and provided explicit feedback on every trial. This extended block length and consistent feedback may have provided enough time for the gradual accumulation of feedback to drive criterion shifts; whereas the block lengths and intermittent feedback used by Hicks and Starns may have been insufficient for accumulated feedback to impact criterion.

This is consistent with the incremental nature of RL and false-feedback based criterion shifts, as demonstrated here and previously (Han & Dobbins, 2008; 2009). Together with the positive results of combining feedback and base-rate differences (e.g.

Kantner & Lindsay, 2010), this suggests that veridical feedback can be a useful source of information when there are inherent strength-based or base-rate differences in a recognition test, and when there is sufficient time to accumulate evidence from feedback.

However, many recognition tests do not manipulate strength distributions or base-rate differences and report no impact of veridical feedback (Kantner & Lindsay, 2010; Starns, White, & Ratcliff, 2010). Why does veridical feedback not impact criterion under these conditions? The RL framework provides a potential explanation for this pattern of results. Learning in RL is largest when outcomes significantly deviate from expectations, resulting in large PEs and large learning signals. If feedback does not meaningfully deviate from expectation, learning is unlikely to occur. As an example, the average criterion reported by Kantner and Lindsay (2010) was approximately 0.08 for participants receiving feedback and 0.07 for participants not receiving feedback (an optimal, neutral criterion is 0.00). Similarly, participants in the Veridical conditions of the present study adopted a slightly conservative criterion. This may be because the learning signals provided by veridical feedback point, on balance, to a neutral or slightly conservative criterion that matches the participant's default criterion. This speculation is supported by the individual analyses correlating net PE with criterion. In these analyses, the accumulated PE encountered by participants in the Veridical condition is consistent with a relatively neutral criterion (see **Figure 2.4**).

Thus, this RL framework proposes two conditions that promote dynamic criterion shifts. First, criterion shifts are more likely to occur when there is enough time to accumulate sufficient evidence in favor of a criterion shift. Second, criterion shifts are more likely to occur when outcomes (e.g. external feedback) sufficiently deviate from

underlying expectations and the existing criterion. In the present study we used false positive feedback to bias the net PE experienced by individuals, but this framework predicts that any manipulation of outcomes (such as selectively omitting positive feedback; e.g. Han & Dobbins, 2008) that sufficiently biases the net PE should show the same pattern of shifts.

Relationship between the RL framework and existing models of recognition memory

Our framework proposes that individuals are learning to behave according to decision policies for a given level of retrieved memory strength evidence. In this sense, the absolute level of retrieved memory strength serves as a context for the decision. Thus, the framework aims to describe recognition memory decisions in the same way as conventional value-based or perceptual decision making models. In these models, individuals must learn the appropriate decision policies for a given context, wherein the context is typically a specific perceptual stimulus, such as a particular shape or symbol. However, in recognition memory decisions the relevant context is not a discrete stimulus, but instead, it is a latent state drawn from a continuous abstract dimension: the level of memory evidence (memory strength) evoked by a stimulus. In a typical recognition test, test items are only presented once and multiple test items will result in similar levels of retrieved memory evidence. Therefore, the individual must learn the values and actions that are appropriate for specific levels of memory strength (Scimeca & Badre, 2012). For example, an individual might learn that the appropriate policy for a high level of memory strength is to act as if the item is old (e.g. make an "old" decision; Cox & Dobbins, 2011) because this policy is more likely to result in a positive outcome given this level of memory strength. Likewise, through experience, an individual might learn that in the

context of low memory strength, acting as if the item is old is unlikely to result in a positive outcome.

In this "Reinforcement Ratio" framing, when an individual retrieves a given level of evidence, the individual acts based on her/his previous reinforcement history under similar levels of evidence in the past (Wixted & Gaitan, 2002). This framework is analogous to likelihood-ratio models of recognition memory in which a memory decision is based on the relative probability that a level of memory evidence would be produced by an old versus a new item (Cox & Dobbins, 2011; Glanzer et al., 2009). In a basic likelihood ratio model, an individual considers an item as "old" whenever this ratio is greater than 1, which occurs above the intersection of the new and old distributions. This pattern of responding is akin to using a neutral criterion in a signal-detection framework. Similarly, our Reinforcement Ratio model proposes that individuals learn the reinforcement contingencies that occur for a given level of evidence. Thus, as an example, when biased feedback results in a more liberal criterion, an individual has learned that there are more levels of memory evidence over which she should act as if the item is old. Previous theorists have proposed similar models of recognition memory in (G. S. Brown & White, 2009) and in humans (Cox & Dobbins, 2011; Koop et al., 2014; Mickes et al., 2011; Ratcliff & Starns, 2013; Wixted & Gaitan, 2002) whereby an individual learns evidence-to-decision policies through a history of reinforcement.

The Reinforcement Ratio framework predicts that individuals will learn a response profile similar to that predicted by a likelihood ratio model. Indeed, Wixted and Gaitan (2002) have proposed that reinforcement histories provide the underlying learning mechanism by which humans learn to perform like likelihood ratio computers. This is

illustrated in the following example. Consider the level of evidence just above the intersection of the new and old distributions. Veridical feedback following correct "new" and correct "old" decisions made for this level of evidence will result in the same magnitude positive PE (because PE is scaled relative to memory strength) but "old" decisions are more likely to result in positive feedback (because old items are more likely at this level of evidence). Therefore, based on this pattern of reinforcement, an individual will learn that the appropriate decision policy given this level of evidence is to act as if the item is old. This will also hold for all levels of memory evidence above the intersection of the two distributions because p(old|strength) > p(new|strength) above the intersection. Conversely, "new" responses will be more frequently reinforced for all levels of memory evidence below the intersection of the two distributions because p(new|strength) > p(old|strength) below the intersection.

This learned response profile can explain the seemingly-puzzling results observed when individuals are given a recognition memory tests that consists either entirely of targets (a target-pure test) or entirely of unstudied lures (a target-free test). When participants complete these tests without any feedback, the criterion for target-pure tests and target-free tests are strikingly similar (Cox & Dobbins, 2011; Koop et al., 2014). Although this result is challenging for strength-based models and formal likelihood-ratio models, it is easily explained by a Reinforcement Ratio model. Responses on both target-pure and target-free lists are essentially driven by habits, or the learned decision policies for different levels of memory evidence (Cox & Dobbins, 2011). When participants are provided feedback on target-pure or target-free lists, they are able to use this feedback to adjust their responding to the drastic differences in base rates (Koop et al., 2014). The

idea that recognition decisions are shaped by learned, habitual decision policies is also consistent with work showing that an individual's criterion shows trait-like stability across time, test material, and test format (Kantner & Lindsay, 2012).

The Reinforcement Ratio framework also accommodates the Memory Strength alternative that is supported by the present results. Under the MS alternative, correct "new" responses are more likely to receive large positive PEs compared to "old" responses, because new items are generally associated with lower memory strength. Thus, one might predict that individuals will become increasingly conservative over the course of a recognition test with veridical feedback. However, just as this deviation does not occur in empirical data, this sort of deviation does not occur in the Reinforcement Ratio framework. Rather, the distribution of value over memory strength will primarily affect the rate at which the appropriate decision policy for the continuum of states along the memory strength dimension is learned. As an example, consider again the level of evidence just above the intersection of the new and old memory strength distributions. Veridical feedback following correct "new" and correct "old" decisions made for this level of evidence will result in the same magnitude positive PE (because PE is scaled relative to memory strength) but "old" decisions are more likely to result in positive feedback (because old items are more likely at this level of evidence). Therefore, based on this pattern of reinforcement, an individual will learn that the appropriate decision policy given this level of evidence is to act as if the item is old. This will also hold for all levels of memory evidence above the intersection of the two distributions because p(old|strength) > p(new|strength) above the intersection. Because the converse is true

below the intersection, the individual will learn that the appropriate decision policy for these low levels of evidence is to act as if the item is new.

However, the Memory Strength alternative is not easily accommodated by global-matching or strength-based models in which an individual maintains and updates a single criterion value as a reference and compares retrieved memory evidence to this reference. In these models, adjusting the criterion based on PEs scaled relative to a memory strength value function would result in systematic deviations such that an individual making recognition decisions with veridical feedback would become increasingly more conservative. This would occur because a basic version of this model would use the PEs experienced across all levels of memory evidence to update a single threshold value. That is, in a recognition task with veridical feedback, the average positive PE experienced for correct "new" decisions is larger than the average positive PE experienced for correct "old" decisions, averaging across all levels of memory evidence. If these PEs were combined to update a single criterion threshold, the larger positive PEs for "new" decisions would result in an ever-increasing conservative threshold.

Conclusion

In conclusion, we have demonstrated that targeting false feedback to either high or low confidence memory responses impacts the magnitude of dynamic shifts in recognition decision criterion. This is novel evidence for a role of prediction errors and reinforcement learning in the control of recognition memory decisions and provides direct empirical support for the final hypothesis laid out in Chapter 1. Reinforcement learning theory has proven a fruitful framework for understanding probabilistic learning,

decision making, and action selection, and the present work extends this framework to the domain of recognition memory decisions. This not only provides a new domain in which to test predictions of reinforcement learning, but provides a concrete mechanism for how individuals evaluate and adjust control processes involved in declarative memory decisions.

Supplemental Material

Section I: False feedback manipulation checks in main experiment

The feedback manipulation in the main experiment was successful at targeting false feedback to high or low confidence responses. The count of false feedback trials was submitted to a feedback group (FF-Old / FF-New) by targeted confidence (Low Confidence / High Confidence) ANOVA. The average targeted confidence for Low Confidence groups was significantly lower than for the High Confidence groups (Means = 0.26, 0.62; $F(1,60) = 95.860$, $p < 0.001$). The experimental script was designed to provide 30 false feedback trials to each participant; the mean number of false feedback trials was 33.1. This deviation was likely to do participants proclivity to make increased errors (due to the false feedback manipulation) over the course of the experiment, such that the experimental script initially underestimated participants' error rates and provided additional false feedback trials.

There was a significant difference in the number of false feedback trials between the Low and High Confidence groups ($F(1,60) = 6.927$, $p < 0.05$) such that participants in the Low Confidence groups were provided more false feedback than the High Confidence groups (Means: FF-Old – Low Confidence: 37.4; FF-Old – High Confidence: 29.4; FF-New – Low Confidence: 34.1; FF-New – High Confidence: 31.6). Note that this difference alone cannot explain the pattern we report in the decision criterion, as the larger criterion shifts were shown by the FF-Old – Low Confidence group and the FF-New – High Confidence group. Furthermore, looking specifically at the two FF-New groups, which showed the greatest difference in the magnitude of criterion, there was not

a significant difference in the number of false feedback trials provided ($t(1,30) = 0.0941$, $p = 0.354$).

Section II. Supplemental data: overview and experimental methods

Here we report data from two pilot studies completed using the false feedback recognition paradigm. The results from these experiments provide important insight into the efficacy of the false feedback manipulation and provide independent replications of many of the important results reported in the main chapter.

Method

Methods were identical to the main experiment except where noted below.

Participants. Fifty-three participants (34 females; age, 18-22 years; mean, 19.5 years) participated in Experiment A and fifty-three participants (36 females; age, 18-23 year; mean, 19.0 years) participated in Experiment B in exchange for course credit or payment. All participants had normal or corrected-to-normal vision and were native English speakers. Participants gave written informed consent according to guidelines established and approved by the Human Research Protections Office of Brown University. After completing the experimental session, participants were thoroughly debriefed regarding the false feedback manipulation. Participants were assigned randomly to experimental group.

Procedure. The stimulus set consisted of the same 772 words as the main experiment. Participants studied 672 words during the semantic judgment task, divided into six blocks. Participants proceeded to the recognition test immediately following the study phase. The recognition test phase consisted of 100 target items and 100 lure items in random order, divided into four blocks of 50 trials each. Because the recognition test was

appended to an ongoing study of semantic decisions (which we do not discuss here), the list of 100 lure items was fixed across participants.

False Feedback Manipulation. A Veridical Feedback (VF) group provided a baseline control for the experimental groups in both Experiments A and B. The false feedback in Experiment A was manipulated between two groups: a FF-Old and FF-New group. The FF-Old group received false positive feedback on 70% of their "old" errors, independent of their response confidence, and veridical feedback on all other trials. The FF-New group received false positive feedback on 70% of their "new" errors, independent of their response confidence, and veridical feedback on all other trials.

In Experiment B, the validity of feedback was again manipulated between a FF-Old and a FF-New group. Additionally, we initially aimed to target false positive feedback to either low or high confidence responses. For half of participants in each of the FF-Old and FF-New groups, false feedback was more likely to be provided to errors made with confidence ratings above the midpoint of the confidence scale: 100% of targeted errors (either false alarms or misses) made above the midpoint of the scale received false positive feedback, 40% of targeted errors made below the midpoint of the scale received false positive feedback, and all other trials received veridical feedback. For the other half of participants in each group, false feedback was more likely to be provided to errors made with confidence ratings below the midpoint (100% for targeted errors below the midpoint, 40% for targeted errors above the midpoint).

However, this experimental manipulation was unsuccessful: the mean confidence of false feedback trials did not differ significantly between these high and low confidence subgroups. In a 2 [Group: FF-Old / FF-New) x 2 [Targeted Confidence: Low / High]

ANOVA on the average confidence of trials that received false feedback, there was no main effect of condition ($F(1,49) = 0.773$, $p = 0.384$) and no significant main effect of targeted confidence ($F(1,49) = 3.145$, $p = 0.082$). Similarly, post-hoc tests revealed that the manipulation was not significant at either level of condition: a *t*-test between FF-Old-HC and FF-Old-LC did not pass conventional levels of significance ($p = 0.051$); a t-test between FF-New-HC and FF-New-LC was not significant ($p = 0.590$). Given this failure of the targeted confidence manipulation, we collapsed across this manipulation for all further analyses.

Inspection of the data suggested that the false feedback manipulation failed because of individual differences in the distribution of responses over the range of the confidence scale. This observation motivated the adaptive algorithm used in the main experiment, which was designed to tailor false feedback to each individual's response tendencies, therefore maximizing the difference in targeted confidence across participants while providing an approximately equal number of false feedback trials to each individual. The manipulation check analyses reported in Section I of the *Supplemental Material* indicate that the manipulation used in the main experiment was successful.

Results

SDT and confidence analyses. Criterion, d-prime, hit rate, and false alarm rate as a function of group are provide in **Supplemental Table 2.4**. We replicated the main effect of false feedback on criterion in both Experiment A ($F(2,50) = 13.415$, $p < 0.001$, $\eta_p^2 = 0.349$) and in Experiment B ($F(2,74) = 6.00$, $p < 0.005$, $\eta_p^2 = 0.140$). Confidence ratings as a function of accuracy are shown in **Supplemental Table 2.5**. We replicated the finding that correct responses are made with higher confidence than incorrect responses

in both ($F(1,105) = 667.511$, $p < 0.001$, $\eta_p^2 = 0.864$). Due to space limitations, we do not consider any other statistical tests.

Net PE and criterion correlational analyses. Using the procedure outlined in the main chapter, we computed the Net PE that each participant encountered across all trials of the recognition test. We computed separate Net PE scores under the assumptions of the Memory Strength alternative (an MSNPE score) and under the Expected Response Outcome alternative (an ERONPE score). We collapsed across all participants from Experiments A and B for the correlational analysis. Both the MSNPE ($R = 0.44$, $p < 0.001$) and ERONPE ($R = 0.46$, $p < 0.001$) were significantly and positive correlated with terminal criterion. As in the main experiment, one source of this correlation could simply be a main effect difference between groups. To account for this effect, we performed a multiple regression analysis that included three predictors: Group (FF-Old, FF-New, VF), Memory Strength Net PE, and Expected Response Outcome Net PE. In this regression model, both the Memory Strength Net PE and Expected Response Outcome Net PE remained significant predictors of criterion (both p's < 0.05).

Drift Diffusion Model analyses. Using the same methods described in the main chapter, we replicated the drift diffusion modeling analyses on the supplemental data from Experiments A and B. We focused on the comparison between the *drift criterion* model and the *response bias* model. We fit each model separately for Experiment A and Experiment B. The DIC scores for each model are shown in **Supplemental Table 2.6**. For both Experiment A and B, the *drift criterion* model provided a better fit to the data (indicated by lower DIC score), replicating the result from the main experiment in independent data sets.

Measure	Group	Time	
		First Half	Second Half
Criterion	Veridical Feedback	0.05 (0.04)	0.10 (0.04)
	False Feedback - New (Exp A)	0.35 (0.08)	0.24 (0.05)
	False Feedback - New (Exp B)	0.17 (0.06)	0.21 (0.05)
	False Feedback - Old (Exp A)	-0.04 (0.07)	-0.15 (0.09)
	False Feedback - Old (Exp B)	0.00 (0.07)	-0.09 (0.06)
D-Prime	Veridical Feedback	1.48 (0.10)	1.67 (0.09)
	False Feedback - New (Exp A)	1.19 (0.24)	1.07 (0.23)
	False Feedback - New (Exp B)	1.34 (0.12)	1.36 (0.12)
	False Feedback - Old (Exp A)	1.24 (0.15)	1.13 (0.14)
	False Feedback - Old (Exp B)	1.43 (0.08)	1.41 (0.08)
Hit Rate	Veridical Feedback	0.75 (0.02)	0.76 (0.02)
	False Feedback - New (Exp A)	0.59 (0.04)	0.61 (0.04)
	False Feedback - New (Exp B)	0.67 (0.03)	0.67 (0.02)
	False Feedback - Old (Exp A)	0.73 (0.03)	0.75 (0.03)
	False Feedback - Old (Exp B)	0.75 (0.02)	0.77 (0.02)
FA Rate	Veridical Feedback	0.23 (0.02)	0.18 (0.01)
	False Feedback - New (Exp A)	0.20 (0.04)	0.23 (0.05)
	False Feedback - New (Exp B)	0.22 (0.02)	0.21 (0.02)
	False Feedback - Old (Exp A)	0.29 (0.03)	0.35 (0.04)
	False Feedback - Old (Exp B)	0.26 (0.03)	0.28 (0.02)

Supplemental Table 2.4. Criterion, D-Prime, Hit Rate, and False Alarm Rates for Supplemental Data.
Values in parentheses represent standard error of the mean. FA: false alarms.

Group	Response Type	
	Correct	Incorrect
Veridical Feedback	0.70 (0.03)	0.47 (0.03)
False Feedback - New (Exp A)	0.69 (0.03)	0.50 (0.05)
False Feedback - New (Exp B)	0.68 (0.03)	0.46 (0.02)
False Feedback - Old (Exp A)	0.70 (0.03)	0.53 (0.03)
False Feedback - Old (Exp B)	0.67 (0.03)	0.47 (0.03)

Supplemental Table 2.5. Confidence Ratings for Correct and Incorrect Responses.
Values in parentheses represent standard error of the mean.

	Model	DIC	deviance	pD
Experiment 1A	**Drift Criterion**	12642	12420	222
	Response Bias	12700	12488	213
Experiment 1B	**Drift Criterion**	18050	17725	325
	Response Bias	18117	17799	318

Supplemental Table 2.6. DDM model comparison in Experiments A and B. Lower DIC values indicate better fitting model. In each experiment, the drift criterion model better fit the data. DIC: deviance information criterion.

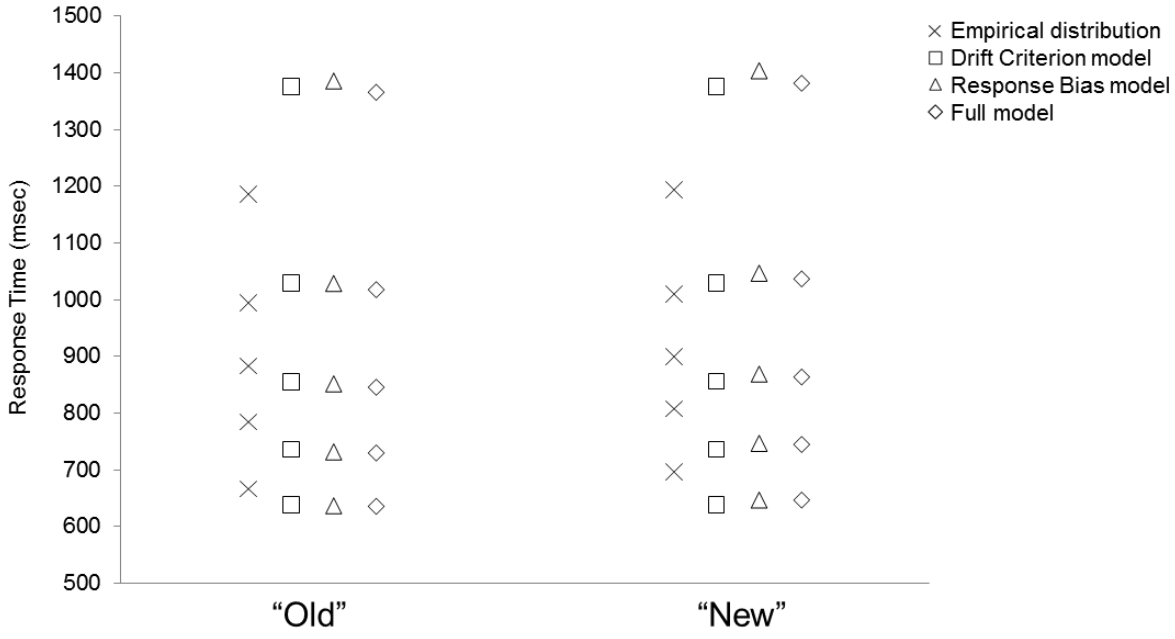

Supplemental Figure 2.7. Quantile plots of observed and simulated RT distributions. Five RT quantiles (.1, .3, .5, .7, .9) from the empirical RT distributions from the main experiment (X's) are plotted for "old" and "new" responses. Simulated RTs from the drift criterion DDM model (squares), response bias DDM model (triangles), and full model (diamonds) are shown for comparison. For all quantiles for all models, the respective empirical quantile fell within the credible 95% interval of the simulated data.

CHAPTER 3

Striatal Prediction Errors Support Dynamic Control of Declarative Memory Decisions[‡]

Adaptive memory requires control over how retrieved evidence is interpreted given the context. Yet the neural substrate that supports adjustments to memory decisions remains unknown. In Chapter 2, I showed that adjustments to recognition memory decisions follow the behavioral hallmarks of PE-based learning, and that these PEs are scaled relative to a value signal derived from the underlying memory strength associated with a decision. In this chapter, I use functional neuroimaging to show that predictions errors coded by the striatum update control over a memory decision.

Human participants completed a recognition memory test that incorporated biased feedback to drive adjustments of the recognition criterion. Using model-based fMRI, I found that PEs – the deviation between the outcome and expected value of a memory decision – correlated with striatal activation and predicted individuals' final criterion. Importantly, unlike conventional reinforcement learning, striatal PEs were scaled relative to memory strength rather than the expected trial outcome. Modeling further indicated that learning adjusted the evaluation of mnemonic evidence rather than response bias, consistent with the results of Chapter 2. Finally, I observed that the learned criterion transferred to a free recall task. Thus, I provide convergent evidence that declarative memory decisions can be adjusted via striatal learning signals.

[‡] This chapter has been adapted from a manuscript submitted for publication:
Scimeca, J. M., Katzman, P. L., & Badre, D. (under review). Striatal prediction errors support dynamic control of declarative memory decisions.

Introduction

The human brain is capable of efficiently retrieving and acting on information from memory that has high utility given the current circumstances. Intrinsic features of our memory systems related to basic encoding, retrieval, and consolidation partly confront this "information retrieval problem" (J. R. Anderson & Milson, 1989; Schacter & Addis, 2007). However, strategic cognitive control processes supported by frontoparietal brain networks further guide retrieval toward high utility outcomes given our current goals (Badre, Lebrecht, Pagliaccio, Long, & Scimeca, 2014; Badre & Wagner, 2007; Barredo, Oztekin, & Badre, 2013; Dobbins, Simons, & Schacter, 2004; Mitchell & Johnson, 2009; Wagner et al., 2005). It follows that memory disorders arising in patients with compromised cognitive control systems are often traceable to a failure of strategic guidance and use of memory (Aly et al., 2011; Schacter et al., 1996; Wheeler et al., 1995). Yet a striking observation has been that healthy individuals apply these memory retrieval strategies efficiently, balancing costs and benefits, and further, that they do so without explicit instruction (J. R. Anderson & Milson, 1989; Badre & Wagner, 2007; Benjamin, 2007; Moscovitch, 1992; Rugg & Wilding, 2000). Thus, a fundamental question concerns how the brain acquires and adapts memory control processes.

Outside of the memory domain, reinforcement learning (RL) mechanisms are thought to contribute to the acquisition and adaptation of policies for strategic action selection. In the brain, frontostriatal circuits are central to this process, with frontal cortex supporting maintenance of task relevant information and the nigra-striatal dopamine system mapping value to the selection of appropriate actions and policies. Ventral striatum, particularly the nucleus accumbens, has been widely associated with receipt of

primary and secondary rewards (Haber & Knutson, 2010; Knutson, Adams, et al., 2001), as well as reward prediction error (PE), which reflects the degree to which an experienced outcome value deviates from the expected value (EV). Across many experiments, these striatal value signals have been shown to drive learning of cognitive control for action selection and nondeclarative learning (Badre & Frank, 2012; Frank et al., 2004; Gläscher et al., 2010; Haber & Knutson, 2010; Maia & Frank, 2011; Schultz et al., 1997).

Thus, one hypothesis is that cognitive control of memory is similarly acquired through a reinforcement learning (i.e., nondeclarative learning) process, wherein positive PEs reinforce a reliance on a prevailing memory strategy and negative PEs punish it (Badre et al., 2014; Scimeca & Badre, 2012). Importantly, however, the application of RL mechanisms to the learning of cognitive control over memory is complicated by several unknowns regarding how value is assigned to memory decisions and outcomes (Elward et al., 2014; Han et al., 2010; Schwarze, Bingel, Badre, & Sommer, 2013; Scimeca & Badre, 2012).

We focus here on how human participants learn to relate evidence from memory to a decision and response; specifically, we address adjustments of the recognition memory criterion. Recognition memory involves classifying a stimulus as novel (new) or previously encountered (old). Recognition memory theory commonly assumes that all items will elicit at least some match to long-term memory and thus will provide some evidence of oldness, termed "memory strength" (Eichenbaum et al., 2007; Malmberg, 2008; Mitchell & Johnson, 2009; Norman & O'Reilly, 2003; Ratcliff & McKoon, 2008; Stretch & Wixted, 1998a). The memory strength distribution for old items (i.e., items

studied during learning) will be higher than the distribution for new items (i.e., unstudied items), although these distributions can overlap (**Figure 3.1a**). The recognition decision criterion specifies the threshold level of evidence above which an "old" decision is made (**Figure 3.1a**). Criterion setting is a form of mnemonic control in that it relates retrieved evidence to a task-dependent decision and response.

Having a dynamic recognition criterion provides the primary means by which an individual can balance errors due to misidentifying a new item as old (false alarms) or failing to recognize an old item (misses). A neutral criterion falls at the intersection of the studied and unstudied distributions (**Figure 3.1a)**. A conservative criterion requires more evidence for an "old" decision, while a liberal criterion requires less evidence for an "old" decision. The optimal criterion placement depends on an individual's context and goals (Rotello & Macmillan, 2007). For example, depending on whether you are at a conference or in a foreign country, you may be more or less likely to approach someone who is only vaguely familiar. Both the expectation that the person is actually someone you have met and the expected value of speaking to them or the costs of being wrong differ across these two scenarios. Thus, expectation and value arise at multiple stages in the retrieval and decision process.

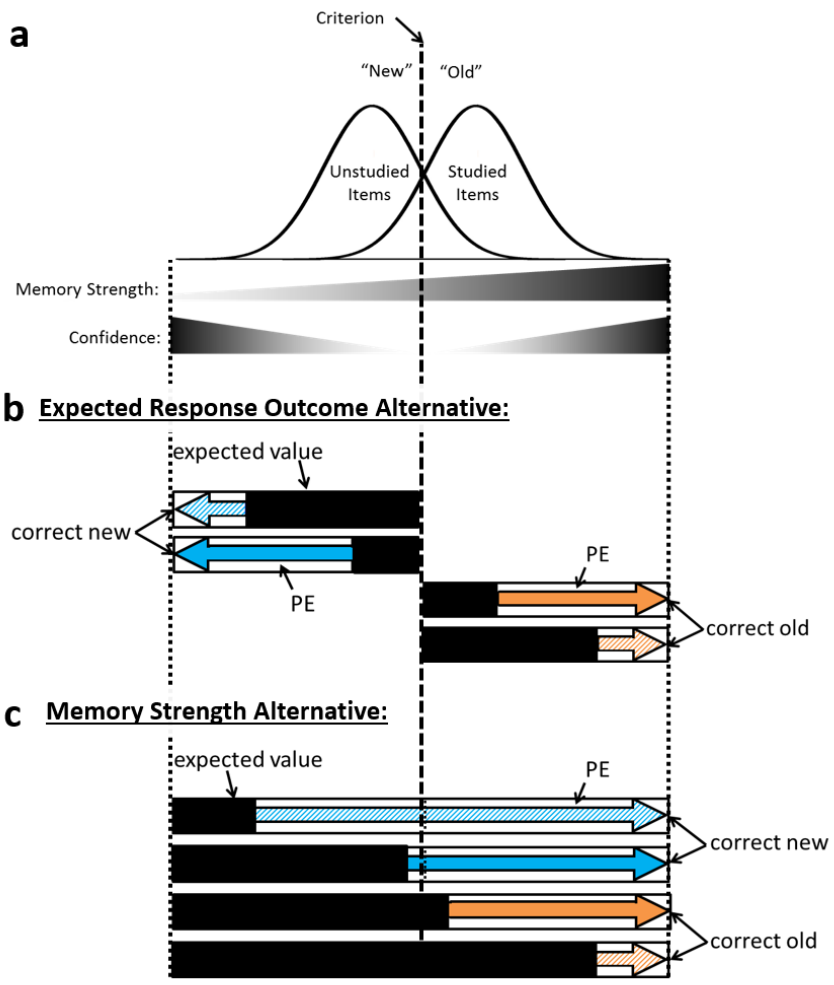

Figure 3.1. Two alternative models of computing PEs for memory decisions. (a) Standard signal detection model of recognition memory decisions. Memory strength increases along the x-axis from left to right; confidence increases for responses further from the criterion. The neutral criterion shown here falls at the intersection of the memory strength distributions for unstudied and studied items. The scale of responses depicted in panels b and c is directly mapped onto the scale depicted in panel a. (b) The expected values (EV) and prediction errors (PE) associated with positive feedback following four representative responses under the Expected Response Outcome alternative. (c) The EVs and PEs associated with positive feedback following four representative responses under the Memory Strength alternative. Blue arrows correspond to positive PEs following a "new" decision and thus reinforce a more conservative criterion, whereas orange arrows correspond to positive PEs following an "old" decision and thus reinforce a more liberal criterion.

One implicit manipulation that reliably produces dynamic adjustments in recognition criterion is biased feedback (Han & Dobbins, 2009). For example, providing false positive feedback following incorrect "old" responses drives individuals to adopt a more liberal criterion. We used this paradigm to induce adjustments to decision criterion and identify the underlying neural mechanisms. We hypothesized that feedback will elicit memory-related PE signals in striatum, and these PEs will drive adjustments in recognition criterion by reinforcing or punishing the preceding memory decision. However, a challenge to testing this hypothesis is that the source of EV associated with a memory decision is unknown. We thus considered two alternative models of how the EV and PE are determined for memory decisions.

The first alternative, which we term the Expected Response Outcome (ERO) alternative, holds that the EV of a memory decision is directly related to the expected outcome of the full trial following a response. As such, this EV can be estimated directly from the "confidence" in the memory decision. That is, a high confidence rating reflects a high expectation of getting the trial correct, and so receiving positive feedback. This relationship holds equivalently for both "new" and "old" responses. In signal processing terms, higher confidence ratings correspond to memory decisions for items further from the criterion (**Figure 3.1a;** Stretch & Wixted, 1998). **Figure 3.1b** shows four possible responses under this alternative. For example, positive feedback following a high confidence "new" response would result in a relatively small positive PE and a modest conservative shift in criterion (**Figure 3.1b**, top row).

The second alternative, which we term the Memory Strength (MS) alternative, holds that the EV of a memory decision is linked to the retrieved memory strength that

led to that decision, such that greater retrieved memory strength is linked to greater EV; i.e., "oldness" confers higher EV regardless of the response that follows. Under this alternative, the EV of a decision is related directly to memory strength, and so while EV will be positively correlated with confidence ratings for "old" decisions, it will be negatively correlated with confidence for "new" decisions (**Figure 3.1a**). **Figure 3.1c** shows four possible responses under this alternative. For example, positive feedback following an item with low memory strength (*i.e.* a high confidence "new" response) would result in a relatively large positive PE and thus a large conservative shift in criterion (**Figure 3.1c**, top row).

Both of these alternatives have plausible support from the existing literature. Individuals in recognition tasks typically demonstrate reliable metacognitive accuracy: high confidence responses are more likely to be followed by a positive outcome than are low confidence responses (Chua, Schacter, Rand-Giovannetti, & Sperling, 2006; Stretch & Wixted, 1998a). This is consistent with the ERO alternative and is largely analogous to the conceptualization of EV in conventional RL domains. The MS alternative is consistent with theories that argue that successful memory retrieval has utility for retrieval monitoring and for memory-guided action, and is therefore rewarding (Elward et al., 2014; Gallo, 2013; Han et al., 2010; Johnson et al., 1993; Mitchell & Johnson, 2009); as well as theories that posit that retrieval elicits approach/further processing (e.g., familiarity gated retrieval; Hayama & Rugg, 2009; Kahn, Davachi, & Wagner, 2004).

We used the logic described in **Figure 3.1** to test the hypothesis that striatal PEs drive adjustments in the recognition criterion. First, using trial-by-trial estimates of MS-PE and ERO-PEs, we test the hypothesis that striatum will track PE and these PEs will

drive shifts in memory criterion. Then, using drift diffusion (DDM) modeling and a free recall transfer protocol, we demonstrate that learning is localized to the evaluation of retrieved evidence, as opposed to a response bias.

Results

Biased feedback elicits dynamic adjustments in recognition memory decisions

Participants completed a recognition memory test during fMRI scanning (**Figure 3.2a**; see *Methods*). Overall, participants were able to accurately perform the task (mean proportion correct = 76.41%; SEM = 1.21%) and fewer than 5% of trials were non-response errors. The mean response time was 993 ms (SEM: 12 ms).

Following a jittered delay, feedback concerning the accuracy of the response was provided on each trial. False positive feedback was provided on approximately 70% of incorrect "old" responses (false alarms). All other responses received veridical feedback. This biased feedback drove participants to adopt a more liberal criterion over the course of the experiment (**Figure 3.2b**), consistent with previous behavioral work (Han & Dobbins, 2009; *Chapter 2*). The criterion shift effect was supported by a main effect of time on criterion ($F(5,90) = 2.938$, $p = .017$, $\eta_p^2 = .140$), as well as by planned *t*-tests comparing the terminal criterion against the initial criterion ($t(18) = 2.848$, $p = .011$, Cohen's $d = 0.692$) and against a neutral criterion of zero ($t(18) = 3.730$, $p = .002$, Cohen's $d = 0.859$).

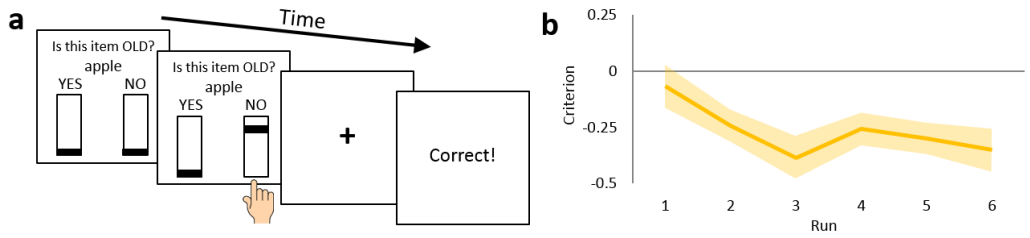

Figure 3.2. Behavioral task stimuli and results. (a) Participants indicated their old/new responses using a continuous confidence scale and received feedback following a jittered delay. (b) The biased feedback paradigm elicited a liberal shift in criterion (i.e., more negative criterion values) over successive runs of the recognition test. Error bars are SEM.

Striatum codes prediction errors for recognition memory decisions

To identify PE signals in the brain, we calculated trial-by-trial EVs and PEs to use in model-based fMRI analyses. We constructed a separate general linear model (GLM) for each of the ERO and MS alternatives using the respective EV and PE values (**Figure 3.1**; see *Methods*). This approach allows the regressors in each model to capture both the shared and unique variance for each alternative.

We first considered the neural circuitry that tracked trial-by-trial Memory Strength PEs (MS-PEs) at feedback. As predicted, the striatum, along with other brain areas, tracked MS-PEs following memory decisions (**Figure 3.3b**; **Supplemental Figure 3.6**; **Supplemental Table 3.1**). The signal across an anatomical region-of-interest (ROI) encompassing the whole striatum (see *Methods*) was significantly associated with MS-PE ($t(18) = 4.858$, $p < .001$, Cohen's $d = 1.114$). Moreover, this striatal signal predicted individuals' terminal criterion ($R = -0.547$, $p = .015$; **Figure 3.4a**), such that individuals with greater PE-related activation of striatum showed larger adjustments to their recognition criterion.

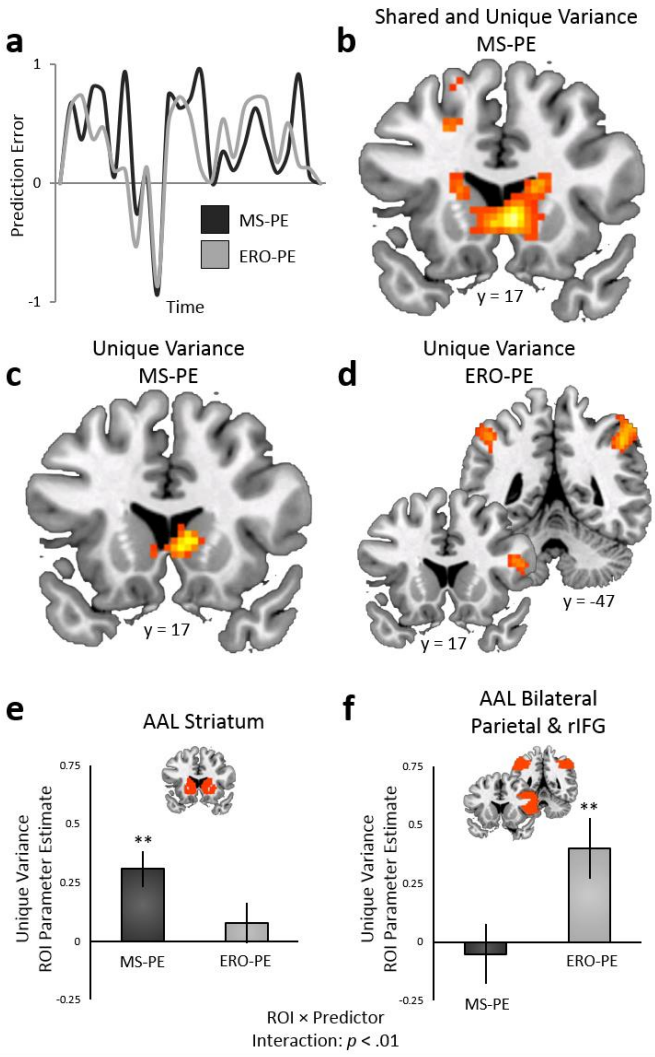

Figure 3.3. FMRI responses to trial-by-trial PE estimates. (a) Representative example of PEs predicted by the Memory Strength (MS-PE) and Expected Response Outcome (ERO-PE) alternatives. The two alternatives make correlated but dissociable predictions. (b) Striatum tracks the trial-by-trial PE associated with the shared and unique variance of the Memory Strength alternative. (c, d) When the predictors are restricted to unique variance, MS-PE is coded in striatum and ERO-PE is coded in right inferior frontal gyrus and bilateral inferior parietal cortex. (e, f) Signal extracted from anatomical ROIs (inset; see text for definitions) reveal a significant ROI × predictor interaction. Whole brain statistical maps are thresholded at $p < .05$, FDR cluster corrected. Color maps represent t statistics with a range from 0 to 8. Parameter estimates are in arbitrary units. Error bars are SEM. $^{**}p < .01$ versus zero.

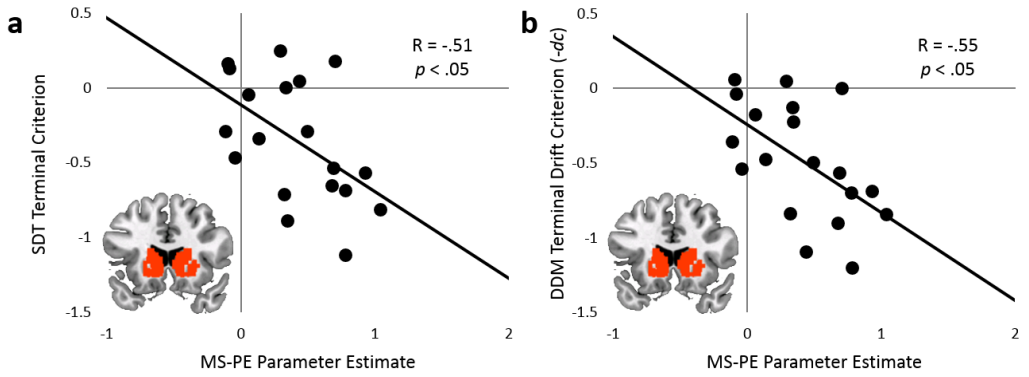

Figure 3.4. Correlation of striatal PE with terminal SDT criterion and DDM drift criterion. MS-PE parameter estimates from striatum predict individuals' terminal criterion estimated from signal-detection theory (SDT; a) and terminal drift criterion estimated from drift diffusion modeling (DDM; b). More liberal criterion values are plotted down in both panels. Parameter estimates are in arbitrary units and reflect the contribution of both shared and unique variance for MS-PE to BOLD signal in striatum.

We found a similar pattern of results for the neural correlates of Expected Response Outcome PEs (ERO-PEs). ERO-PE also activated striatum, among other brain regions (**Supplemental Figure 3.6**; **Supplemental Table 3.2**). However, as will be discussed below, this similar pattern is likely due to the large degree of shared variance between the MS-PE and ERO-PE predictors (**Figure 3.3a**).

Memory Strength Prediction Errors better explain striatal signal

We constructed a third GLM that included both MS-PE and ERO-PE regressors (see *Methods*). This approach allows the regressors to compete for variance and thus activity associated with either regressor in this model will be due to the unique variance above and beyond their shared variance. We found that the unique variance of MS-PE was tracked by striatum (**Figure 3.3c**; **Supplemental Figure 3.6**; **Supplemental Table 3.1**). The unique variance of ERO-PE was not significantly associated with any activity

in striatum. Instead, the unique variance of ERO-PE was associated with activation in bilateral inferior parietal cortex and right inferior frontal gyrus (**Figure 3.3d**; **Supplemental Figure 3.6**; **Supplemental Table 3.1**). A direct contrast of the unique variance of MS-PE > ERO-PE, though conservative, revealed a marginally significant cluster in right striatum (k = 4; $p < .001$, uncorrected; peak coordinates: x=9, y=17, z=4; p = .053 FWE SVC; see *Methods*). The reverse contrast (ERO-PE > MS-PE) did not reveal reliable activation in striatum or elsewhere.

To assess the pattern of unique variance across the striatum relative to other regions in the brain, we complemented this voxelwise analysis with ROI analyses and sought to test a predictor × region interaction (Henson, 2006). We first extracted the signal from the anatomical striatum ROI (**Figure 3.3e**). Activity across the entire striatal ROI was significantly associated with unique MS-PE ($t(18) = 4.172$, $p < .001$, Cohen's d = 0.957) but not with unique ERO-PE ($t(18) = 1.162$, $p = .260$, Cohen's d = 0.266), although the difference between the two predictors did not reach significance ($t(18)$ = 1.558, $p = .137$, Cohen's d = 0.357). Then, guided by the functional activations observed for unique ERO-PE, we used the AAL database to construct a second anatomically defined ROI by combining bilateral inferior parietal masks with a mask of right inferior frontal gyrus opercularis (**Figure 3.3f**). Activity in this anatomical ROI was significantly associated with unique ERO-PE ($t(18) = 3.355$, $p = .004$, Cohen's d = 0.770) but not with unique MS-PE ($t(18) = -0.224$, $p = .825$, Cohen's d = -0.051) and there was a reliable difference between ERO-PE and MS-PE ($t(18) = 2.114$, $p = .049$, Cohen's d = 0.827). We submitted the parameter estimates for unique MS-PE and unique ERO-PE from these two ROIs to a 2×2 ANOVA and found a significant predictor × ROI interaction ($F(1,18)$

= 13.224, p = .002, η_p^2 = .424). Taken together, these results show that the signal in striatum is better explained by the MS-PE predictor than by the ERO-PE predictor, relative to the opposite pattern observed in frontoparietal regions.

Expected Value is broadly coded by both retrieval and value networks

We found that the MS-EV regressor of shared and unique variance was coded in a broad network that included precuneus, medial PFC, caudate, posterior cingulate, and amygdala (**Supplemental Figure 3.7**; **Supplemental Table 3.2**). This network overlapped with the "core" recollection network (Elward et al., 2014; Rugg & Vilberg, 2013; Spaniol et al., 2009) and regions typically exhibiting "retrieval success" effects (Scimeca & Badre, 2012; Spaniol et al., 2009; Wagner et al., 2005). Many of these regions have also been implicated in representing subjective value during choice tasks (De Martino, Fleming, Garrett, & Dolan, 2013; Knutson et al., 2005; Roy, Shohamy, & Wager, 2012). The ERO-EV regressor of shared and unique variance was correlated with a similar network of regions (**Supplemental Figure 3.7**; **Supplemental Table 3.2**), overlapping with regions identified in previous studies of recognition confidence (Chua et al., 2006).

When we included MS-EV and ERO-EV in the same analysis (GLM 4; see *Methods;* **Supplemental Figure 3.7**; **Supplemental Table 3.2**), we found that the unique variance of these regressors was represented in overlapping areas of ventromedial PFC and precuneus (De Martino et al., 2013). ERO-EV, but not MS-EV, was represented in the left anterior frontal pole, left superior frontal gyrus, and left putamen. MS-EV, but not ERO-EV, was represented in a more posterior portion of vmPFC.

The psychological locus of learning is in memory evaluation

Participants made more "old" responses over the course of the recognition test, indicating a liberal shift in criterion. As such, there are (at least) two possible psychological loci where learning could have occurred. One explanation is that the adjustment simply involved an increased propensity to make the "old" response because "old" responses more frequently led to positive feedback outcomes. Thus, individuals learned to bias their responding towards "old" responses regardless of the retrieved memory evidence, akin to the learning observed in typical RL and decision-making paradigms. We term this "response-level learning". An alternative explanation is that the more liberal criterion reflects adjustments to the evaluation of retrieved memory evidence. That is, individuals learned to adjust the threshold above which they considered a given level of memory evidence as "old". We term this "memory-level learning". Thus, though the imaging data suggest that memory-level PE is the learning signal driving criterion shifts, it remains important to establish that these shifts actually reflect the adjustment of memory process (i.e., the interpretation of evidence from memory), as opposed to motor or response learning.

To probe the psychological locus of learning, we first applied computational modeling to the behavioral data collected during the fMRI experiment. A liberal criterion shift due to either memory-level or response-level learning will have the same result on response proportions (increased proportion of "old" responses) and thus standard signal-detection theory analyses cannot distinguish between these two alternatives. However, the DDM framework accounts for both response proportions and the distribution of

response times associated with response types (Ratcliff & McKoon, 2008), and includes separate parameters that correspond to the memory-level learning and response-level learning alternatives (Ratcliff & McKoon, 2008).

In the DDM, drift criterion (*dc*) determines whether a given level of retrieved memory evidence is considered evidence of oldness or evidence of newness. The response bias (*z*) parameter determines whether an "old" or a "new" response is more likely, regardless of the level of memory evidence. As such, shifts in *dc* or *z* map in a straightforward way to the memory-level versus response-level accounts of criterion shift, respectively. These parameters are identifiable because they make distinct predictions regarding the response time distributions for "old" and "new" responses (Ratcliff & McKoon, 2008).

We first compared two DDM models: one in which the drift criterion was allowed to vary across participants and over time (memory-level learning), and one in which response bias was allowed to vary across participants and over time (response-level learning). The Drift Criterion model provided a better fit to the data (Drift Criterion model DIC: 6457.97; Response Bias model DIC: 6570.35; see *Methods*; **Supplemental Table 3.3**; **Supplemental Figure 3.8**). This suggests that the adjustments made during the recognition test involved memory-level learning and independently replicates previous behavioral work (*Chapter 2*). Finally, the MS-PE striatal signal (shared and unique variance) across participants was correlated with the terminal drift criterion values estimated from the Drift Criterion model ($R = 0.55$, $p = .016$; **Figure 3.4b**).

Although the Response Bias model did not fit the data as well as the Drift Criterion model, we note that the correlation between striatal MS-PE and the terminal

response bias values from the Response Bias model approached significance (R = 0.45, *p* = .052). Thus, we fit a final model that included both drift criterion and response bias and therefore allowed the parameters to compete to explain the patterns in the behavioral data. In this combined model, striatal MS-PE reliably predicted terminal drift criterion (R = 0.47, *p* = .041) but there was no relationship with terminal response bias (R = 0.07, *p* = .790). Taken together, these results show that shifts in drift criterion better fit the behavioral data and are closely linked to the MS-PE signals in the brain.

Transfer of Recognition Memory Criterion to Free Recall

To further distinguish between the response-level and memory-level alternatives and to demonstrate the generalizability of learning, we conducted a separate behavioral experiment to test whether the adjustments made during the recognition test extended to mnemonic evaluation processes involved in monitoring free recall performance. An individual's recognition criterion is predictive of that individual's false recall rate (i.e., critical lure recall) in the Deese-Roediger-McDermott (DRM) false memory paradigm (Kantner & Lindsay, 2012). This suggests that evaluating memory evidence during recognition memory decisions and monitoring freely recalled items for output can tap a shared underlying evaluative process. For example, free recall participants may generate candidate responses and then assess their associated memory strength, verbally reporting those responses that have sufficient evidence for oldness.

If the criterion shifts we observe in the biased-feedback recognition paradigm are due simply to shifts in response bias (response-level learning), then terminal recognition criterion does not reflect memory evaluation, per se, and therefore would not predict

DRM false recall. However, because the neuroimaging and computational modeling indicate that the criterion shift is due to adjustments in a memory evaluation process (memory-level learning), we predicted that terminal criterion following biased-feedback-induced shifts would successfully predict DRM false recall rates.

Two new groups of participants completed a recognition memory task. One group received biased feedback that induced a liberal criterion and one group received biased feedback that induced a conservative criterion (see *Methods*; **Supplemental Figure 3.9**). Immediately following the recognition test, participants completed a DRM free recall task for ten DRM lists, each comprising words semantically related to a critical lure word that was omitted from the study list.

We found that terminal criterion in the biased-feedback recognition task predicted DRM critical lure recall: participants who adopted a more liberal criterion recalled more critical lures ($R = -0.357$, $p = .022$; **Figure 3.5**). Additionally, we found that average lure recall differed between groups as predicted: the group biased towards a liberal criterion was more likely to recall critical lures than the group biased towards a conservative criterion ($t(39) = 1.706$, one-tailed $p = .048$, Cohen's $d = 0.531$; **Figure 3.5**).

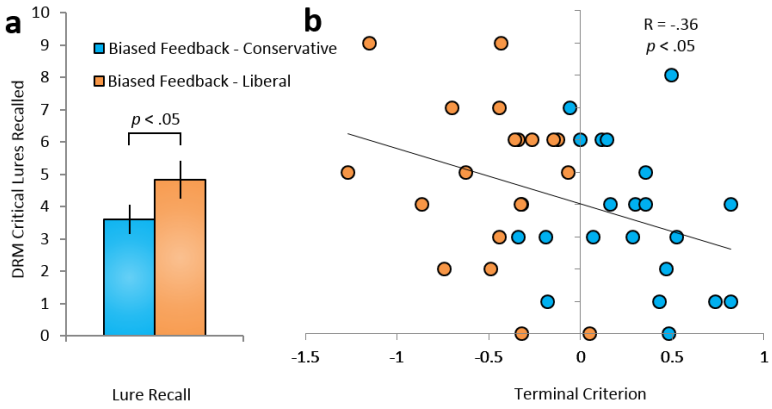

Figure 3.5. Free recall transfer experiment performance. (a) Average critical lure recall for two groups of independent participants that completed a free recall DRM paradigm immediately after a recognition test with biased feedback that induced either a conservative (Biased Feedback - Conservative) or liberal (Biased Feedback - Liberal) criterion. The group that learned a conservative recognition criterion recalled fewer critical lures than the group that learned a liberal recognition criterion. *P* value shown for one-tailed *t*-test between groups. Error bars are SEM. (b) Criterion at the end of the biased-feedback recognition paradigm predicted the number of critical lures recalled.

Discussion

Here, we provide convergent evidence that memory strategies can be learned through RL mechanisms that rely on EV and PE signals deriving from memory retrieval outcomes themselves. Specifically, we found that striatum tracked a trial-by-trial PE signal following feedback in a recognition memory test, with the PE signal scaled relative to the strength of retrieved memory evidence. This PE signal was predictive of individuals' final criterion in the recognition test. Next, we showed that the behavioral results were best fit by a DDM in which learning occurred at the level of the evidence evaluation rather than a specific response bias. Finally, we observed that the criterion

learned during recognition memory transferred to free recall. The computational modeling and free recall results are both consistent with the learning of a general criterion by which mnemonic evidence is evaluated rather than a specific response tendency.

These results advance an emerging literature on facilitative interactions between nondeclarative and declarative memory systems (Scimeca & Badre, 2012; Shohamy & Adcock, 2010). Although previous work has focused primarily on interactions during encoding (Shohamy & Adcock, 2010; Wittmann et al., 2005), the current study provides empirical evidence that the nondeclarative system also influences cognitive control of declarative memory retrieval. Specifically, memory outcome-derived value could be one means by which mnemonic control strategies are learned. Previous work has proposed that PEs may serve to evaluate and update the retrieval strategies that give rise to clustering patterns in free recall tasks (Becker & Lim, 2003; Long et al., 2010). And, in the domain of semantic retrieval, a recent study of rule-guided retrieval found that striatum was sensitive to unexpected outcomes and tracked the degree to which relying on the rule benefitted subsequent memory decisions (Badre et al., 2014). However, the present study provides the first direct evidence that learning of memory control processes can be mediated by PEs coded by the nigra-striatal system.

The present results further inform the functional form by which memory outcomes confer value. In the context of the present study, the strength of retrieved memory was inherently valued, and indeed, PEs derived from this memory strength signal drove striatal responses. By contrast, there was little evidence that the overt response outcome drove PE or striatal responding beyond the variance this signal shared with MS-PE. Though perhaps counterintuitive from the perspective of conventional RL,

the idea that match to memory would be inherently valued is plausible given prior research. For example, mere exposure effects have been widely observed in social psychology, wherein people show preference for objects, people, and images they have consistently encountered previously (Zajonc, 1968). Further, general recognition of an object often triggers further retrieval, so called familiarity-gated recollection (Hayama & Rugg, 2009; Kahn et al., 2004). Put another way, evidence of oldness elicits a type of mnemonic approach response.

Further consistent with the conclusion that value was mnemonic in origin, the MS-EV signal correlated with a set of regions that overlapped with regions in the "core retrieval network" (Elward et al., 2014; Rugg & Vilberg, 2013; Spaniol et al., 2009). Among these were activations in ventromedial frontal areas previously associated with meta-mnemonic "feelings-of-knowing" (Schnyer et al., 2004) and memory schemas (Warren, Jones, Duff, & Tranel, 2014). An intriguing, though speculative, possibility is that initial retrieval of evidence from long-term memory results in a schema match or schema activation, and it is this congruency between retrieval outcome and a prevailing schema representation that is valued. This hypothesis requires testing, but has clear implications for how learning and valuation may occur in more complex tasks and situations.

Outside of the absence of robust ventral striatal activation, the regions correlating with the EV signals also overlapped with the network of regions that have been shown to encode subjective value; and indeed some of these regions are the same as those in the "core retrieval network" discussed above (De Martino et al., 2013; Knutson et al., 2005; Roy et al., 2012). In particular, overlapping areas in medial PFC and precuneus correlated

with unique signals from both MS-EV and ERO-EV. It is a notable that a recent decision-making study showed a similar pattern of results: in an economic choice task, both value and subjective confidence were separately identifiable in overlapping areas of ventromedial PFC and precuneus (De Martino et al., 2013). This decision-making study also found that subjective confidence, but not value, was represented in a more anterior region in right rostrolateral prefrontal cortex. We similarly observed activation for ERO-EV, but not MS-EV, in anterior PFC, suggesting that more anterior regions of PFC may have access to metacognitive confidence signals but not value signals in both economic and mnemonic decisions.

It is important to emphasize that the present results do not necessarily indicate that response outcome has no impact on learning or recognition criterion. We also observed activation that tracked ERO-PE in right inferior frontal gyrus and bilateral inferior parietal cortex. Although we did not have *a priori* predictions regarding this signal, an intriguing post hoc account comes from research in the decision-making literature that has distinguished between reward prediction errors (coded in striatum) and state prediction errors (Gläscher et al., 2010). State prediction errors (SPEs) occur when observed state transitions deviate from an internal model of the environment, and these SPEs have been identified in lateral prefrontal cortex and inferior parietal cortex (Gläscher et al., 2010). The ERO-PE signal may reflect participants' surprise when feedback deviates from metacognitive expectations (Badre et al., 2014), e.g. positive feedback following a low confidence response. However, additional work is needed to explore the potential role of SPEs in declarative memory.

The present work not only informs the field of declarative memory research but also provides a new domain in which to generalize RL theory. Recent years have seen a growing understanding of the contributions of striatum and RL to higher-level cognition (Badre & Frank, 2012; Frank et al., 2004; Gläscher et al., 2010; Haber & Knutson, 2010; Maia & Frank, 2011), and the dynamics of RL in declarative memory potentially differ in important ways from typical RL tasks. Many RL tasks involve learning simple response-outcome mappings for concrete stimuli, although these principles have been extended to more complex actions and policies (Badre & Frank, 2012; Gläscher et al., 2010). Our results suggest that RL taking place during recognition memory is not at the level of a simple stimulus-response-outcome mapping. Instead, the DDM results and the demonstration of successful transfer to free recall suggest that the RL mechanisms update an evaluative control process that operates on the latent and multidimensional contents of retrieved memory representations. In other words, RL principles generalize beyond overt actions and action policies, to "cognitive" actions like memory retrieval. Further, the fact that MS-PE rather than ERO-PE drove striatal responses and learning in the current task demonstrates that depending on the task and type of learning, value can derive from complex latent processes, instead of (or in addition to) the overt outcome of a plan or response.

The current study provides a clear motivation to assess striatal contributions to declarative memory in patients with nigra-striatal dysfunction, such as Parkinson's or Huntington's disease. Although these patient populations show profound impairments in implicit learning tasks (Knowlton et al., 1996), they also show more subtle deficits in declarative memory retrieval similar to the deficits seen in patients with frontal lobe

damage, and this pattern may be due to their inability to appropriately evaluate and update retrieval strategies (Scimeca & Badre, 2012). More generally, our findings suggest that successful control of declarative memory retrieval can be thought of as a skill that is learned and updated through the non-declarative reinforcement learning system. This novel perspective, combined with the observation of successful transfer of learning across retrieval tasks, may have important implications beyond the lab, such as for memory training interventions in the clinic or the classroom.

Methods

Experiment 1: FMRI study of recognition criterion learning from biased feedback

Participants. Nineteen right-handed adults (age = 18-29 years, mean = 22.8 years; 7 women) with normal or corrected-to-normal vision completed the neuroimaging study and are included in the analyses. All participants were without psychiatric or neurological conditions, contraindications for MRI, or medications affecting the CNS. All participants gave written informed consent and were compensated $20/hour according to guidelines established and approved by the Institutional Review Board of the Research Protections Office at Brown University. Four additional participants completed the experiment but were excluded *a priori* from all analyses because of excessive head movement detected during preprocessing (> 3mm translation in a single run). We did not predetermine sample size, but our sample size is within the range used in previous behavioral studies of this paradigm (Han & Dobbins, 2008, 2009; *Chapter 2*), as well as previous fMRI studies of memory from our lab (e.g. Badre et al., 2014).

Task procedure and analysis. The stimulus set consisted of 640 nouns naming concrete objects. For each participant, 280 words were randomly selected to be used in the study phase, 240 of which were included as studied items in the test phase. An additional 240 words were randomly selected to be used as unstudied (lure) items in the test phase.

The study task took place on a laptop computer outside the scanner. Participants were instructed to make one of two semantic judgments about each presented item and were not forewarned of the upcoming recognition test. On each trial, a word was presented for 300 ms, followed by a fixation screen for 1 s, followed by the response prompt. The prompt indicated which semantic judgment to make by providing the possible answer choices: “organic/inorganic” or “small/large” (relative to a typical-sized

shoebox). Participants had 1s to indicate their response with a key press and failure to respond on time elicited an auditory tone as feedback; no other feedback was provided for any responses. Participants completed 280 trials with an inter-trial interval of 1.2 s. Self-paced breaks were provided after the 96th and 192nd trials. The order of the study words was randomized for each participant.

The recognition test was completed within the MRI scanner. The test phase consisted of 480 trials divided into six runs each seven minutes and ten seconds in length; participants were allowed a self-paced break between runs. Participants were told that they would be presented with a mix of studied (old) and unstudied (new) items and to respond to the test prompt "Is this item old?" by indicating "YES" with their index finger and "NO" with their middle finger on the response box (**Figure 3.2a**). We refer to "YES" responses as "old" responses and "NO" responses as "new" responses throughout the chapter. Participants had 1.5 seconds to initiate their response from the time the test item appeared on the screen; response time (RT) was coded based on the elapsed time between item presentation and the initiation of the button response.

Participants were further instructed to indicate their confidence by holding down the response button. This caused a bar to move along an unfilled rectangle underneath the corresponding "YES" or "NO" response (**Figure 3.2a**). Participants were told that the top of the rectangle corresponded to high confidence and the bottom of the rectangle corresponded to low confidence. They were told to position the bar to match their confidence and encouraged to use the whole range of confidence ratings for their responses. The bar moved at a speed such that it took 1.5 seconds to move across the entire rectangle, which was 300 pixels high. To deconfound confidence rating from the

movement of the bar, the bars started at the top of the rectangles on half of trials and the bottom of the rectangles on the other half. This positioning was always visible to the participant before they initiated their response. If participants wanted to indicate a confidence rating that corresponded to the end of the rectangle where the bar started, they were told to press the response button very briefly (e.g. a low confidence response when the bar starts at the bottom of the rectangle). The test prompt, probe word, and answer choices remained on screen for the duration of the confidence rating.

A fixation screen of variable duration followed the response/confidence rating and was followed by a feedback screen indicating either "Correct!" or "Incorrect!" (**Figure 3.2a**). False positive feedback was provided following approximately 70% of incorrect "old" responses to new items (false alarms). For each participant, the experimental script tracked the range of confidence responses provided for false alarms to determine the median false alarm confidence rating before each trial. 100% of false alarms made with a confidence rating below this median received false positive feedback, and approximately 40% of false alarms made with a confidence rating above this median received false positive feedback. All other trials received veridical feedback.

False feedback was targeted to low confidence false alarms because both the MS-PE and ERO-PE alternatives predict that this pattern of feedback will result in a large behavioral effect: a large liberal shift in criterion (see *Chapter 2*). After feedback, a filler fixation screen was presented such that the total length of the trial, not counting the variable inter-stimulus interval between response and feedback, summed to 4 seconds. This filler fixation screen transitioned seamlessly to the inter-trial interval fixation screen of variable duration.

Design efficiency was optimized using a combination of Optseq2 (http://surfer.nmr.mgh.harvard.edu/optseq/) and custom-made scripts. An equal number of studied and unstudied items were presented in each run; the probe word presented on each studied/unstudied trial was randomly drawn from the studied and unstudied items for each participant. The order of studied and unstudied items, the timing of the inter-trial interval, and the timing of the inter-stimulus interval were optimized to maximize detection power and allow us to separably estimate effects associated with response from those associated with feedback. The duration of the inter-trial interval fixation had a mean of 2 seconds and a range of 1-11 seconds. The duration of the inter-stimulus fixation between response and feedback had a mean of 2.85 seconds and a range of 1-6 seconds.

To acquaint participants with the response method used for the recognition test, they completed an ostensibly unrelated estimation task immediately following the study phase and before the test phase. The display was visually similar to the display used in the recognition test, but the test word was replaced with either "right" or "left" and a number between 0 and 100. Participants responded with either the left or right arrow key on the laptop and moved the bar to the corresponding position in the appropriate rectangle. They completed 40 trials of this practice task. Feedback was only provided on non-response trials.

We computed the recognition criterion for each run (80 trials) of the scanning session using the standard signal-detection approach (Macmillan & Creelman, 2005; Wickens, 2002). Criterion was calculated as:

$$\frac{z(Hit\ Rate) - z(False\ Alarm\ Rate)}{-2}$$

Eq. 1

Initial criterion refers to criterion for the first run and terminal criterion refers to criterion for the final run. These criterion values were submitted to a repeated-measures ANOVA with Run as a factor, and follow-up *t*-tests as described in the Results. Note that we use the terminal criterion values, instead of a difference score between terminal and initial criterion, for our correlation analyses because false positive feedback was possible immediately from the onset of the experiment, and thus initial criterion does not provide a clean baseline to compute a difference score.

fMRI data collection and preprocessing. Whole-brain imaging was performed on a Siemens 3T TIM Trio MRI system. High-resolution T1-weighted (multi-echo MP-RAGE) anatomical images were acquired for visualization (TR = 2200 ms; echo times = 1.54 ms, 3.36 ms, 5.18 ms, 7.00 ms; flip angle = 7º; 144 sagittal slices; 1.2 × 1.2 × 1.2 mm). Functional images were collected over six runs using a gradient-echo echo-planar sequence (TR = 2 s; echo times = 28.0 ms; flip angle = 90º; 40 axial slices; 3 × 3 × 3 mm). Head motion was restricted using firm padding surrounding the head. Stimuli were projected onto a screen and viewed through a mirror attached to a 32-channel head coil. Responses were provided with the right hand through a Mag Design and Engineering MRI-compatible four-button response pad. The first five volumes of each functional run were discarded to allow for T1 stabilization.

Preprocessing and data analysis were performed using SPM8 (www.fil.ion.ucl.ac.uk/spm). Image data quality was first assured via visual inspection and using customized versions of TSDiffAna (http://sourceforge.net/projects/spmtools/) and ArtRepair software (http://cibsr.stanford.edu/tools/human-brain-project/artrepair-

software.html). Functional images were then corrected for differences in slice acquisition timing by resampling all slices to match the first slice. Images were then motion corrected across all runs using B-Spline interpolation. For four participants that had head motion greater than 3mm across the entire session but not within a single run, motion-correction was applied separately for each run. As already noted, participants with movement of greater than 3mm within a run (N=4) were excluded. Nuisance regressors were included for all participants to account for run-to-run variance. Data were then normalized based on MNI stereotaxic and then spatially smoothed with an 8-mm FWHM isotropic Gaussian kernel.

fMRI data analysis and visualization. Data analysis was conducted under the assumptions of the general linear model as implemented in SPM8. Single subject effects were estimated using a fixed-effect model. All regressors were generated by convolving stick functions (duration = 0) with a canonical hemodynamic response function and its temporal derivative. Nuisance regressors were included to account for run-to-run variance and low-frequency signal components. Linear contrasts at each voxel were used to obtain subject-specific estimates for each effect. These estimates were entered into second-level analyses treating subjects as a random effect, using a one-sample *t* test against a contrast value of 0 at each voxel.

Unless otherwise noted, activations detected with whole-brain analyses were considered statistically reliable to the extent that they survived a False Discovery Rate (FDR)-corrected threshold of $p < .05$ at the cluster level. Whole-brain maps were initially thresholded at $p < .001$, uncorrected, and cluster corrected to $p < .05$ using SPM's FDR algorithm. The critical cluster extent for each contrast is listed in each corresponding

table. Because the analysis contrasting unique MS-PE versus unique ERO-PE is particularly conservative, we report activation using an uncorrected $p < .001$ threshold. Given our *a priori* hypotheses regarding effects in striatum, we also used small volume correction (SVC) using an anatomical striatum mask that combined the Automated Anatomical Labeling database (AAL; Tzourio-Mazoyer et al., 2002) definitions of bilateral caudate and putamen, applying Family-Wise Error (FWE) height correction across this small volume. We explicitly note this alternative approach in the Results. We note that applying this alternative correction method (SVC using $p < .05$ FWE height correction within the striatal mask) across all of our reported whole-brain analyses provides a pattern of results that is qualitatively similar to the FDR cluster-corrected results that we report. Such consistency between whole-brain cluster correction and SVC height correction suggests robustness to our key striatal results.

All figures are from group contrasts rendered on the MNI canonical brain in neurological convention. The statistical thresholds used for display purposes are listed in the figure captions. All coordinates are in MNI space and correspond to the peak voxel in the cluster. Local maxima reported by SPM within a cluster are also reported. Microanatomical labels were determined based on inspection of the Talairach and Tournoux atlas (Talairach & Tournoux, 1988), and macroanatomical labels were determined using the Harvard-Oxford Probabilistic Atlas implemented in FSL (www.fmrib.ox.ac.uk/fsl) and prior literature.

Whole-brain analyses were complemented by Region of Interest (ROI) analyses. ROIs were defined using anatomical masks from the AAL database (Tzourio-Mazoyer et al., 2002). As already noted, our striatum ROI was constructed by combining the AAL

definitions of bilateral caudate and putamen (total volume: 32,232 mm^3), as our lab has used previously (Badre et al., 2014). We also constructed an anatomical ROI guided by the functional activations associated with our ERO-PE regressor, in order to test for an ROI × effect interaction, as described in the Results. This ROI was constructed by combining the AAL definitions of bilateral inferior parietal cortex and right inferior frontal gyrus opercularis (total volume: 41,528 mm^3). To assess ROI activation, the time-series signal averaged across the entire ROI was extracted using the MarsBaR toolbox (Brett, Anton, Valabreague, & Poline, 2002), and the respective GLM was regressed against this time-series. The parameter estimates (beta values) associated with each regressor were computed and the parameter estimates for regressors of interest were subjected to repeated-measures analysis of variance, *t* tests, and correlation analyses as described in the Results. All reported *p* values are two-tailed unless explicitly noted otherwise, and considered significant at $p < .05$. All error bars indicate standard error of the mean.

fMRI general linear models. Parametric regressors for fMRI analyses were calculated using trial-by-trial Expected Values (EVs) and Prediction Errors (PEs) under both the Memory Strength and Expected Response Outcome alternatives. For the Memory Strength alternative, we transformed response confidence on each trial to derive a Memory Strength EV (MS-EV) as depicted in **Figure 3.1**, such that high confidence "new" responses were coded as the lowest memory strength and high confidence "old" responses were coded as the highest memory strength. For "new" responses, maximum confidence corresponded to an MS-EV of 0 and minimum confidence corresponded to an

MS-EV of 0.5. For "old" responses, minimum confidence corresponded to an MS-EV of 0.5 and maximum confidence corresponded to an MS-EV of 1. For each trial, we calculated the MS-PE by taking the difference between the presented outcome (veridical and false positive feedback: 1; negative feedback: 0) and MS-EV for that trial. For the Expected Response Outcome alternative, we derived an Expected Response Outcome EV (ERO-EV) on each trial that was directly related to response confidence: for both "new" and "old" responses, minimum confidence responses corresponded to an ERO-EV of 0 and maximum confidence responses corresponded to an ERO-EV of 1, as depicted in **Figure 3.1**. For each trial, we calculated the ERO-PE by taking the differences between the presented outcome and ERO-EV for that trial.

We first constructed two versions of our GLM: a Memory Strength GLM and an Expected Response Outcome GLM. The GLMs were identical except for the parametric regressor that was included. We included a condition regressor for trial onset and a condition regressor for the presentation of feedback. In each GLM, the trial onset regressor was modulated by the respective EV parametric regressor and the feedback presentation regressor was modulated by the respective PE parametric regressor. These parametric regressors were convolved with the HRF and temporal derivative basis functions. We further included separate condition regressors for the trial onset and feedback events of non-response trials on which the participant made an error of omission. These regressors were treated as nuisance regressors; they were not modulated by a parametric regressor and were not analyzed. We included identical nuisance regressors in our third and fourth GLMs described below.

In each of these models the parametric regressor captures all variance associated with that regressor above and beyond the variance captured by the primary condition regressor. For example, the MS-PE regressor in the MS GLM would include only variance above and beyond the main effect of feedback presentation, and would include variance shared with the ERO-PE values as well as any variance unique to MS-PE. Likewise, the ERO-PE regressor in the ERO GLM would include variance shared with the MS-PE values and its unique variance. The average correlation between the MS-PE and ERO-PE predictors after convolution with the canonical HRF was $R^2 = 0.52$ (the correlation was first determined within each participant, transformed to a z-score, averaged across all participants, and transformed back to a correlation coefficient.) The average correlation between the convolved MS-EV and ERO-EV predictors was $R^2 = 0.19$.

Our third GLM included the MS-PE and ERO-PE parametric regressor in the same model. This model included a condition regressor for trial onset and a condition regressor for feedback presentation; these were not modulated by any parametric regressor. Instead, we convolved the MS-PE and ERO-PE values with the HRF and input these values as two user-entered regressors. This allows the two PE regressors to compete for variance in the model and capture only unique variance(Henson, 2007). This is similar to what happens when two parametric regressors are included as modulators on a condition regressor in SPM. By default, SPM performs serial orthogonalization such that the first parametric regressor includes both shared variance and its unique variance, and the second parametric regressor includes only its unique variance. For example, if a hypothetical GLM included MS-PE as the first parametric regressor and ERO-PE as a

second parametric regressor, the ERO-PE regressor would only capture the unique variance.

Finally, including both MS-PE and ERO-PE as user-defined regressors within the same model allows them to directly compete for variance and therefore allows us to perform a direct contrast in the GLM framework. Thus, our fourth GLM included both MS-EV and ERO-EV parametric regressors in the same model. Like our third GLM, this model included a regressor for trial onset and a regressor for feedback presentation; these were not modulated by any parametric regressor. Instead, we convolved the MS-EV and ERO-EV values with the HRF and input these values as two user-entered regressors.

Drift diffusion modeling analysis. In the DDM, the decision process arises from noisy evidence accumulation (or drift) towards one of two response bounds(Ratcliff, 1978). For recognition decisions, one bound corresponds to an "old" response and the other bound corresponds to a "new" response. Thus, over time, evidence accumulation drifts towards one of these response bounds at a rate (i.e., the drift rate; *v*) that is determined by the strength of evidence recovered from memory. Evidence that is evaluated as signaling "old" leads to a positive drift rate, towards the "old" boundary. Evidence that is evaluated as signaling "new" leads to a negative drift rate, towards the "new" boundary. The accumulation process is terminated when the drift reaches either boundary, and the corresponding response is made. Stronger evidence would lead to a larger drift rate, and thus, a faster response.

The response bias, *z*, determines the starting point of the diffusion process relative to either of the response boundaries. When bias is closer to the "old" response boundary, participants are overall more likely (requiring less evidence) to respond "old"

independently of the strength of the evidence (drift rate). Shifts in response bias result in a larger effect on the leading edge of the RT distribution relative to shifts in the drift criterion (Ratcliff & McKoon, 2008).

The drift criterion parameter, *dc*, in the DDM determines whether observation of a signal of a given strength is more likely to elicit drift towards the "old" or "new" boundary. This is not thought to change how a memory strength or match signal is retrieved from long term memory (which is impacted by factors such as strength of studied items) but instead how this match signal is evaluated and mapped onto drift rate (Criss, 2010; Ratcliff & McKoon, 2008; Ratcliff & Starns, 2009; Starns, Ratcliff, & McKoon, 2012; White & Poldrack, 2014). From this perspective, the drift process can be thought of as a series of small signal detection processes, where at each time step, the evidence for an item is sampled. If that sample is greater than the old/new criterion then the drift is incremented toward "old," and is otherwise incremented toward "new." The criterion of those signal detection processes conducted at each time step is the drift criterion. Changes in drift criterion have a larger impact on the tail than do changes in bias (Ratcliff & McKoon, 2008).

The recognition decision was modeled using the DDM, following prior work (Ratcliff et al., 2004). For all model variants, parameters were fit to the accuracy and RT data (from both correct and error trials). We estimated bias (*z*) and drift criterion (*dc*); along with parameters for an initial non-decision interval (*t*), thought to reflect processes like encoding of the target; and the distance between the boundaries (*a*). Response time data were examined for outliers: any response faster than 300ms was considered a guess or an error and excluded from the analyses (Criss, 2010; Ratcliff et al., 2004; Starns,

Ratcliff, & White, 2012). Non-response errors (when participants failed the response deadline) were excluded from analysis. These constituted fewer than 5% of trials on average across participants.

Model fitting was conducted using the Hierarchical Drift Diffusion Modeling (HDDM) module (Wiecki et al., 2013). This Python module uses a hierarchical Bayesian estimation procedure that fits the DDM parameters based on all participant data simultaneously. This is analogous to random effects estimation in that it treats between-subject variance as a random variable while fitting within-subject parameters simultaneously. A Markov-chain Monte-Carlo (MCMC) procedure estimated the DDM parameters' posterior distributions. 10,000 samples from the distributions were estimated. The first 3,000 samples were discarded (burn in), and of the remaining samples, every tenth sample was retained (thinning). Model convergence was assessed based on Monte-Carlo error (MC error) and visual assessments of chain convergence (Gelman et al., 2004). Model selection was based on minimization of the deviance information criterion (DIC), which is more readily compatible with MCMC estimation than Akaike information criterion or Bayesian information criterion. In our implementation of the model, more liberal values of *z* and *dc* correspond to more positive values. For visual comparison to SDT plots, we have plotted liberal down in all figures.

Experiment 2: Transfer of Memory Criterion to Free Recall

Participants. Forty-six participants (mean age: 20.2 years, range 18-28 years; 28 female) who did not participate in the fMRI experiment participated in the recognition-free recall experiment. Five participants were excluded from all analyses because they indicated on a post-test questionnaire that they were familiar with the DRM paradigm(Kantner &

Lindsay, 2012). All participants gave written informed consent and were compensated $10/hour according to guidelines established and approved by the Institutional Review Board of the Research Protections Office at Brown University. Experimenters were blind to the learning condition assignment of each participant during the experimental session and during coding of free recall responses. We note that a subset of the participants included in this experiment and analysis are included in the analysis of recognition memory criterion described in *Chapter 2*. The two groups included in the present experiment correspond to the FF–Old – Low Confidence and FF–New – High Confidence groups. However, the free recall data reported here have not previously been reported.

Task Procedure and Analysis. Participants first completed a false feedback recognition paradigm similar to the one used for the fMRI experiment. They studied 240 words and completed a recognition test composed of 144 studied words and 144 unstudied words. The test was divided into 6 blocks of 48 trials separated by self-paced breaks. One group, the Biased Feedback – Liberal group, received false feedback on approximately 70% of false alarms and veridical feedback on all other responses. The false feedback was targeted to low confidence responses to maximize the effectiveness of the liberal criterion shift. The second group, the Biased Feedback – Conservative group, received false feedback on approximately 70% of "new" errors (targeted to high confidence responses to maximize the effectiveness of the conservative criterion shift) and veridical feedback on all other responses.

Immediately following the recognition test, participants completed study/test cycles for ten DRM lists (Roediger & McDermott, 1995). The ten lists with the highest

rate of free recall critical lure intrusions from Stadler, Roediger, and McDermott (1999) were selected. Within each list, the words were presented in the order they are listed in Stadler, Roediger, and McDermott (1999). Eight permutations of list order were created and each participant received one of these orders at random. The fifteen study words were presented one at a time on the screen and the participant read the word aloud at a rate of one per 2 seconds. Immediately following study, a prompt initiated the recall phase and participants had two minutes to verbally recall as many words as they "were reasonably sure had been presented on the list".

Raters blind to the experimental condition listened to the recall verbalizations of each participant using the Penn TotalRecall toolbox (http://memory.psych.upenn.edu/TotalRecall) and coded for the presence or absence of the critical lure on each list. The total number of critical lures recalled by each participant was calculated and submitted to *t*-test and correlation analyses. Based on previous results (Kantner & Lindsay, 2012), we predicted that the Biased Feedback – Liberal group would recall more critical lures than the Biased Feedback – Conservative group. Thus, we performed a directional *t*-test and report the one-tailed p-value. For the correlation analysis, we computed the recognition criterion for each block (48 trials) and used the terminal criterion (from the final block) for the correlation with critical lure recall, paralleling the correlational approach used in Experiment 1.

Supplemental Material

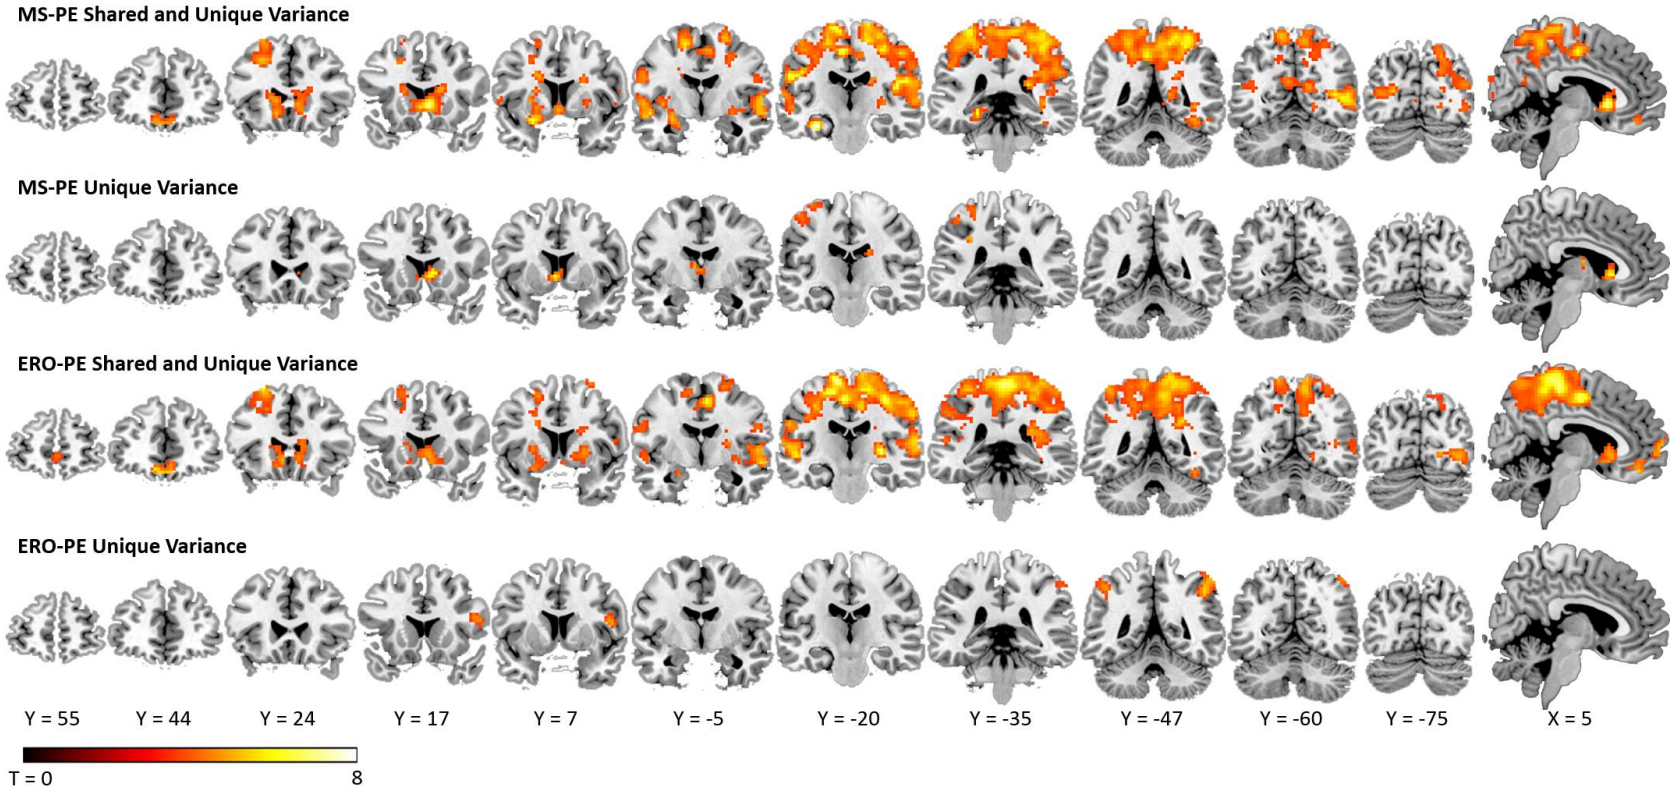

Supplemental Figure 3.6. Whole brain activation maps showing correlates of Prediction Error signals. The first row shows the shared and unique variance associated with the Memory Strength PE (MS-PE); the second row shows just unique variance associated with MS-PE. The third row shows the shared and unique variance associated with the Expected Response Outcome PE (ERO-PE); the fourth row shows just the unique variance associated with ERO-PE. Whole brain statistical maps are thresholded at $p < .05$, FDR cluster corrected. Color maps represent *t* statistics with a range from 0 to 8.

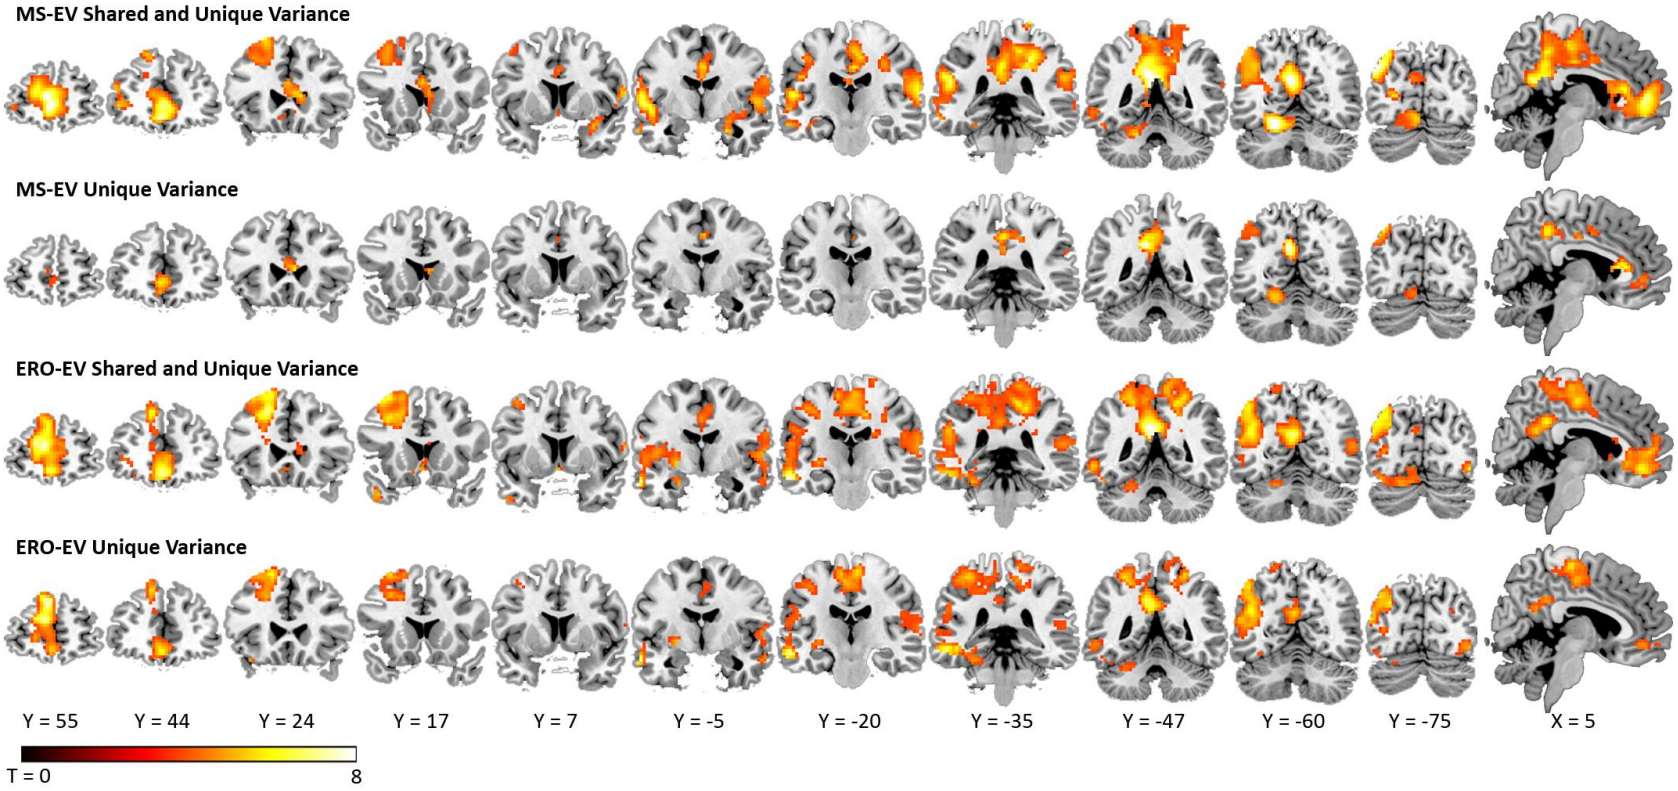

Supplemental Figure 3.7. Whole brain activation maps showing correlates of Expected Value signals. The first row shows the shared and unique variance associated with the Memory Strength EV (MS-EV); the second row shows just unique variance associated with MS-EV. The third row shows the shared and unique variance associated with the Expected Response Outcome EV (ERO-EV); the fourth row shows just the unique variance associated with ERO-EV. Whole brain statistical maps are thresholded at $p < .05$, FDR cluster corrected. Color maps represent *t* statistics with a range from 0 to 8.

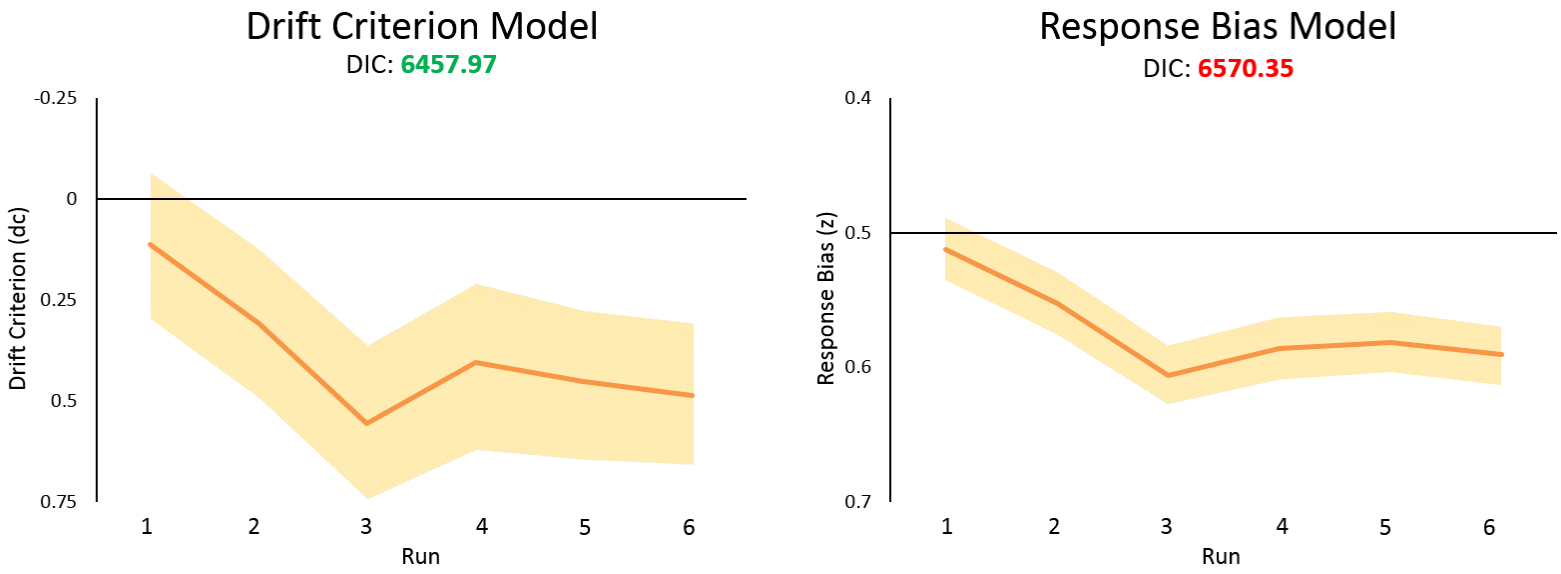

Supplemental Figure 3.8. DDM drift criterion and response bias estimates as a function of time. The left panel plots the drift criterion (dc) values as a function of time derived from the Drift Criterion Model. The right panel plots the response bias (z) values derived from the Response Bias model. More liberal values are plotted down. The DIC (deviance information criterion) values were used for model-selection; the preferred model is the one with a lower DIC. Error bars represent 95% confidence intervals.

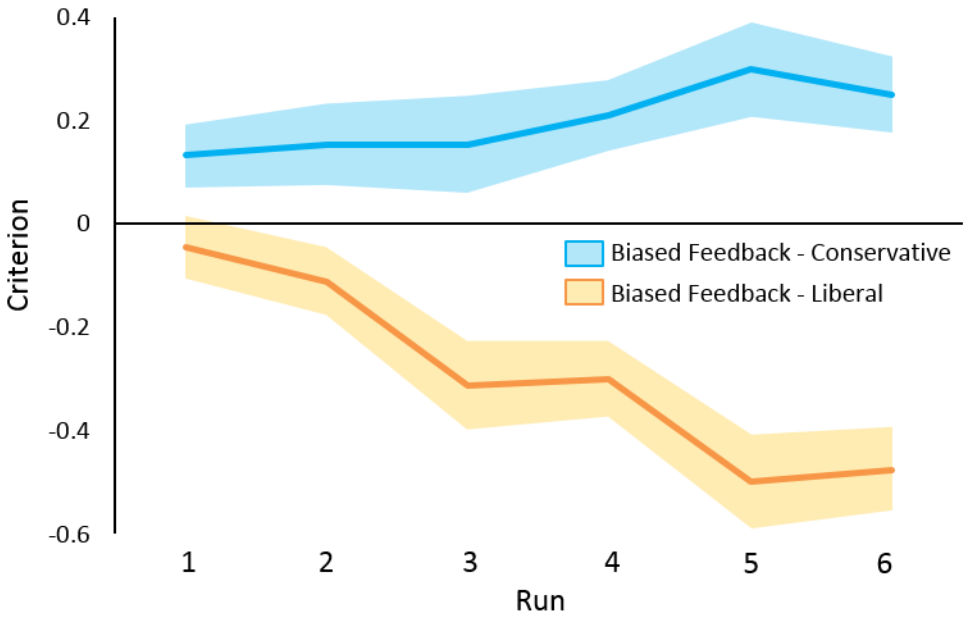

Supplemental Figure 3.9. Recognition criterion in the free recall transfer experiment. Participants in the Biased Feedback – Conservative and Biased Feedback – Liberal adopted a conservative and liberal criterion, respectively, over the course of the experiment. More liberal criterion values are plotted down. Error bars represent SEMs.

Supplemental Table 3.1. fMRI Activations for Prediction Errors from Memory Strength and Expected Response Outcome Alternatives

Region	~Brodmann's Area	MNI Coordinates			Number of voxels	Peak Statistics	
		x	y	z		T	Z
MS-PE Shared and Unique Variance							
Critical cluster extent: 399							
Left hippocampus		-33	-19	-20	8542	7.95	5.15
Left precentral gyrus	6	-6	-16	49		7.32	4.92
Right caudate		6	17	1		7.18	4.87
Left amygdala		-24	2	-17		7.03	4.82
Right parietal operculum cortex	41	45	-19	19		6.89	4.76
Right supramarginal gyrus	40	57	-28	31		6.74	4.70
Left middle frontal gyrus	9	-30	26	43		6.73	4.70
Left postcentral gyrus	1	-57	-19	28		6.67	4.68
Right superior parietal lobe	5	21	-49	61		6.61	4.65
Occipital pole	18	-12	-85	16	399	6.14	4.45
MS-PE Unique Variance							
Critical cluster extent: 157							
Left caudate		-3	8	1	157	6.34	4.54
Right caudate		9	20	4		6.16	4.46
Left thalamus		-3	-7	10		5.76	4.28
Left postcentral gyrus	40	-45	-28	52	252	5.35	4.09
Left precentral gyrus	4	-33	-22	67		4.79	3.80
ERO-PE Shared and Unique Variance							

Critical cluster extent: 126

Right superior temporal gyrus	22	57	-10	-8	7097	8.03	5.17
Left superior temporal gyrus	22	-51	-16	-2		7.74	5.07
Left precentral gyrus	6	-12	-19	70		7.59	5.02
Right precentral gyrus	6	9	-25	64		7.29	4.91
Right putamen		30	-19	4		6.98	4.80
Right cingulate gyrus	24	12	-16	43		6.71	4.69
Left medial frontal cortex	11	-9	44	-14	172	5.75	4.28
Right frontal pole	10	9	62	4		5.35	4.09
Left middle temporal gyrus	21	42	-52	10	126	5.75	4.28
Right occipital pole	18	21	-79	1	165	5.14	3.98

ERO-PE Unique Variance

Critical cluster extent: 77

Right supramarginal gyrus	40	51	-40	49	199	7.12	4.85
Right angular gyrus	40	57	-55	43		4.75	3.78
Right inferior frontal gyrus, pars opercularis	44	48	11	13	77	5.83	4.31
Left angular gyrus	40	-51	-52	49	103	5.53	4.17
Left supramarginal gyrus	40	-45	-49	37		4.35	3.55

All reported clusters were significant at a False Discovery Rate (FDR)-corrected threshold of $p < .05$ at the cluster level. Whole-brain maps were initially thresholded at $p < .001$, uncorrected, and cluster corrected to $p < .05$ using SPM's FDR algorithm. The critical cluster extent for each contrast is listed above.

Rows that denote the peak voxel within a cluster include a value for "Number of voxels". Local maxima within the cluster reported by SPM are listed on subsequent rows.

Supplemental Table 3.2. fMRI Activations for Expected Value from Memory Strength and Expected Response Outcome Alternatives

Region	~Brodmann's Area	MNI Coordinates			Number of voxels	Peak Statistics	
		x	y	z		T	Z
MS-EV Shared and Unique Variance							
Critical cluster extent: 523							
Left precuneus	30	-9	-52	19	5051	10.67	5.92
Left frontal pole	10	-3	56	1		8.89	5.44
Left paracingulate gyrus	10	-6	56	10		8.63	5.36
Left posterior cingulate gyrus	31	-6	-43	31		8.35	5.28
Left inferior parietal lobule	7	-45	-73	31		8.12	5.20
Left medial frontal cortex	10	-3	47	-11		7.71	5.06
Left cerebellum		-18	-64	-17	523	9.31	5.56
Left inferior temporal gyrus	21	-48	-40	-14	1083	7.68	5.05
Left middle temporal gyrus	21	-57	-22	-17		7.28	4.91
Left superior temporal gyrus	22	-48	-1	-5		7.07	4.83
Left parietal operculum cortex	40	-51	-34	22		7.03	4.82
Right amygdala		30	2	-23	943	7.22	4.89
Right postcentral gyrus	40	63	-16	19		6.32	4.53
Right inferior temporal gyrus	20	57	-31	28		6.13	4.45
MS-EV Unique Variance							
Critical cluster extent: 42							
Left precuneus	31	-6	-64	28	566	9.79	5.70
Left posterior cingulate gyrus	31	-12	-46	34		7.99	5.16
Right posterior cingulate gyrus	31	9	-40	40		6.17	4.46
Left inferior parietal lobule	7	-42	-73	46	163	7.05	4.82

Left cerebellum		-12	-67	-14	195	7.02	4.82
Left paracingulate gyrus	32	-3	47	-2	175	6.86	4.75
Right anterior cingulate gyrus	24	6	29	10	122	6.31	4.53
Left anterior cingulate gyrus	24	-12	38	13		4.38	3.57
Anterior cingulate gyrus	32	0	-1	40	48	5.56	4.19
Right parietal operculum cortex	40	57	-31	25	42	4.60	3.69

ERO-EV Shared and Unique Variance

Critical cluster extent: 130

Left frontal pole	10	-12	62	19	2122	9.91	5.73
Left medial frontal gyrus	11	-3	47	-14		9.01	5.48
Left superior frontal gyrus	8	-15	29	64		8.67	5.37
Left middle frontal gyrus	8	-39	14	58		6.69	4.68
Paracingulate gyrus	10	0	50	4		6.30	4.52
Left orbitofrontal gyrus	11	-33	38	-14		5.64	4.23
Subcallosal cortex	32	0	26	-11		5.25	4.04
Anterior cingulate gyrus	32	0	35	4		5.05	3.93
Left posterior cingulate cortex	23	-6	-46	25	4908	9.00	5.47
Left middle temporal gyrus	21	-54	-22	-17		8.18	5.22
Left precuneus	31	-3	-64	25		7.43	4.96
Left lateral occipital cortex	18	-24	-88	1		6.88	4.76
Left superior parietal lobule	7	-24	-43	55		6.35	4.54
Left precentral gyrus	6	-3	-16	52		6.19	4.47
Right postcentral gyrus	5	12	-31	55		5.96	4.37
Right superior parietal lobule	5	21	-49	55		5.91	4.35
Left putamen		-24	-1	-2		5.72	4.26
Left inferior parietal lobule	40	-48	-67	43	666	7.60	5.02
Left inferior parietal lobule	39	-45	-70	25		6.95	4.79
Right lateral occipital cortex	19	48	-70	-8	149	7.03	4.81
Right superior temporal gyrus	21	60	-10	-8	708	6.12	4.45

Right parietal operculum cortex	13	54	-28	19		6.03	4.40
Right precentral gyrus	6	63	5	13		5.34	4.08
Right optic pole	17	18	-94	10	130	5.57	4.20
ERO-EV Unique Variance							
Critical cluster extent: 87							
Left frontal pole	10	-15	59	19	1277	10.18	5.80
Left medial frontal cortex	11	-3	47	-14		9.12	5.51
Left superior frontal gyrus	8	-12	29	52		6.49	4.60
Left orbital frontal cortex	11	-36	26	-20		5.73	4.27
Left orbital frontal cortex	47	-24	35	-14		5.11	3.97
Left middle temporal gyrus	21	-57	-22	-14	2590	7.79	5.09
Left posterior cingulate gyrus	23	-9	-43	34		7.40	4.95
Left superior parietal lobule	5	-30	-43	58		6.02	4.40
Left precuneus	23	-3	-64	25		5.75	4.28
Right precentral gyrus	6	6	-19	64		5.70	4.25
Right superior parietal lobule	5	27	-46	61		5.66	4.24
Left occipital pole	18	-24	-94	-5	231	7.48	4.98
Left inferior parietal lobule	39	-45	-67	40	664	7.23	4.89
Left lateral occipital cortex	39	-48	-73	19		7.23	4.89
Right lateral occipital cortex	18	48	-70	-8	306	6.54	4.62
Right parietal operculum cortex	13	54	-28	19	422	5.40	4.11
Left putamen		-24	-4	-5	87	5.16	3.99
Left cerebellum		-36	-70	-23	91	4.35	3.55

All reported clusters were significant at a False Discovery Rate (FDR)-corrected threshold of $p < .05$ at the cluster level. Whole-brain maps were initially thresholded at $p < .001$, uncorrected, and cluster corrected to $p < .05$ using SPM's FDR algorithm. The critical cluster extent for each contrast is listed above.

Supplemental Table 3.3. Comparison Between Drift Diffusion Models

Model	DIC	deviance	pD
Drift Criterion	6457.98	6313.06	144.91
Response Bias	6570.35	6447.43	122.92

Lower DIC values indicate better fitting model.
DIC: deviance information criterion.

CHAPTER 4

Remembering to Learn: Perceived Oldness and Task Goals Modulate Encoding at Retrieval

Successful memory retrieval has been linked to activity in brain systems associated with value and reward, such as striatum. However, the factors that contribute to these putative value signals and the behavioral consequences of these signals remain unclear. Based on the established role of reward and value in enhancing episodic encoding, we hypothesized that value signals during memory retrieval support adaptive re-encoding of retrieved information (the first hypothesis outlined in Chapter 1). Here we use a variant of the testing effect procedure to test the behavioral predictions of this hypothesis. In an initial test, we manipulated the relationship between perceived oldness and task goals, two factors which have been linked to striatal activity during memory retrieval (Han et al., 2010). Crossing these factors in an initial retrieval test allowed us to assess the independent impact of these factors on performance in a final test. We found that both perceived oldness and goal-attainment support retrieval-based re-encoding, when task goals are made sufficiently salient. We interpret the re-encoding effects of the current studies as a behavioral consequence of putative striatal value signals. Thus, these behavioral results provide converging evidence consistent with multiple co-occurring striatal value signals during retrieval and suggest that an important function of these signals is to support re-encoding of retrieved information.

Introduction

A core function of adaptive memory is to encode information that may prove valuable in future contexts and to retrieve information that is valuable in a current context (J. R. Anderson & Milson, 1989; Scimeca & Badre, 2012; Shohamy & Adcock, 2010). This requires that the memory system be able to appropriately index valuable information such that it is prioritized in memory. The nigra-striatal dopaminergic value system is known to be critical to value-based learning in action selection and decision-making (Cools, 2011; Daw & Tobler, 2014; Haber & Knutson, 2010; O'Doherty et al., 2004), and recent work has begun to explore the role of striatal engagement in declarative memory retrieval (Badre et al., 2014; Elward et al., 2014; Han et al., 2010) and encoding (Murty & Adcock, 2013; Shohamy & Adcock, 2010; Wittmann et al., 2005). However, the functional form by which striatum assigns value to different retrieval outcomes and the subsequent impact of these value signals on future memory remains unknown. Here we use a variant of the testing effect paradigm to test the hypothesis that value signals during episodic memory track task-independent oldness and task-specific goal-attainment, and that these value signals in turn support re-encoding of retrieved information (Scimeca & Badre, 2012).

Functional magnetic resonance imaging (fMRI) studies of human memory have found that striatal activation is commonly observed during simple recognition tasks when contrasting hits (correct responses to studied target items) versus correct rejections (correct responses to unstudied lure items). This finding is typically referred to as the "retrieval success" effect (for relevant meta-analyses, see Kim, 2011; Scimeca & Badre, 2012; Spaniol et al., 2009; Wagner et al., 2005). Despite the ubiquity of this finding, the

factors that drive the striatal response remain unresolved. Recently, Han and colleagues (Han et al., 2010) proposed two non-exclusive accounts for the retrieval success effect. The "retrieval-dependent account" proposes that the striatal response is linked to the perceived oldness of retrieved information (*i.e.* ecphory, *cf.* Schacter, Eich, & Tulvlng, 1978) which underlies "old" responses but not "new" responses. Under this account, the striatal response marks successful recovery of information from long-term memory as a valued retrieval outcome, independent of the goals of the current task. In contrast, the "goal-attainment account" proposes that the striatal response is linked to the motivational significance associated with "old" responses. As Han and colleagues noted, a typical memory test in a laboratory setting will often explicitly or implicitly indicate that the goal of the test is to remember previously studied items, and participants may thus be more motivated to make "old" responses than "new" responses.

Thus, perceived oldness and task goals are aligned in a typical memory test, and dissociating between these accounts of striatal activity therefore requires that task goals be experimentally crossed with perceived oldness. Han and colleagues (Han et al., 2010) accomplished this by incentivizing either "old" or "new" responses with a monetary reward during separate recognition tests. Using fMRI, they observed increased striatal activity for the incentivized response, consistent with the goal-attainment account. However, they also observed above-baseline striatal activity for "old" responses even when "new" responses were incentivized, at a level comparable to the activity observed in a recognition test without incentives. This suggests that striatal activity may reflect both goal-attainment and perceived oldness.

A related fMRI study that combined a source memory task with incentives identified voxels in striatum that exclusively showed memory effects, separate voxels that exclusively showed reward effects, and voxels that showed both effects of memory and reward (Elward et al., 2014). As the authors note, this pattern is not easily accounted for by a single construct (that is, by perceived oldness or goal-attainment alone), and is instead compatible with multiple striatal value signals that co-occur during retrieval.

These previous studies used the pattern of striatal activity associated with different types of retrieval outcomes to make forward inferences about the value of perceived oldness and goal-attainment (Elward et al., 2014; Han et al., 2010; Henson, 2006; Moran & Zaki, 2013). In the present study, we sought to assess putative striatal value signals associated with perceived oldness and goal-attainment by measuring their influence on future behavior: specifically, their influence on subsequent memory for items encountered during retrieval. In studies of encoding, recruitment of the nigra-striatal system by reward and reward anticipation is associated with enhanced encoding (Adcock et al., 2006; Murty & Adcock, 2013; Shohamy & Adcock, 2010; Wittmann et al., 2005). We hypothesized that recruitment of this value system by perceived oldness or goal-attainment would result in a similar enhancement of retrieval-based re-encoding.

To test this hypothesis, we utilized a variant of the testing effect procedure. In a typical testing effect experiment, participants first study a list of items and then one group of participants completes a memory test for those items while the other group completes a second study phase of the same items. On a final test of memory, the group that completed the initial memory test performs better than the group that completed two study phases (P. C. Brown, Roediger, & McDaniel, 2014; Roediger & Butler, 2011;

Roediger & Karpicke, 2006a, 2006b). In the present experiments, all participants completed an initial and final test. Following from the logic of Han and colleagues (Han et al., 2010), we crossed perceived oldness and task-goals during the initial test in order to identify the impact of these factors on performance in the final test.

Although an expanding literature has begun to delineate which manipulations give rise to the testing effect, there currently exists little mechanistic theory regarding the neurocognitive basis for these effects (Roediger & Butler, 2011). Therefore, the current study allows us to address two open and related questions: (1) Is perceived oldness and/or goal-attainment associated with value during retrieval? (2) Do these value signals support retrieval-based re-encoding?

Experiment 1

All participants completed three tasks: a study task, Test 1, and Test 2. We manipulated task goals in Test 1 through a simple instructional manipulation. One group of participants was instructed that Test 1 was a memory test and that their task was to recognize old words that had been presented in the study task (Memory Group). A second group of participants was instructed that Test 1 was a novelty detection test and that their task was to identify new words not previously encountered in the study task (Novelty). All participants were told that Test 2 was a word identification task in which their job was to indicate whether each word was old or new.

Both Test 1 and Test 2 had the same task structure as typical recognition tests: half of the items were targets (items encountered in a previous task) and half of the trials were lures (items not previously encountered). Critically, the target items in Test 2 were

composed of both the targets and lures encountered in Test 1 (see **Figure 4.1** and *Logic and Design* below for more details). This allowed us to sort targets in Test 2 based on two factors: whether they were perceived as old or new in Test 1; and whether or not they were associated with goal-attainment in Test 1. For the Memory Group, perceived oldness and task goals were aligned: "old" responses in Test 1 were perceived as old and associated with goal-attainment. For the Novelty Group, perceived oldness and task goals were crossed: "old" responses in Test 1 were associated with perceived oldness, while "new" responses were associated with goal-attainment. Therefore, the specific pattern of performance for different types of targets in Test 2 allows us to identify the impact of putative value signals on retrieval-based re-encoding during Test 1.

The retrieval-dependent account predicts that both the Memory Group and the Novelty Group will show a re-encoding benefit for perceived oldness. The goal-attainment account predicts that the Memory Group will show a perceived oldness benefit while the Novelty Group will show a perceived newness benefit. Finally, both perceived oldness and task goals may be associated with value, as suggested by the imaging results described above (Elward et al., 2014; Han et al., 2010). If this is the case, then the Memory Group will show a large benefit of perceived oldness (which is aligned with task goals) while the Novelty Group will show an intermediate effect (because perceived oldness and task goals will have opposing effects).

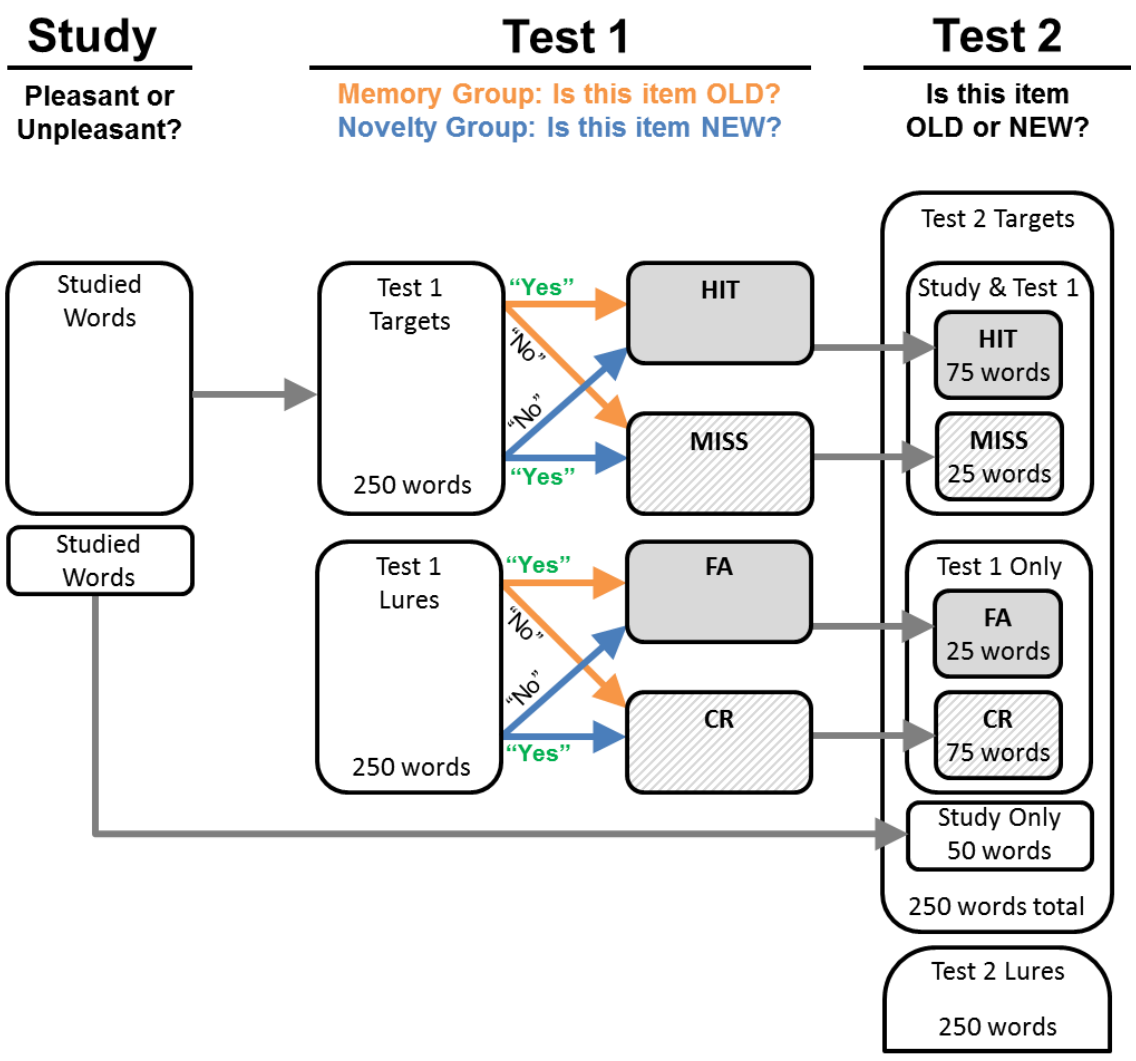

Figure 4.1 Experimental Design for Experiment 1. The task prompts and flow of words across the three tasks in Experiment 2. In the Study task, all participants made a pleasantness judgement. In Test 1, the Memory Group (orange) and Novelty Group (blue) were provided different task prompts. All responses in Test 1 were classified as either Hits, Misses, False Alarms (FAs), or Correct Rejections (CRs), depending on the participant's responses (colored arrows and response in quotes). Items perceived as old (Hits, FAs) are shown with a filled gray background; items perceived as new (Misses, CRs) are shown with a patterned background. Goal-congruent responses are shown in green font. Critically, perceived oldness and task goals are aligned for the Memory Group but crossed for the Novelty Group. In Test 2, targets consisted of items encountered in Test 1 (labeled here by their Test 1 response type) and items encountered only at study. An equal number of new words were used as lures. All participants made an "OLD or NEW?" word identification judgment in Test 2. The primary difference in Experiment 2 is that the ratio of targets to lures in Test 1 was changed from 1:1 to 5:1. See *Logic and Design* in the text for more details.

Method

Participants. Thirty four participants (21 female; age, 18-23 years; mean = 19.6 years) participated in exchange for payment. All participants had normal or corrected-to-normal vision and were native English speakers. Participants gave written informed consent according to guidelines established by the human Research Protections Office of Brown University. Two participants (one from the Memory Group and one from the Novelty Group) were excluded from all analyses because they had d-prime scores more than two standard deviations below the mean, resulting in 16 participants in each group.

Procedure. The experiment consisted of three tasks across two consecutive days. On the first day, participants completed a Study task and Test 1. During the Study task, participants made a pleasantness judgment for a series of 300 words. On each trial, the word was presented centrally and the answer choices (pleasant / unpleasant) were presented on either side of the screen, counterbalanced across participants. Participants had 1.5 seconds to respond by pressing either the "1" key (choice on the left) or "2" key (choice on the right). Following the response, the answer choices would disappear and the word would remain on the screen for the remainder of the trial. A central fixation cross was presented for an inter-trial interval of 0.5 seconds.

Test 1 followed immediately after the Study task. Participants were randomly assigned to either the Memory Group or the Novelty Group at the beginning of the experiment, and the framing of Test 1 differed between the two groups: instructions for each condition were identical except for the name given to the task (memory test or novelty detection test) and the phrase instructing them to either "identify old items" or

"identify new items". To remind participants of the task instructions, the task prompt was presented at the top of the screen for each trial: "Is this item OLD?" or "Is this item NEW?", respectively.

On each trial, the task prompt appeared at the top of the screen with a word presented centrally below. The answer choices ("yes" / "no)" were presented on either side of the screen, counterbalanced across participants, along with an unfilled rectangle used as a visual analog scale of confidence (see Chapters 2 and 3). Participants made their response with the left and right arrow keys and indicated their confidence on each trial by holding down the key: as they held down the key, a small bar moved up the rectangle underneath the corresponding response. Participants were instructed that moving the bar higher up in the rectangle corresponded to higher confidence, and to release the key when the bar was at the appropriate height. If participants did not release the key, the bar stopped at the top of the rectangle and it was coded as maximum confidence. This confidence rating is similar to previous studies in our lab (such as the experiments described in Chapters 2 and 3). All responses were included in analyses, irrespective of participants' reported confidence.

Participants had 2 seconds to make their initial response; the confidence rating took a maximum of 1.5 seconds to make a maximum confidence response. Following the response, the task prompt and word remained on the screen such that each trial was fixed to 3.5 seconds. A central fixation cross was presented for an inter-trial interval of 1 second. Test 1 consisted of 500 trials: 250 target words and 250 lure words (see **Figure 4.1** and *Logic and Design* for more details). Participants were allowed a self-paced break every 50 trials.

Approximately 24 hours later, participants returned to the lab to complete Test 2. The 24 hour delay was used because previous studies of reward-related enhanced encoding have found that reward-related encoding effects are larger after longer delays (e.g. Wittmann et al., 2005), likely due to the role of dopamine in long-term potentiation and consolidation (Huang & Kandel, 1995; Jay, 2003; Lemon & Manahan-Vaughan, 2006; Li et al., 2003). For all participants, this task was referred to as a "Word Identification Task". Participants were informed that some of the words they would be asked to identify were encountered in one or both of the tasks from the first day, while some words were not encountered on the first day. All participants were instructed to "identify whether the word is old or new". The task prompt "OLD or NEW?" or "NEW or OLD?" appeared at the top of the screen, counterbalanced across participants, and participants indicated their response using the left or right shift key (using their index finger on separate hands). The rest of the task timing and stimulus screen was identical to Test 1, with participants again indicating their response confidence using a visual analog scale. Test 2 consisted of 500 trials: 250 target words and 250 lure words (see **Figure 4.1** and *Logic and Design)*. No feedback was provided in any task.

In order to prevent confounds associated with motor responses, separate button responses were used for each of the three tasks. Additionally, the two answer choices for each task (Study: pleasant / unpleasant; Test 1: yes / no; Test 2: old / new) could appear on either left or right side of the screen. The presentation was fixed within each participant: that is, the side of a response choice did not vary from trial to trial. Instead, the side of the responses for each task was counterbalanced across participants, such that these three factors were fully crossed every eight participants.

Logic and Design. **Figure 4.1** shows the flow of words across the Study task, Test 1, and Test 2. The total stimulus set consisted of 800 adjectives. The studied items, Test 1 lures, and Test 2 lures were randomly drawn from this set for each participant. The Study task consisted of 300 words: 250 words were also used as Test 1 targets and 50 words were not included in Test 1 but were included as Test 2 targets ("Study only" items). Words in the Study task were presented in pseudo-random order such that every 30 trials was composed of 25 words to be used as Test 1 targets and 5 words to be used as "Study only" items in Test 2.

Test 1 consisted of 500 words: the 250 studied words served as targets, and 250 unstudied words served as lures. Based on a participant's responses, all words in Test 1 were classified as Hits (correct responses to Test 1 targets), Misses (incorrect responses to Test 1 targets), false alarms (FAs; incorrect responses to Test 1 lures), or correct rejections (CRs; correct responses to Test 1 lures). In **Figure 4.1**, items that were perceived as old (Hits, FAs) are shown with a filled gray background and items that were perceived as new (Misses, CRs) are shown with a patterned gray background. Words in Test 1 were presented in pseudo-random order such that every 50 trials was composed of half targets and half lures.

The colored arrows and quoted responses under Test 1 in **Figure 4.1** indicate the corresponding response for each group that resulted in the four possible Test 1 outcomes (Memory Group: orange arrows; Novelty Group: blue arrows). Responses congruent with task goals are shown in green font. Critically, perceived oldness and task goals are aligned for the Memory Group (Hits and FAs) but these factors are crossed for the Novelty Group such that items perceived as new (Misses and CRs) are congruent with

task goals. This crossing unconfounds the impact of perceived oldness and task goals during Test 1 retrieval-based re-encoding.

A subset of items from each of the four Test 1 response types (Hits, Misses, FAs, CRs) were used as Test 2 targets. 50 words that were encountered during the Study task by not in Test 1 were included as targets in Test 2, for a total of 250 target words. 250 words not encountered during the Study task or during Test 1 served as lures in Test 2. Words in Test 2 were presented in pseudo-random order such that every 100 trials consisted of 50 lures and 50 target items (15 Hits, 5 Misses, 5 FAs, 15 CRs, 10 Study only words).

To ensure that the composition of Test 2 targets would be approximately equal for each participant, only a subset of the Hits (75 words), Misses (25 words), FAs (25 words), and CRs (75 words) from Test 1 were included as Test 2 targets. The specific words from within each response type were chosen at random. If a participant did not have a sufficient number of trials for one response type (e.g. a participant with high accuracy who only made 23 misses in Test 1), these items were replaced with another Test 1 item such that there were exactly 100 targets encountered at Test 1 & Study (Hits and Misses) and exactly 100 targets encountered at Test 1 Only (FAs and CRs) included in Test 2 for all participants (e.g. a Hit item was used in place of a Miss item for participants with an insufficient number of misses in Test 1.)

Analysis approach. The primary measure of interest was accuracy in Test 2 as a function of response type in Test 1. We used the standard signal-detection theory measure d-prime (Macmillan & Creelman, 2005; Wickens, 2002) which is defined as the difference

between z-transformed hit rate and z-transformed false alarm rate. False alarm rate was calculated from "old" responses to Test 2 lures. Hit rates were calculated separately for "old" responses to Hits, Misses, FAs, and CRs (defined by Test 1 response) and Study only items, which were in turn used to generate five d-prime measures, as shown in **Figure 4.2** and **Figure 4.3**.

Results and Discussion

Test 1 Performance. There was no group difference in d-prime (memory group: M=1.50, SD=0.45; novelty group: M=1.69, SD=0.40; $t(30) = 1.27$).

Test 2 Performance: The Role of Perceived Oldness and Task Goals. To assess the role of perceived oldness and task goals on retrieval-based re-encoding, d-prime scores from the four conditions of interest (Hits, Misses, FAs, CRs as defined by Test 1 response) were submitted to a between-subjects repeated measures ANOVA with factors of repetitions (encountered at study and test: Hits, Misses; encountered at test only: FAs, CRs), perceived oldness (perceived as old: Hits, FAs; perceived as new: Misses, CRs), and group. Accuracy rates for these four conditions and for studied-only control items are shown for each group in **Figure 4.2**.

There was a main effect of perceived oldness across groups ($F(1,30) = 98.035$, $p < 0.001$) such that items perceived as old during Test 1 were better recognized during Test 2. To test whether this effect persists even in the absence of veridical memory, we specifically compared the items encountered at test only (FAs and CRs). Items incorrectly perceived as old (FAs) were better recognized than items correctly classified as new (CRs) in both the memory ($t(15) = 3.08$, $p < 0.01$) and novelty groups ($t(15) =$

4.61, $p < 0.001$). This is consistent with studies that have identified increased striatal responses for false alarms compared to correct rejections (e.g. Abe et al., 2008) and suggests that perceived oldness is associated with re-encoding regardless of the veridical content of memory retrieval.

We did not find evidence that task goals modulate re-encoding at retrieval: the statistical interaction between group and perceived oldness was not significant ($F(1,30) = 0.020$). Instead, the effect of perceived oldness was observed within both the memory group and the novelty group, as depicted by the difference scores shown in the bottom row of **Figure 4.2**. The nonsignificant interaction may reflect a true null effect of task goals, or it may reflect a failure of the instructional manipulation. Although the instructions for the novelty group framed Test 1 as a novelty detection task, participants may have still treated the task as a memory test and simply adjusted their responses to match the task prompt. The aim of Experiment 2 was to address this potential issue by increasing the saliency of novelty during Test 1, akin to typical oddball studies of novelty detection.

The Novelty Group had greater overall accuracy compared to the Memory Group (main effect of group, $F(1,30) = 5.910$, $p = 0.021$). It is possible that the novelty detection task supported encoding of items during Test 1 such that overall memory was enhanced for Test 2, despite the absence of a specific effect of task goal. However, an alternative explanation is that the participants randomly assigned to the Novelty Group had better memories than the participants assigned to the Memory Group independent of the group manipulation.

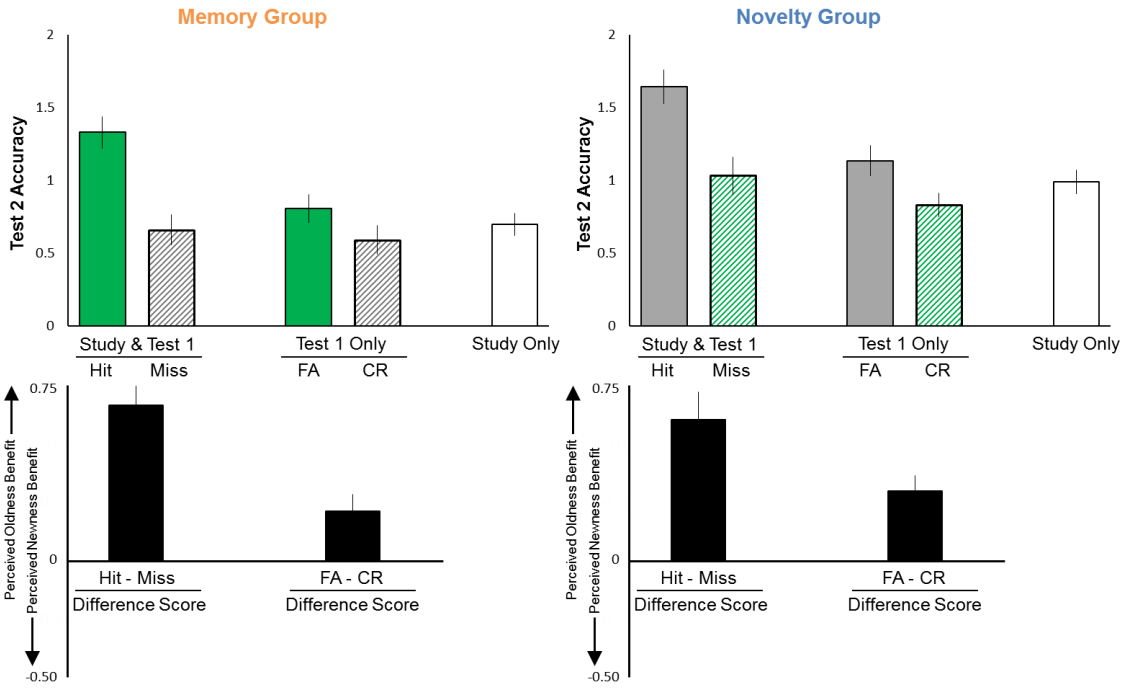

Figure 4.2. Accuracy for All Test 2 Targets in Experiment 1. The top row shows Test 2 accuracy for all Test 2 target items. Items are sorted by their Test 1 response type, as indicated by the horizontal axis labels. Filled bars indicate items perceived as old in Test 1; patterned bars indicate items perceived as new in Test 1; green indicates items that were goal-congruent in Test 1 (see main text and Figure 4.1 for more details). Items encountered during the Study phase only are shown for comparison. The bottom row shows difference scores to summarize the retrieval-based re-encoding benefit of perceived oldness (plotted upward) or perceived newness (plotted downward). Both the Memory Group and Novelty Group showed a perceived oldness benefit such that those items perceived as old in Test 1 were better recognized in Test 2.

Experiment 2

In Experiment 1, we did not find evidence that task goals influenced retrieval-based re-encoding. It is possible that the instructional manipulation was not sufficient to get participants in the Novelty Group to treat Test 1 like a novelty detection test, and instead they simply treated it like a memory test and adjusted their responses to match the test prompt. To make the manipulation of task goals more salient in Experiment 2, we additionally manipulated the base rate of target and lure items in Test 1. A common paradigm used to study novelty detection is the "oddball paradigm", in which novel events are rare events (oddballs). We therefore changed the ratio of target items to lure items to 5:1 in Test 1, in the range of proportions typically used in studies of the oddball paradigm (Braver, Barch, Gray, Molfese, & Snyder, 2001; Linden et al., 1999; Marois, Leung, & Gore, 2000; Mayer, Harrington, Adair, & Lee, 2006; Tricomi et al., 2004). The proportion of targets and lures was changed to 5:1 for both the Memory Group and the Novelty Group. Importantly, this task design still allows us to identify separate effects of perceived oldness and task goals because these factors remain crossed between the two groups.

Method

Participants. Thirty three participants (21 female; age, 18-32 years; mean = 20.1 years) participated in exchange for payment. All participants had normal or corrected-to-normal vision and were native English speakers. Participants gave written informed consent according to guidelines established by the human Research Protections Office of Brown

University. One participant was excluded from all analyses for failing to respond on over 20% of trials in Test 2, resulting in 16 participants in each group.

Procedure. The timing and task structure of the Study, Test 1, and Test 2 tasks were identical to Experiment 1 except for the number of items used in each test, as described below. Critically, the proportion of lure items in Test 1 was reduced from 50% to 17% (100 lure words out of 600 total words), similar to the proportion used in oddball tasks used to study novelty detection. All participants were explicitly instructed that the majority of items included in Test 1 would be items encountered during the initial (study) task. Importantly, the Memory Group was still instructed that they would be completing a memory test and that their job was to identify old words; the Novelty Group was still instructed that they would be completing a novelty detection test and that their job was to identify new words.

Logic and Design. Words in Experiment 2 were drawn from the same pool of 800 words used in Experiment 1; a random subset of only 760 words were used for each participant. The studied items, Test 1 lures, and Test 2 lures were randomly drawn from this subset for each participant. The overall flow of words in Experiment 2 was similar to that depicted in **Figure 4.1**, but a different number of words was associated with each box.

The Study task consisted of 520 words: 500 words were also used as Test 1 targets and 20 words were not included in Test 1 but were included as Test 2 targets. Test 1 consisted of 600 words: 500 studied words served as targets, and 100 unstudied words served as lures. Based on a participant's responses, all words in Test 1 were classified as Hits, Misses, FAs, or CRs. As in Experiment 1, perceived oldness and task goals are

aligned for the Memory Group but crossed for the Novelty Group. Words in Test 1 were presented in pseudo-random order such that every 60 words was composed of 50 targets and 10 lures. To make the high proportion of target items especially salient, the first ten trials were always target words for all participants.

A subset of items from each of the four Test 1 response types (Hits, Misses, FAs, CRs) were used as Test 2 targets. An additional 20 words that were encountered only during the Study task were included as targets, for a total of 140 Test 2 targets. 140 words not encountered during the Study task or Test 1 served as lures in Test 2. To ensure that the composition of Test 2 targets would be approximately equal for each participant, only a subset of the Hits (40 words), Misses (20 words), FAs (20 words), and CRs (40 words) from Test 1 were included as Test 2 targets. The specific words from with each response type were chosen at random. Like Experiment 1, if a participant did not have a sufficient number of trials for one response type, these items were replaced with another Test 1 item such that there were exactly 100 Test 1 & Study items (Hits and Misses) and exactly 100 Test 1 Only items (FAs and CRs) included in Test 2 for all participants.

Results and Discussion

Test 1 Performance. There was no group difference in d-prime (Memory Group: M=2.32, SD=0.10; Novelty Group: M=2.37, SD=0.12; $t(30) = -0.32$).

Test 2 Performance: The Role of Perceived Oldness and Task Goals. To assess the role of perceived oldness and task goals on retrieval-based re-encoding, we performed the same analysis as in Experiment 1: d-prime scores from the four conditions of interest were submitted to a between-subjects repeated-measure ANOVA with factors of repetitions,

perceived oldness, and group. Accuracy rates for these four conditions and for studied-only control items are shown for each group in **Figure 4.3**.

There was a main effect of perceived oldness ($F(1,30) = 26.705$, $p < 0.001$) and an effect of task goals (interaction between group and perceived oldness: $F(1,30) = 4.55$, $p = 0.041$) on Test 2 memory performance. This pattern of results is consistent with two co-occurring value effects contributing to retrieval-based re-encoding during the first retrieval test: a value signal associated with perceived memory evidence and a value signal associated with retrieval outcomes that are consistent with task goals. The reliable effect of task goals stands in contrast to Experiment 1 and is likely due to the manipulation of base rates during the first retrieval test: while participants in the Memory Group maintained memory retrieval as their task goal, the rarity of new items combined with the instructional manipulation to make the task goal of detecting novel items more salient for the Novelty Group.

Follow-up analyses constrained to each group separately revealed that the effect of perceived oldness was reliable for the Memory Group ($F(1,15) = 48.984$, $p < 0.001$) but not for the Novelty Group ($F(1,15) = 3.161$, $p = 0.096$), although there remained a qualitative benefit for perceived oldness in both groups, as shown in the bottom row of **Figure 4.3**. Thus, it appears that the value associated with perceived oldness has a larger impact on retrieval-based re-encoding than the value associated with task-congruent retrieval outcomes, at least given our task manipulations. We cannot rule out the possibility that additional manipulations which further enhanced the value associated with detecting novel items (e.g. by incentivizing "new" responses; or by providing a cover story which asserted that an individual's ability to detect novel events is predictive

of general intelligence; *cf*. Han et al., 2010) might result in a larger effect of task goals such that the Novelty Group showed enhanced memory for perceived newness.

There was no difference in overall memory between the two groups ($F(1,30) = 0.079$). This suggests that the effect of group observed in Experiment 1 was not due to the group manipulation but instead due simply to individual differences in overall memory between the participants assigned to each group.

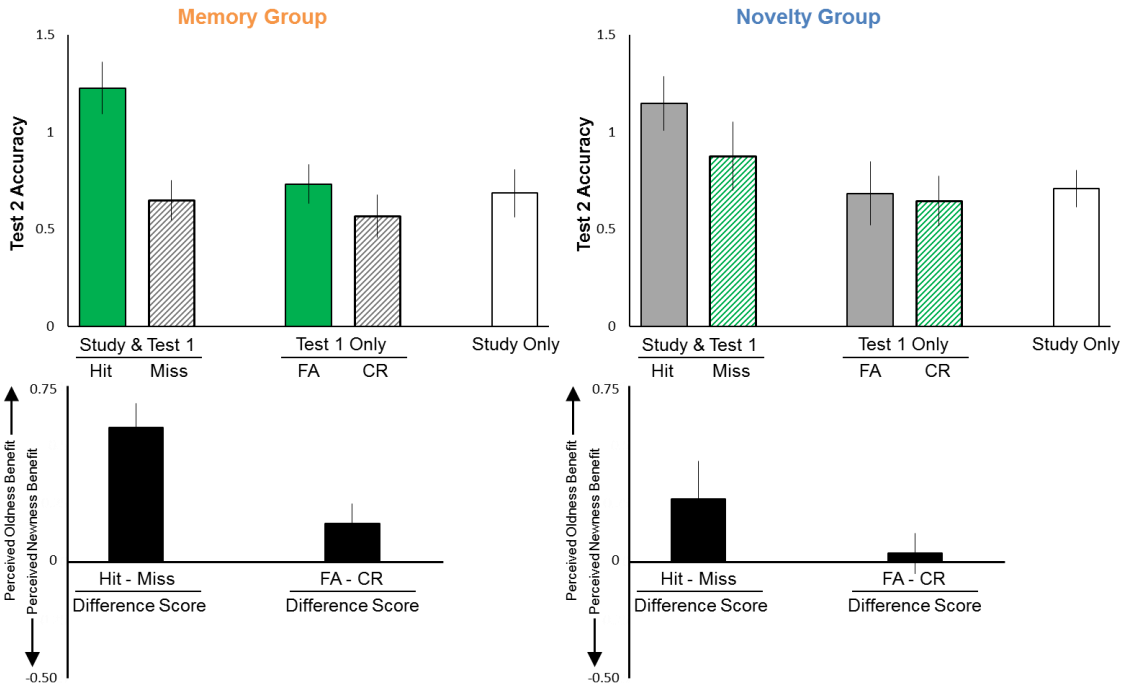

Figure 4.3. Accuracy for All Test 2 Targets in Experiment 2. The top row shows Test 2 accuracy for all Test 2 target items. Items are sorted by their Test 1 response type, as indicated by the horizontal axis labels. Filled bars indicate items perceived as old in Test 1; patterned bars indicate items perceived as new in Test 1; green indicates items that were goal-congruent in Test 1 (see text and Figure 4.1 for more details). Items encountered during the Study phase only are shown for comparison. The bottom row shows difference scores to summarize the retrieval-based re-encoding benefit of perceived oldness (plotted upward) or perceived newness (plotted downward). There was a main effect of perceived oldness (all difference scores are positive) but this effect was modulated by task goals (the positive difference scores in the Novelty Group are significantly reduced). This suggests that both perceived oldness and task goals impacted retrieval-based re-encoding during Test 1.

General Discussion

We hypothesized that one function of striatal activity typically observed during memory retrieval is the adaptive re-encoding of valued information (Scimeca & Badre, 2012). The current studies provide behavioral evidence consistent with this hypothesis: we found that both perceived oldness and goal-attainment support retrieval-based re-encoding. Previous neuroimaging studies have identified these two factors as potential sources of value-related activity in the nigra-striatal system during retrieval (Elward et al., 2014; Han et al., 2010). We interpret the re-encoding effects of the current studies as a behavioral consequence of these putative value signals. Thus, these behavioral results provide converging evidence consistent with multiple co-occurring striatal value signals during retrieval.

Using fMRI, Elward and colleagues (2014) identified distinct striatal signals associated with successful memory (due primarily to recollection) and external reward (due to monetary incentive). Here we identify distinct re-encoding effects associated with perceived oldness (likely due to a combination of familiarity and recollection) and goal-attainment (due to task goals). Together, these findings point to an emerging consensus that striatal signals during retrieval arise from multiple sources of value. However, future work that combines more elaborated behavioral paradigms and neuroimaging will be required to determine which manipulations tap into the same underlying construct, and which constructs are represented in overlapping or distinct neural populations within the nigra-striatal system. For example, the monetary incentives used in previous studies (Elward et al., 2014; Han et al., 2010) represent an exogenous reward, whereas the manipulation of task goals in the current studies represents a more endogenous source of

motivational reward. This distinction between exogenous and endogenous rewards may result in recruitment of distinct neural substrates and thus could exert an influence on re-encoding through distinct neural pathways.

During initial encoding, value signals have been shown to support enhanced subsequent memory (Adcock et al., 2006; Murty & Adcock, 2013; Shohamy & Adcock, 2010; Wittmann et al., 2005), and the neurobiological mechanisms that link value to enhanced memory during initial encoding provide a set of candidate mechanisms for retrieval-based re-encoding. One possibility is that interactions between the nigra-striatal system and hippocampus directly support enhanced memory (Adcock et al., 2006; Shohamy & Adcock, 2010; Wolosin, Zeithamova, & Preston, 2012), potentially through dopamine-mediated long-term potentiation or synaptic tagging (Dunsmoor, Murty, Davachi, & Phelps, 2015; Huang & Kandel, 1995; Jay, 2003; Lemon & Manahan-Vaughan, 2006; Li et al., 2003). Another possibility is that nigra-striatal value signals modulate a broad cortical network, such as frontal and visual cortices, to enhance processing of value-relevant information and thus indirectly support hippocampal-dependent encoding (Murty & Adcock, 2013). Given the widespread projections of the dopaminergic value system (Cools, 2011; Fallon & Cools, 2014; Lisman & Grace, 2005; Shohamy & Adcock, 2010), value signals during encoding and retrieval likely exert their downstream effects through a variety of pathways.

In the current studies, we used recognition tests because this afforded us a tightly controlled paradigm in which to demonstrate the initial behavioral effect of retrieval-based re-encoding. Specifically, we used recognition probes from an initial recognition test as probes on a final recognition test. However, it is possible that value signals during

retrieval support re-encoding more broadly: for example, perceived oldness (or goal-attainment) associated with a probe on the initial test may also enhance encoding of test-irrelevant information that is presented simultaneously (such as a novel paired-associate or a source detail). This broader effect has been found in studies of reward-enhanced memory during initial encoding (e.g. Murty & Adcock, 2013). Furthermore, item recognition provides a relatively constrained retrieval environment, relative to more effortful memory tests such as source tasks or free recall. Given that the current studies establish a basic link between value and retrieval-based re-encoding, future work can seek to delineate the scope and generality of value-based re-encoding across different forms of retrieval.

Because the current studies used recognition, one alternative explanation of the perceived oldness effect is the effect of retrieval fluency. Fluency describes the ease with which information is processed or accessed, and this fluency can influence memory decisions and metacognitive judgments (e.g. Benjamin, Bjork, & Schwartz, 1998). Although all words were randomized across participants, it is possible that idiosyncratic patterns of participant-specific fluency drove "old" responses during both Test 1 and Test 2. For example, certain words may be processed more fluently by a participant for idiosyncratic reasons, and this fluency could make the participant more likely to respond "old" to these words during Test 1. During Test 2, the participant would again be more likely to respond "old" to these items. However, this effect may not be due to the re-encoding benefit of perceived oldness during Test 1 but instead is again due to the high levels of fluency the participant experiences for these specific words. Critically, this fluency explanation cannot account for the effect of task goals that we observe in

Experiment 2 (reflected in the statistical interaction between perceived oldness and group). Therefore, while we cannot rule out a possible contribution from retrieval fluency, this alternative explanation cannot account for the full pattern of observed results.

Finally, the retrieval-based re-encoding results described here provide one potential mechanism for commonly observed testing effects, although it is almost certainly the case that other cognitive mechanisms also play a role (Roediger & Butler, 2011). The present results also have potentially important implications for testing that is used in educational settings and for the implementation of memory rehabilitation programs (Middleton & Schwartz, 2012), highlighting the importance of considering motivation, goals, and value in these contexts.

CHAPTER 5

Conclusions

The central thesis of this dissertation is that the nigra-striatal dopaminergic value system makes important contributions to declarative memory retrieval, in support of the flexible and adaptive nature of memory. Intelligent and flexible behavior requires that we bring to bear information from past experience that is currently relevant in order to guide future thought and action. It is apparent from everyday life, and from over a century of psychological research into memory processes (Benjamin, 2007; James, 1890; Schacter et al., 1978), that healthy individuals can readily exert this form of strategic access to memory. Research in recent decades has focused on the important role of a frontal-parietal brain network in implementing this control (Badre & Wagner, 2007; Moscovitch, 1992, 1994; Thompson-Schill, D'Esposito, Aguirre, & Farah, 1997; Wagner et al., 2005; Wheeler, Stuss, & Tulving, 1997). This dissertation provides foundational empirical evidence that striatum, likely in concert with frontally-mediated control processes, is also critical to flexible control of declarative memory retrieval.

As reviewed in Chapter 1, the existing neuroimaging and neuropsychology literatures are consistent with a role of striatum in retrieval. Although striatal activity is a consistent finding in fMRI studies of memory (Spaniol et al., 2009), the striatum has rarely been the primary focus of these studies. However, our review identified a pattern in that striatal activation was often observed in studies that placed control demands on retrieval. The specific patterns of memory deficits observed in patients with dysfunction of the nigra-striatal system, such as Parkinson's disease and Huntington's disease

patients, provided converging evidence and further implies a causal link between striatum and effective control of retrieval. Importantly, we note that controlled processing is involved not only in the processes involved in accessing information from long-term memory, but also in the processes that monitor and evaluate the content retrieved from memory (Benjamin, 2007).

Building from this review, we leveraged existing theories of striatal function that have emerged from studies of value-based learning, decision-making, and action selection (Cools, 2011; Daw & Tobler, 2014; Doll & Frank, 2009; Haber & Knutson, 2010; Shohamy & Adcock, 2010) to propose three hypotheses regarding the specific contributions of striatum to declarative memory: (1) striatum marks valued retrieval outcomes and supports adaptive re-encoding of retrieved information; (2) striatum gates into working memory a goal or cue that biases memory access towards relevant information; (3) striatum supports reinforcement learning of control policies that operate during retrieval. The empirical work contained in this dissertation was designed to specifically address the first and third hypotheses.

The series of studies discussed in Chapters 2 and 3 addressed the role of striatally-mediated reinforcement learning processes during recognition memory decisions. In recognition memory, an individual must evaluate memory evidence in order to decide whether the level of evidence indicates that a stimulus has been previously encountered. In this sense, a decision policy operates on the contents retrieved from memory to determine the most appropriate decision ("old" or "new"). This decision policy is summarized by the signal detection theory measure of the *decision criterion* (Macmillan & Creelman, 2005; Wickens, 2002). We manipulated the degree to which feedback

deviated from an individual's expectations by targeting false feedback to either high or low confidence responses, with the intent of manipulating the magnitude of the experienced Prediction Errors. This manipulation impacted the degree to which individuals adjusted their decision criterion, implicating PE-based learning in the regulation of memory decisions.

We combined this behavioral paradigm with fMRI to investigate the neural substrates of PE-based learning in recognition memory. Reinforcement learning has been linked to the nigra-striatal dopaminergic system by converging evidence from neuroimaging, neurophysiology, computational modeling, patient studies, and genetics studies (Doll & Frank, 2009; Frank et al., 2007, 2004; Haber & Knutson, 2010; Schultz, 1998) and the striatum in particular is known to track reward PE signals (Badre & Frank, 2012; Gläscher et al., 2010). We found that striatum tracked PEs during recognition memory and that the magnitude of this neural signal predicted individual differences in adjustments to criterion (mirroring similar findings from the value-based learning literature, Schonberg, Daw, Joel, & O'Doherty, 2007). The combination of behavioral and neuroimaging results provide strong converging evidence that a striatally-mediated RL system regulates control processes in declarative memory retrieval. This fills a fundamental gap in theories of cognitive control of memory by implicating a concrete mechanism and well-characterized neural system as the means by which healthy individuals can dynamically evaluate and adjust mnemonic control processes.

Critically, in both the behavioral and neuroimaging results, we found that PEs in recognition memory are scaled relative to a value signal that is linked to the memory strength underlying the memory decision, as opposed to a value signal tied to response

confidence. In making our initial predictions about the role of PEs in recognition memory, we turned to analogous experiments conducted in standard RL tasks. RL principles have been extended to a variety of different experimental domains that fall under the umbrella of value-based learning and decision-making. However, this particular result is not easily predicted from the existing RL literature. Instead, this finding serves to highlight the importance of explicitly testing RL principles in other domains, like declarative memory.

The underlying reason why PE signals are scaled relative to memory strength remains unclear and will require additional experimental work. One salient difference between the current tasks and many RL tasks is that we required participants to determine the appropriate decision for a given level of memory evidence, which is an abstract latent state. This stands in contrast to the more concrete stimuli that often provide the learning context in typical RL tasks. The expected value of memory decisions may be linked to the level of memory evidence because the current task required control processes to operate on the latent memory signal.

An alternative explanation is that individuals may have learned, through a lifetime of experience, that higher levels of memory strength are typically associated with positive outcomes and thus should be inherently valued. A related possibility is that the specific framing of the recognition test is the source of the link between memory evidence and value. That is, because the task was framed as a "memory test" in which the implicit goal was to remember information from long-term memory, memory evidence may have been ascribed motivational value. As noted in Chapter 1, striatum is responsive not only to item familiarity but also to novelty, depending on goals and context. It may be that the

relationship between memory evidence and value would be reversed if the experimental task were framed as a novelty-detection task (that is, signals of novelty or low levels of memory evidence would be associated with high expected value).

The experiments described in Chapter 4 begin to address these questions while simultaneously testing the first hypothesis regarding striatal contributions to memory retrieval: that striatum supports re-encoding of retrieved information. By framing a retrieval task as either a "memory test" or a "novelty-detection test", we can begin to determine whether the striatal valuation system assigns value to task-independent evidence of oldness or to retrieval outcomes consistent with task-goals (or some combination of both). By analogy to effects typically observed during initial encoding (Shohamy & Adcock, 2010), the retrieval-based re-encoding hypothesis proposes that value associated with retrieval outcomes during retrieval should result in enhanced memory on subsequent memory tests. Using a modified testing effect procedure, we thus measured behavioral performance on a final memory test to infer the value of retrieval outcomes during an initial test. Consistent with the hypothesis, we found evidence that both task-independent perceived evidence of oldness and goal-congruent retrieval outcomes supported retrieval-based re-encoding.

Whereas the current studies used the behavioral consequences of value as the primary outcome measure, the findings make straightforward predictions regarding the underlying neural substrates: both perceived oldness and goal-attainment should result in striatal activity during retrieval, which in turn should be predictive of subsequent memory on a final test. However, future neuroimaging work can provide additional insights into the specific nature of retrieval-related striatal value signals and their mechanism of action

on enhancing subsequent memory. Perceived oldness and goal-attainment may tap into overlapping or distinct neural populations in striatum (Elward et al., 2014). Likewise, these factors may modulate retrieval-based re-encoding through overlapping or distinct interactions with other brain structures like frontal cortex and hippocampus (Murty & Adcock, 2013).

In this dissertation, I used fMRI to measure neural activity linking our experimental manipulations to changes in behavior. This approach identified the striatum as the source of mnemonic PEs. However, fMRI is an observational method, and demonstrating a causal role for the nigra-striatal dopaminergic system in regulating retrieval control policies will require interventional approaches such as pharmacological manipulations or studies of patient populations. For example, the present framework predicts that patients with nigra-striatal dysfunction, such as Parkinson's or Huntington's disease patients, should show marked deficits in their ability regulate recognition decisions based on feedback. This deficit in the patients' ability to dynamic adjust memory control processes should be present despite the patients' generally typical performance in recognition tasks. Intriguingly, this account also predicts that patients who show decreased overall memory performance but relatively spared metacognitive accuracy (which has been observed in e.g. Alzheimer's disease patients; Gallo, Cramer, Wong, & Bennett, 2012) should be able to use feedback-based cues to regulate recognition decisions, so long as their nigra-striatal system is intact. Future work with patient populations will provide critical insight into effective memory training and rehabilitation approaches (Dunlosky, Rawson, Marsh, Nathan, & Willingham, 2013;

Lustig, Shah, Seidler, & Reuter-Lorenz, 2009; McDaniel & Bugg, 2012; Middleton & Schwartz, 2012).

Finally, this research program describes an important link between declarative and nondeclarative memory systems, often studied in isolation. Striatum and associated structures have been linked to a wide-range of psychological processes from Pavlovian conditioning to abstract rule-learning and working memory management (Badre & Frank, 2012; Chatham, Frank, & Badre, 2014; O'Doherty et al., 2003). Here we demonstrate an important role for this system in episodic memory. This informs a more nuanced understanding of the relationship between memory systems in the brain – a modern model that is replacing the long-dominant taxonomy of memory systems that posited rigidly demarcated declarative and nondeclarative systems (e.g. Scoville & Milner, 1957; Squire & Wixted, 2011; see also Cabeza & Moscovitch, 2013; Dew & Cabeza, 2011; Dudai & Morris, 2013). Research that provides a better understanding of how these various brain systems interact during the control of memory will advance our understanding of both cognitive control and memory systems, in general.

Bibliography

Abe, N., Okuda, J., Suzuki, M., Sasaki, H., Matsuda, T., Mori, E., … Fujii, T. (2008). Neural correlates of true memory, false memory, and deception. *Cerebral Cortex (New York, N.Y. : 1991)*, *18*(12), 2811–9. doi:10.1093/cercor/bhn037

Adcock, R. A., Thangavel, A., Whitfield-Gabrieli, S., Knutson, B., & Gabrieli, J. D. E. (2006). Reward-motivated learning: mesolimbic activation precedes memory formation. *Neuron*, *50*(3), 507–17. doi:10.1016/j.neuron.2006.03.036

Alexander, G. E., DeLong, M. R., & Strick, P. L. (1986). Parallel organiztion of functionally segregated circuits linking basal ganglia and cortex. *Annual Review of Neuroscience*, *9*, 357–381.

Aly, M., Yonelinas, A. P., Kishiyama, M. M., & Knight, R. T. (2011). Damage to the lateral prefrontal cortex impairs familiarity but not recollection. *Behavioural Brain Research*, *225*(1), 297–304. doi:10.1016/j.bbr.2011.07.043

Anderson, J. R., Fincham, J. M., Qin, Y., & Stocco, A. (2008). A central circuit of the mind. *Trends in Cognitive Sciences*, *12*(4), 136–43. doi:10.1016/j.tics.2008.01.006

Anderson, J. R., & Milson, R. (1989). Human memory: An adaptive perspective. *Psychological Review*, *96*(4), 703–719.

Anderson, M. C., Ochsner, K. N., Kuhl, B., Cooper, J., Robertson, E., Gabrieli, S. W., … Gabrieli, J. D. E. (2004). Neural systems underlying the suppression of unwanted memories. *Science (New York, N.Y.)*, *303*(5655), 232–5. doi:10.1126/science.1089504

Ashby, F. G., & Maddox, W. T. (2005). Human category learning. *Annual Review of Psychology*, *56*, 149–78. doi:10.1146/annurev.psych.56.091103.070217

Badre, D., Doll, B. B., Long, N. M., & Frank, M. J. (2012). Rostrolateral Prefrontal Cortex and Individual Differences in Uncertainty-Driven Exploration. *Neuron*, *73*(3), 595–607. doi:10.1016/j.neuron.2011.12.025

Badre, D., & Frank, M. J. (2012). Mechanisms of hierarchical reinforcement learning in cortico-striatal circuits 2: evidence from FMRI. *Cerebral Cortex*, *22*(3), 527–536. doi:10.1093/cercor/bhr117

Badre, D., Lebrecht, S., Pagliaccio, D., Long, N. M., & Scimeca, J. M. (2014). Ventral Striatum and the Evaluation of Memory Retrieval Strategies. *Journal of Cognitive Neuroscience*, *26*(9), 1928–1948. doi:10.1162/jocn

Badre, D., Poldrack, R. A., Paré-Blagoev, E. J., Insler, R. Z., & Wagner, A. D. (2005). Dissociable Controlled Retrieval and Generalized Selection Mechanisms in Ventrolateral Prefrontal Cortex. *Neuron*, *47*(6), 907–918. doi:10.1016/j.neuron.2005.07.023

Badre, D., & Wagner, A. D. (2007). Left ventrolateral prefrontal cortex and the cognitive control of memory. *Neuropsychologia*, *45*(13), 2883–901. doi:10.1016/j.neuropsychologia.2007.06.015

Banks, W. P. (1970). Signal Detection Theory and Human Memory. *Psychological Bulletin*, *74*, 81–99.

Barnes, J., Boubert, L., Harris, J., Lee, A., & David, A. S. (2003). Reality monitoring and visual hallucinations in Parkinson ' s disease. *Neuropsychology*, *41*, 565–574.

Barredo, J., Oztekin, I., & Badre, D. (2013). Ventral Fronto-Temporal Pathway Supporting Cognitive Control of Episodic Memory Retrieval. *Cerebral Cortex*. doi:10.1093/cercor/bht291

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*(6), 821–32. doi:10.1162/089892903322370744

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* (Volume 48., Vol. 48, pp. 175–223). Academic Press. doi:10.1016/S0079-7421(07)48005-7

Benjamin, A. S., Bjork, R. a, & Schwartz, B. L. (1998). The mismeasure of memory: when retrieval fluency is misleading as a metamnemonic index. *Journal of Experimental Psychology. General*, *127*(1), 55–68. doi:10.1037/0096-3445.127.1.55

Berridge, K. C. (2007). The debate over dopamine's role in reward: the case for incentive salience. *Psychopharmacology*, *191*(3), 391–431. doi:10.1007/s00213-006-0578-x

Bornstein, A. M., & Daw, N. D. (2011). Multiplicity of control in the basal ganglia: computational roles of striatal subregions. *Current Opinion in Neurobiology*, *21*(3), 374–380. doi:10.1016/j.conb.2011.02.009

Botvinick, M. M., Niv, Y., & Barto, A. C. (2009). Hierarchically organized behavior and its neural foundations: a reinforcement learning perspective. *Cognition*, *113*(3), 262–80. doi:10.1016/j.cognition.2008.08.011

Braver, T. S., Barch, D. M., Gray, J. R., Molfese, D. L., & Snyder, A. (2001). Anterior Cingulate Cortex and Response Conflict: Effects of Frequency, Inhibition and Errors. *Cerebral Cortex*, *11*(9), 825–836. doi:10.1093/cercor/11.9.825

Braver, T. S., & Cohen, J. D. (2000). On the Control of Control : The Role of Dopamine in Regulating Prefrontal Function and Working Memory.

Brett, M., Anton, J.-L., Valabreague, R., & Poline, J.-B. (2002). Region of interest analysis using an SPM toolbox. In *NeuroImage* (Vol. 16, p. S497).

Brønnick, K., Alves, G., Aarsland, D., Tysnes, O.-B., & Larsen, J. P. (2011). Verbal memory in drug-naive, newly diagnosed Parkinson's disease. The retrieval deficit hypothesis revisited. *Neuropsychology*, *25*(1), 114–24. doi:10.1037/a0020857

Brown, G. S., & White, K. G. (2009). Reinforcer probability, reinforcer magnitude, and the reinforcement context for remembering. *Journal of Experimental Psychology. Animal Behavior Processes*, *35*(2), 238–249. doi:10.1037/a0013864

Brown, J. W., Bullock, D., & Grossberg, S. (2004). How laminar frontal cortex and basal ganglia circuits interact to control planned and reactive saccades. *Neural Networks*, *17*(4), 471–510. doi:10.1016/j.neunet.2003.08.006

Brown, P. C., Roediger, H. L., & McDaniel, M. A. (2014). *Make It Stick*. Harvard University Press.

Brown, S., Steyvers, M., & Hemmer, P. (2007). Modeling experimentally induced strategy shifts. *Psychological Science*, *18*(1), 40–5. doi:10.1111/j.1467-9280.2007.01846.x

Bruno, D., Higham, P. a, & Perfect, T. J. (2009). Global subjective memorability and the strength-based mirror effect in recognition memory. *Memory & Cognition*, *37*(6), 807–18. doi:10.3758/MC.37.6.807

Buckner, R. L. (1996). Beyond HERA: Contributions of specific prefrontal brain areas to long-term memory retrieval. *Psychonomic Bulletin & Review*, *3*(2), 149–158. doi:10.3758/BF03212413

Buckner, R. L., Koutstaal, W., Schacter, D. L., Wagner, a D., & Rosen, B. R. (1998). Functional-anatomic study of episodic retrieval using fMRI. I. Retrieval effort versus retrieval success. *NeuroImage*, *7*(3), 151–62. doi:10.1006/nimg.1998.0327

Bunzeck, N., Dayan, P., Dolan, R. J., & Duzel, E. (2010). A common mechanism for adaptive scaling of reward and novelty. *Human Brain Mapping*, *31*(9), 1380–94. doi:10.1002/hbm.20939

Bunzeck, N., & Düzel, E. (2006). Absolute coding of stimulus novelty in the human substantia nigra/VTA. *Neuron*, *51*(3), 369–79. doi:10.1016/j.neuron.2006.06.021

Burrell, Q. (1980). A simple stochastic model for library loans. *Journal of Documentation*, *36*(2), 115–132.

Cabeza, R., & Moscovitch, M. (2013). Memory Systems, Processing Modes, and Components: Functional Neuroimaging Evidence. *Perspectives on Psychological Science*, *8*(1), 49–55. doi:10.1177/1745691612469033

Cary, M. (2003). A dual-process account of the list-length and strength-based mirror effects in recognition. *Journal of Memory and Language*, *49*(2), 231–248. doi:10.1016/S0749-596X(03)00061-5

Chatham, C. H., Frank, M. J., & Badre, D. (2014). Corticostriatal Output Gating during Selection from Working Memory. *Neuron*, *81*(4), 930–942. doi:10.1016/j.neuron.2014.01.002

Chua, E. F., Schacter, D. L., Rand-Giovannetti, E., & Sperling, R. a. (2006). Understanding metamemory: neural correlates of the cognitive process and subjective level of confidence in recognition memory. *NeuroImage*, *29*(4), 1150–60. doi:10.1016/j.neuroimage.2005.09.058

Clark, S. E., & Gronlund, S. D. (1996). Global matching models of recognition memory: How the models match the data. *Psychonomic Bulletin & Review*, *3*(1), 37–60. doi:10.3758/BF03210740

Cohen, M. X., & Frank, M. J. (2009). Neurocomputational models of basal ganglia function in learning, memory and choice. *Behavioural Brain Research*, *199*(1), 141–56. doi:10.1016/j.bbr.2008.09.029

Cohen, N. J., Poldrack, R. A., & Eichenbaum, H. (1997). Memory for items and memory for relations in the procedural/declarative memory framework. *Memory*, *5*, 131–178.

Cohn, M., Moscovitch, M., & Davidson, P. S. R. (2010). Double dissociation between familiarity and recollection in Parkinson's disease as a function of encoding tasks. *Neuropsychologia*, *48*(14), 4142–7. doi:10.1016/j.neuropsychologia.2010.10.013

Cools, R. (2011). Dopaminergic control of the striatum for high-level cognition. *Current Opinion in Neurobiology*, *21*(3), 402–407. doi:10.1016/j.conb.2011.04.002

Cools, R., Clark, L., & Robbins, T. W. (2004). Differential responses in human striatum and prefrontal cortex to changes in object and rule relevance. *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience*, *24*(5), 1129–35. doi:10.1523/JNEUROSCI.4312-03.2004

Cools, R., Ivry, R. B., & D'Esposito, M. (2006). The human striatum is necessary for responding to changes in stimulus relevance. *Journal of Cognitive Neuroscience*, *18*(12), 1973–1983. doi:10.1162/jocn.2006.18.12.1973

Corkin, S. (1968). Acquisition of motor skill after bilateral medial temporal-lobe excision. *Neuropsychologia*, *6*(3), 255–265. doi:10.1016/0028-3932(68)90024-9

Cox, J. C., & Dobbins, I. G. (2011). The striking similarities between standard, distractor-free, and target-free recognition. *Memory & Cognition*, *39*(6), 925–40. doi:10.3758/s13421-011-0090-3

Crescentini, C., Marin, D., Del Missier, F., Biasutti, E., & Shallice, T. (2011). Interference from retrieval cues in parkinson's disease. *Neuropsychology*. doi:10.1037/a0024674

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*(2), 434–47. doi:10.1016/j.neuropsychologia.2007.08.021

Criss, A. H. (2010). Differentiation and response bias in episodic memory: evidence from reaction time distributions. *Journal of Experimental Psychology. Learning, Memory, and Cognition*, *36*(2), 484–99. doi:10.1037/a0018435

Curran, T., DeBuse, C., & Leynes, P. A. (2007). Conflict and criterion setting in recognition memory. *Journal of Experimental Psychology. Learning, Memory, and Cognition*, *33*(1), 2–17. doi:10.1037/0278-7393.33.1.2

Dahlin, E., Neely, A. S., Larsson, A., Bäckman, L., & Nyberg, L. (2008). Transfer of learning after updating training mediated by the striatum. *Science (New York, N.Y.)*, *320*(5882), 1510–2. doi:10.1126/science.1155466

Davidson, P. S. R., Anaki, D., Saint-Cyr, J. a, Chow, T. W., & Moscovitch, M. (2006). Exploring the recognition memory deficit in Parkinson's disease: estimates of recollection versus familiarity. *Brain : A Journal of Neurology*, *129*(Pt 7), 1768–79. doi:10.1093/brain/awl115

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*(6), 1204–1215. doi:10.1016/j.neuron.2011.02.027

Daw, N. D., O'Doherty, J. P., Dayan, P., Seymour, B., & Dolan, R. J. (2006). Cortical substrates for exploratory decisions in humans. *Nature*, *441*(7095), 876–9. doi:10.1038/nature04766

Daw, N. D., & Tobler, P. N. (2014). Value Learning through Reinforcement : The Basics of Dopamine and Reinforcement Learning. In P. W. Glimcher & E. Fehr (Eds.), *Neuroeconomics* (Second Edi., pp. 283–298). Academic Press.

De Martino, B., Fleming, S. M., Garrett, N., & Dolan, R. J. (2013). Confidence in value-based choice. *Nature Neuroscience*, *16*(1), 105–10. doi:10.1038/nn.3279

Delis, D. C., Massman, P. J., Butters, N., & Salmon, D. P. (1987). *California Verbal Learning Test*. New York: Psychological Corporation.

Dennis, S., & Humphreys, M. S. (2001). A context noise model of episodic word recognition. *Psychological Review*, *108*(2), 452–478. doi:10.1037//0033-295X.108.2.452

Desimone, R., & Duncan, J. (1995). Neural mechanisms of selective visual attention. *Annual Review of Neuroscience*, *18*, 193–222.

Dew, I. T. Z., & Cabeza, R. (2011). The porous boundaries between explicit and implicit memory: Behavioral and neural evidence. *Annals of the New York Academy of Sciences*, *1224*(1), 174–190. doi:10.1111/j.1749-6632.2010.05946.x

Dobbins, I. G., Simons, J. S., & Schacter, D. L. (2004). fMRI Evidence for Separable and Lateralized Prefrontal Memory Monitoring Processes. *Journal of Cognitive Neuroscience*, (16), 908–920.

Dobbins, I. G., & Wagner, A. D. (2005). Domain-general and domain-sensitive prefrontal mechanisms for recollecting events and detecting novelty. *Cerebral Cortex (New York, N.Y. : 1991)*, *15*(11), 1768–78. doi:10.1093/cercor/bhi054

Doll, B. D., & Frank, M. J. (2009). The basal ganglia in reward and decision making : computational models and empirical studies, 399–425.

Donaldson, D. I., Petersen, S. E., Buckner, R. L., & Louis, S. (2001). fMRI : Evidence that Priming Does Not Support Recognition Memory, *31*, 1047–1059.

Drag, L. L., Bieliauskas, L. a, Kaszniak, A. W., Bohnen, N. I., & Glisky, E. L. (2009). Source memory and frontal functioning in Parkinson's disease. *Journal of the International Neuropsychological Society : JINS*, *15*(3), 399–406. doi:10.1017/S1355617709090572

Dudai, Y., & Morris, R. G. M. (2013). Memorable Trends. *Neuron*, *80*(3), 742–750. doi:10.1016/j.neuron.2013.09.039

Dunlosky, J., Rawson, K. A., Marsh, E. J., Nathan, M. J., & Willingham, D. T. (2013). Improving Students' Learning With Effective Learning Techniques: Promising Directions From Cognitive and Educational Psychology. *Psychological Science in the Public Interest*, *14*(1), 4–58. doi:10.1177/1529100612453266

Dunsmoor, J. E., Murty, V. P., Davachi, L., & Phelps, E. a. (2015). Emotional learning selectively and retroactively strengthens memories for related events. *Nature*. doi:10.1038/nature14106

Edelstyn, N. M. J., Mayes, A. R., Condon, L., Tunnicliffe, M., & Ellis, S. J. (2007). Recognition, recollection, familiarity and executive function in medicated patients with moderate Parkinson's disease. *Journal of Neuropsychology*, *1*(Pt 2), 131–147. doi:10.1348/174866407X182565

Edelstyn, N. M. J., Shepherd, T. a, Mayes, A. R., Sherman, S. M., & Ellis, S. J. (2010). Effect of disease severity and dopaminergic medication on recollection and familiarity in patients with idiopathic nondementing Parkinson's. *Neuropsychologia*, *48*(5), 1367–75. doi:10.1016/j.neuropsychologia.2009.12.039

Eichenbaum, H., Yonelinas, A. P., & Ranganath, C. (2007). The medial temporal lobe and recognition memory. *Annual Review of Neuroscience*, *30*, 123–52. doi:10.1146/annurev.neuro.30.051606.094328

Elward, R. L., Vilberg, K. L., & Rugg, M. D. (2014). Motivated Memories: Effects of Reward and Recollection in the Core Recollection Network and Beyond. *Cerebral Cortex (New York, N.Y. : 1991)*. doi:10.1093/cercor/bhu109

Fallon, S. J., & Cools, R. (2014). Reward Acts on the pFC to Enhance Distractor REsistance of Working Memory Representations. *Journal of Cognitive Neuroscience*, *In Press*, 1–15. doi:10.1162/jocn

Frank, M. J., & Badre, D. (2012). Mechanisms of hierarchical reinforcement learning in corticostriatal circuits 1: computational analysis. *Cerebral Cortex (New York, N.Y. : 1991)*, *22*(3), 509–526. doi:10.1093/cercor/bhr114

Frank, M. J., Doll, B. B., Oas-Terpstra, J., & Moreno, F. (2009). Prefrontal and striatal dopaminergic genes predict individual differences in exploration and exploitation. *Nature Neuroscience*, *12*(8), 1062–1068. doi:10.1038/nn0510-649a

Frank, M. J., & Fossella, J. a. (2010). Neurogenetics and Pharmacology of Learning, Motivation, and Cognition. *Neuropsychopharmacology*, *36*(1), 133–152. doi:10.1038/npp.2010.96

Frank, M. J., Moustafa, A. a, Haughey, H. M., Curran, T., & Hutchison, K. E. (2007). Genetic triple dissociation reveals multiple roles for dopamine in reinforcement learning. *Proceedings of the National Academy of Sciences of the United States of America*, *104*(41), 16311–6. doi:10.1073/pnas.0706111104

Frank, M. J., & O'Reilly, R. C. (2006). A mechanistic account of striatal dopamine function in human cognition: psychopharmacological studies with cabergoline and haloperidol. *Behavioral Neuroscience*, *120*(3), 497–517. doi:10.1037/0735-7044.120.3.497

Frank, M. J., Seeberger, L. C., & O'Reilly, R. C. (2004). By carrot or by stick: cognitive reinforcement learning in parkinsonism. *Science (New York, N.Y.)*, *306*(5703), 1940–3. doi:10.1126/science.1102941

Gallo, D. A. (2013). Retrieval Expectations Affect False Recollection: Insights From a Criterial Recollection Task. *Current Directions in Psychological Science*, *22*(4), 316–323. doi:10.1177/0963721413481472

Gallo, D. A., Cramer, S. J., Wong, J. T., & Bennett, D. a. (2012). Alzheimer's disease can spare local metacognition despite global anosognosia: revisiting the confidence-accuracy relationship in episodic memory. *Neuropsychologia*, *50*(9), 2356–64. doi:10.1016/j.neuropsychologia.2012.06.005

Gallo, D. A., McDonough, I. M., & Scimeca, J. (2010). Dissociating source memory decisions in the prefrontal cortex: fMRI of diagnostic and disqualifying monitoring. *Journal of Cognitive Neuroscience*, *22*(5), 955–969. doi:10.1162/jocn.2009.21263

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., & Shimamura, A. (1995). Impaired use of organizational strategies in free recall following frontal lobe damage. *Neuropsychologia*, *13*, 1305–1333.

Glanzer, M., & Adams, J. K. (1990). The mirror effect in recognition memory: data and theory. *Journal of Experimental Psychology. Learning, Memory, and Cognition*, *16*(1), 5–16.

Glanzer, M., Hilford, A., & Maloney, L. T. (2009). Likelihood ratio decisions in memory: three implied regularities. *Psychonomic Bulletin & Review*, *16*(3), 431–455. doi:10.3758/PBR.16.3.431

Gläscher, 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*(4), 585–595. doi:10.1016/j.neuron.2010.04.016

Gluck, M. A., Shohamy, D., & Myers, C. (2002). How do People Solve the " Weather Prediction " Task ?: Individual Variability in Strategies for Probabilistic Category Learning. *Learning & Memory*, *9*, 408–418. doi:10.1101/lm.45202.4

Gorski, N. a., & Laird, J. E. (2011). Learning to use episodic memory. *Cognitive Systems Research*, *12*(2), 144–153. doi:10.1016/j.cogsys.2010.08.001

Graybiel, A. M., & Mink, J. (2009). The Basal Ganglie and Cognition. In M. S. Gazzaniga (Ed.), *The Cognitive Neurosciences* (pp. 565–585). MIT Press.

Green, D. M., & Swets, J. A. (1966). *Signal detection theory and psychophysics*. New York: Wiley.

Griffiths, T. L., Steyvers, M., & Firl, A. (2007). Google and the Mind: Predicting Fluency with PageRank. *Psychological Science*, *18*, 1069–1076. doi:10.1111/j.1467-9280.2007.02027.x

Guitart-Masip, M., Bunzeck, N., Stephan, K. E., Dolan, R. J., & Düzel, E. (2010). Contextual novelty changes reward representations in the striatum. *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience*, *30*(5), 1721–6. doi:10.1523/JNEUROSCI.5331-09.2010

Gurney, K., Prescott, T. J., & Redgrave, P. (2001). A computational model of action selection in the basal ganglia. I. A new functional anatomy. *Biological Cybernetics*, *84*(6), 401–10.

Haber, S. N., & Knutson, B. (2010). The reward circuit: linking primate anatomy and human imaging. *Neuropsychopharmacology : Official Publication of the American College of Neuropsychopharmacology*, *35*(1), 4–26. doi:10.1038/npp.2009.129

Han, S., & Dobbins, I. G. (2008). Examining recognition criterion rigidity during testing using a biased-feedback technique: Evidence for adaptive criterion learning. *Memory & Cognition*, *36*(4), 703–715. doi:10.3758/MC.36.4.703

Han, S., & Dobbins, I. G. (2009). Regulating recognition decisions through incremental reinforcement learning. *Psychonomic Bulletin & Review*, *16*(3), 469–74. doi:10.3758/PBR.16.3.469

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? *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience*, *30*(13), 4767–75. doi:10.1523/JNEUROSCI.3077-09.2010

Han, S., O'Connor, A. R., Eslick, A. N., & Dobbins, I. G. (2012). The Role of Left Ventrolateral Prefrontal Cortex during Episodic Decisions: Semantic Elaboration or Resolution of Episodic Interference? *Journal of Cognitive Neuroscience*, *24*(1), 223–234. doi:10.1162/jocn_a_00133

Hayama, H. R., & Rugg, M. D. (2009). Right dorsolateral prefrontal cortex is engaged during post-retrieval processing of both episodic and semantic information. *Neuropsychologia*, *47*(12), 2409–16. doi:10.1016/j.neuropsychologia.2009.04.010

Hazy, T. E., Frank, M. J., & O'Reilly, R. C. (2006). Banishing the homunculus: Making working memory work. *Neuroscience*, *139*(1), 105–118. doi:10.1016/j.neuroscience.2005.04.067

Helkala, E. L., Laulumaa, V., Soininen, H., & Riekkinen, P. J. (1989). Different error pattern of episodic and semantic memory in Alzheimer's disease and Parkinson's disease with dementia. *Neuropsychologia*, *27*(10), 1241–1248.

Henson, R. N. (2006). Forward inference using functional neuroimaging: dissociations versus associations. *Trends in Cognitive Sciences*, *10*(2), 64–9. doi:10.1016/j.tics.2005.12.005

Henson, R. N. (2007). Efficient Experimental Design for fMRI. In K. J. Friston, J. T. Ashburner, S. J. Kiebel, T. E. Nichols, & W. D. Penny (Eds.), *Statistical parametric mapping: The analysis of functional brain images* (pp. 193–210). Amsterdam: Academic Press.

Hicks, J. L., & Starns, J. J. (2014). Strength cues and blocking at test promote reliable within-list criterion shifts in recognition memory. *Memory & Cognition*. doi:10.3758/s13421-014-0397-y

Hildebrandt, H., Brand, a, & Sachsenheimer, W. (1998). Profiles of patients with left prefrontal and left temporal lobe lesions after cerebrovascular infarcations on California Verbal Learning Test-like indices. *Journal of Clinical and Experimental Neuropsychology*, *20*(5), 673–683. doi:10.1076/jcen.20.5.673.1119

Hochreiter, S., & Schmidhuber, J. (1997). Long Short-Term Memory. *Neural Computation*, *9*, 1735–1780.

Hodges, J. R., Salmon, D. P., & Butters, N. (1990). Differential impairment of semantic and episodic memory in Alzheimer ' s and Huntington ' s diseases : a controlled prospective study. *Reproduction*, 1089–1095.

Holroyd, C. B., & Yeung, N. (2012). Motivation of extended behaviors by anterior cingulate cortex. *Trends in Cognitive Sciences*, *16*(2), 121–127. doi:10.1016/j.tics.2011.12.008

Horvitz, J. C. (2000). Mesolimbocortical and nigrostriatal dopamine responses to salient non-reward events. *Neuroscience*, *96*(4), 651–656.

Horvitz, J. C., Stewart, T., & Jacobs, B. L. (1997). Burst activity of ventral tegmental dopamine neurons is elicited by sensory stimuli in the awake cat. *Brain Research*, *759*(2), 251–258. doi:10.1016/S0006-8993(97)00265-5

Huang, Y. Y., & Kandel, E. R. (1995). D1/D5 receptor agonists induce a protein synthesis-dependent late potentiation in the CA1 region of the hippocampus. *Proceedings of the National Academy of Sciences of the United States of America*, *92*(7), 2446–2450. doi:10.1073/pnas.92.7.2446

James, W. (1890). *The Principles of Psychology*. New York: Henry Holt and Company.

Jay, T. M. (2003). Dopamine: A potential substrate for synaptic plasticity and memory mechanisms. *Progress in Neurobiology*, *69*(6), 375–390. doi:10.1016/S0301-0082(03)00085-6

Jetter, W., Poser, U., Freeman, R. B., & Markowitsch, J. H. (1986). A verbal long term memory deficit in frontal lobe damaged patients. *Cortex*, *22*, 229–242.

Johnson, M. K., Hashtroudi, S., & Lindsay, D. S. (1993). Source monitoring. *Psychological Bulletin*, *114*(1), 3–28.

Kafkas, A., & Montaldi, D. (2014). Two separate, but interacting, neural systems for familiarity and novelty detection: A dual-route mechanism. *Hippocampus*, *24*(5), 516–527. doi:10.1002/hipo.22241

Kafkas, A., & Montaldi, D. (2015). Striatal and midbrain connectivity with the hippocampus selectively boosts memory for contextual novelty. *Hippocampus*. doi:10.1002/hipo.22434

Kahn, I., Davachi, L., & Wagner, A. D. (2004). Functional-neuroanatomic correlates of recollection: implications for models of recognition memory. *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience*, *24*(17), 4172–80. doi:10.1523/JNEUROSCI.0624-04.2004

Kakade, S., & Dayan, P. (2002). Dopamine: Generalization and bonuses. *Neural Networks*, *15*(4-6), 549–559. doi:10.1016/S0893-6080(02)00048-5

Kang, S. H. K., McDermott, K. B., & Roediger, H. L. (2007). Test format and corrective feedback modify the effect of testing on long-term retention. *European Journal of Cognitive Psychology*, *19*(4-5), 528–558. doi:10.1080/09541440601056620

Kantner, J., & Lindsay, D. S. (2010). Can corrective feedback improve recognition memory? *Memory & Cognition*, *38*(4), 389–406. doi:10.3758/MC.38.4.389

Kantner, J., & Lindsay, D. S. (2012). Response bias in recognition memory as a cognitive trait. *Memory & Cognition*. doi:10.3758/s13421-012-0226-0

Kim, H. (2011). Differential neural activity in the recognition of old versus new events: An Activation Likelihood Estimation Meta-Analysis. *Human Brain Mapping*, *000*. doi:10.1002/hbm.21474

Knoke, D., Taylor, A. E., & Saint-Cyr, J. A. (1998). The Differential Effects of Cueing on Recall in Parkinson's Disease and Normal Subjects. *Brain and Cognition*, *38*, 261–274.

Knowlton, B. J., Mangels, J. a, & Squire, L. R. (1996). A neostriatal habit learning system in humans. *Science (New York, N.Y.)*, *273*(5280), 1399–402.

Knowlton, B. J., Squire, L. R., & Gluck, M. A. (1994). Probabilistic Classification Learning in Amnesia. *Learning & Memory*, *1*, 106–120. doi:10.1101/lm.1.2.106

Knutson, B., Adams, C. M., Fong, G. W., & Hommer, D. (2001). Anticipation of increasing monetary reward selectively recruits nucleus accumbens. *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience*, *21*(16), RC159.

Knutson, B., Fong, C. A. G. W., Adams, C. M., Varner, J. L., & Hommer, D. (2001). Dissociation of reward anticipation and outcome with event-related fMRI, *12*(17), 3683–3687.

Knutson, B., Taylor, J., Kaufman, M., Peterson, R., & Glover, G. (2005). Distributed neural representation of expected value. *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience*, *25*(19), 4806–12. doi:10.1523/JNEUROSCI.0642-05.2005

Koop, G. J., Criss, A. H., & Malmberg, K. J. (2014). The role of mnemonic processes in pure-target and pure-foil recognition memory. *Psychonomic Bulletin & Review*, *22*(2), 509–516. doi:10.3758/s13423-014-0703-5

Krebs, R. M., Schott, B. H., Schütze, H., & Düzel, E. (2009). The novelty exploration bonus and its attentional modulation. *Neuropsychologia*, *47*(11), 2272–81. doi:10.1016/j.neuropsychologia.2009.01.015

Kuhl, B. a, Dudukovic, N. M., Kahn, I., & Wagner, A. D. (2007). Decreased demands on cognitive control reveal the neural processing benefits of forgetting. *Nature Neuroscience*, *10*(7), 908–14. doi:10.1038/nn1918

Landau, S. M., Lal, R., O'Neil, J. P., Baker, S., & Jagust, W. J. (2009). Striatal dopamine and working memory. *Cerebral Cortex (New York, N.Y. : 1991)*, *19*(2), 445–54. doi:10.1093/cercor/bhn095

Lauwereyns, J., Watanabe, K., Coe, B., & Hikosaka, O. (2002). A neural correlate of response bias in monkey caudate nucleus. *Nature*, *418*, 413–417. doi:10.1038/nature00844.1.

Law, C.-T., & Gold, J. I. (2009). Reinforcement learning can account for associative and perceptual learning on a visual-decision task. *Nature Neuroscience*, *12*(5), 655–663. doi:10.1038/nn.2304

Lemon, N., & Manahan-Vaughan, D. (2006). Dopamine D1/D5 receptors gate the acquisition of novel information through hippocampal long-term potentiation and long-term depression. *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience*, *26*(29), 7723–7729. doi:10.1523/JNEUROSCI.1454-06.2006

Lewis, S. J. G., Dove, A., Robbins, T. W., Barker, R. A., & Owen, A. M. (2004). Striatal contributions to working memory : a functional magnetic resonance imaging study in humans. *Neuroscience*, *19*. doi:10.1111/j.1460-9568.2003.03108.x

Li, S., Cullen, W. K., Anwyl, R., & Rowan, M. J. (2003). Dopamine-dependent facilitation of LTP induction in hippocampal CA1 by exposure to spatial novelty. *Nature Neuroscience*, *6*(5), 526–531. doi:10.1038/nn1049

Liljeholm, M., & O'Doherty, J. P. (2012). Contributions of the striatum to learning, motivation, and performance: an associative account. *Trends in Cognitive Sciences*, *16*(9), 467–475. doi:10.1016/j.tics.2012.07.007

Linden, D. E. J., Prvulovic, D., Formisano, E., Völlinger, M., Zanella, F. E., Goebel, R., & Dierks, T. (1999). The functional neuroanatomy of target detection: An fMRI study of visual and auditory oddball tasks. *Cerebral Cortex*, *9*, 815–823.

Lisman, J. E., & Grace, A. a. (2005). The hippocampal-VTA loop: controlling the entry of information into long-term memory. *Neuron*, *46*(5), 703–13. doi:10.1016/j.neuron.2005.05.002

Long, N. M., Oztekin, I., & Badre, D. (2010). Separable prefrontal cortex contributions to free recall. *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience*, *30*(33), 10967–76. doi:10.1523/JNEUROSCI.2611-10.2010

Lustig, C., Shah, P., Seidler, R., & Reuter-Lorenz, P. a. (2009). Aging, training, and the brain: a review and future directions. *Neuropsychology Review*, *19*(4), 504–22. doi:10.1007/s11065-009-9119-9

Macmillan, N. A., & Creelman, C. D. (2005). *Detection Theory: A User's Guide* (2nd ed.). Mahwah, New Jersey: Lawrence Erlbaum Associates.

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*(2), 303–19.

Maia, T. V, & Frank, M. J. (2011). From reinforcement learning models to psychiatric and neurological disorders. *Nature Neuroscience*, *14*(2), 154–62. doi:10.1038/nn.2723

Malmberg, K. J. (2008). Recognition memory: a review of the critical findings and an integrated theory for relating them. *Cognitive Psychology*, *57*(4), 335–84. doi:10.1016/j.cogpsych.2008.02.004

Marois, R., Leung, H. C., & Gore, J. C. (2000). A stimulus-driven approach to object identity and location processing in the human brain. *Neuron*, *25*(3), 717–728. doi:10.1016/S0896-6273(00)81073-9

Mayer, A. R., Harrington, D., Adair, J. C., & Lee, R. (2006). The neural networks underlying endogenous auditory covert orienting and reorienting. *NeuroImage*, *30*(3), 938–949. doi:10.1016/j.neuroimage.2005.10.050

McClelland, J. L., & Rogers, T. T. (2003). The parallel distributed processing approach to semantic cognition. *Nature Reviews. Neuroscience*, *4*(4), 310–22. doi:10.1038/nrn1076

McDaniel, M. A., & Bugg, J. M. (2012). Memory Training Interventions: What has been forgotten? *Journal of Applied Research in Memory and Cognition*, *1*(1), 58–60. doi:10.1016/j.jarmac.2011.11.002

McNab, F., & Klingberg, T. (2008). Prefrontal cortex and basal ganglia control access to working memory. *Nature Neuroscience*, *11*(1), 103–7. doi:10.1038/nn2024

Mickes, L., Hwe, V., Wais, P. E., & Wixted, J. T. (2011). Strong memories are hard to scale. *Journal of Experimental Psychology. General*, *140*(2), 239–257. doi:10.1037/a0023007

Middleton, E. L., & Schwartz, M. F. (2012). Errorless learning in cognitive rehabilitation: a critical review. *Neuropsychological Rehabilitation*, *22*(2), 138–68. doi:10.1080/09602011.2011.639619

Miller, E. K., & Cohen, J. D. (2001). An Integrative Theory of Prefrontal Cortex Function. *Annual Review of Neuroscience*, *24*, 167–202.

Mink, J. W. (1996). The basal ganglia: Focused selection and inhibition of competing motor programs. *Progress in Neurobiology*, *50*(4), 381–425. doi:10.1016/S0301-0082(96)00042-1

Mitchell, K. J., & Johnson, M. K. (2009). Source monitoring 15 years later: what have we learned from fMRI about the neural mechanisms of source memory? *Psychological Bulletin*, *135*(4), 638–77. doi:10.1037/a0015849

Montague, P. R. (1996). A Framework for Mesencephalic Predictive Hebbian Learning, *76*(5), 1936–1947.

Moran, J. M., & Zaki, J. (2013). Functional neuroimaging and psychology: what have you done for me lately? *Journal of Cognitive Neuroscience*, *25*, 834–42. doi:10.1162/jocn_a_00380

Moscovitch, M. (1992). Memory and Working-with-Memory: A Component Process Model Based on Modules and Central Systems. *Journal of Cognitive Neuroscience*, *4*(3), 257–267. doi:10.1162/jocn.1992.4.3.257

Moscovitch, M. (1994). Cognitive resources and dual-task interference effects at retrieval in normal people: The role of the frontal lobes and medial temporal cortex. *Neuropsychology*, *8*(4), 524–534. doi:10.1037//0894-4105.8.4.524

Moustafa, A. a, Sherman, S. J., & Frank, M. J. (2008). A dopaminergic basis for working memory, learning and attentional shifting in Parkinsonism. *Neuropsychologia*, *46*(13), 3144–56. doi:10.1016/j.neuropsychologia.2008.07.011

Moustafa, A. a., & Gluck, M. a. (2011). Computational cognitive models of prefrontal-striatal-hippocampal interactions in Parkinson's disease and schizophrenia. *Neural Networks*, *24*(6), 575–591. doi:10.1016/j.neunet.2011.02.006

Murty, V. P., & Adcock, R. A. (2013). Enriched Encoding: Reward Motivation Organizes Cortical Networks for Hippocampal Detection of Unexpected Events. *Cerebral Cortex (New York, N.Y. : 1991)*. doi:10.1093/cercor/bht063

Niv, Y., Edlund, J. A., Dayan, P., & O'Doherty, J. P. (2012). Neural prediction errors reveal a risk-sensitive reinforcement-learning process in the human brain. *The Journal of Neuroscience*, *32*(2), 551–62. doi:10.1523/JNEUROSCI.5498-10.2012

Norman, K. a, & O'Reilly, R. C. (2003). Modeling hippocampal and neocortical contributions to recognition memory: a complementary-learning-systems approach. *Psychological Review*, *110*(4), 611–46. doi:10.1037/0033-295X.110.4.611

Nyberg, L., Andersson, M., Forsgren, L., Jakobsson-Mo, S., Larsson, A., Marklund, P., … Bäckman, L. (2009). Striatal dopamine D2 binding is related to frontal BOLD response during updating of long-term memory representations. *NeuroImage*, *46*(4), 1194–9. doi:10.1016/j.neuroimage.2009.03.035

O'Connor, A. R., Han, S., & Dobbins, I. G. (2010). The inferior parietal lobule and recognition memory: expectancy violation or successful retrieval? *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience*, *30*(8), 2924–34. doi:10.1523/JNEUROSCI.4225-09.2010

O'Doherty, J. P., Dayan, P., Friston, K., Critchley, H., & Dolan, R. J. (2003). Temporal difference models and reward-related learning in the human brain. *Neuron*, *38*(2), 329–37.

O'Doherty, J. P., Dayan, P., Schultz, J., Deichmann, R., Friston, K., & Dolan, R. J. (2004). Dissociable roles of ventral and dorsal striatum in instrumental conditioning. *Science (New York, N.Y.)*, *304*(5669), 452–4. doi:10.1126/science.1094285

O'Doherty, J. P., Hampton, A., & Kim, H. (2007). Model-based fMRI and its application to reward learning and decision making. *Annals of the New York Academy of Sciences*, *1104*, 35–53. doi:10.1196/annals.1390.022

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*(2), 283–328. doi:10.1162/089976606775093909

Oostrom, I. Van, Dollfus, S., Brazo, P., Abadie, P., Halbecq, I., Marie, R. M., & Thery, S. (2003). Verbal learning and memory in schizophrenic and Parkinson ' s disease patients. *Psychiatry Research*, *117*, 25–34.

Oztekin, I., & Badre, D. (2011). Distributed Patterns of Brain Activity that Lead to Forgetting. *Frontiers in Human Neuroscience*, *5*(August), 1–8. doi:10.3389/fnhum.2011.00086

Packard, M. G., Hirsch, R., & White, N. M. (1989). Differential Effects of Fornix and Caudate Radial Maze Tasks : Evidence for Multiple Nucleus Lesions on Two Memory Systems. *Science*, *9*, 1465–1472.

Pashler, H., Cepeda, N. J., Wixted, J. T., & Rohrer, D. (2005). When does feedback facilitate learning of words? *Journal of Experimental Psychology. Learning, Memory, and Cognition*, *31*(1), 3–8. doi:10.1037/0278-7393.31.1.3

Poldrack, R. A., Clark, J., Paré-Blagoev, E. J., Shohamy, D., Creso Moyano, J., Myers, C., & Gluck, M. a. (2001). Interactive memory systems in the human brain. *Nature*, *414*, 546–50. doi:10.1038/35107080

Poldrack, R. A., & Foerde, K. (2008). Category learning and the memory systems debate. *Neuroscience and Biobehavioral Reviews*, *32*(2), 197–205. doi:10.1016/j.neubiorev.2007.07.007

Poldrack, R. A., Wagner, a D., Prull, M. W., Desmond, J. E., Glover, G. H., & Gabrieli, J. D. (1999). Functional specialization for semantic and phonological processing in the left inferior prefrontal cortex. *NeuroImage*, *10*(1), 15–35. doi:10.1006/nimg.1999.0441

Postma, A. (1999). The influence of decision criteria upon remembering and knowing in recognition memory. *Acta Psychologica*, *103*(1-2), 65–76.

Ratcliff, R. (1978). A Theory of Memory Retrieval. *Psychological Review*, *85*(2), 59–108.

Ratcliff, R., & McKoon, G. (2008). The diffusion decision model: theory and data for two-choice decision tasks. *Neural Computation*, *20*(4), 873–922. doi:10.1162/neco.2008.12-06-420

Ratcliff, R., Sheu, C. F., & Gronlund, S. D. (1992). Testing global memory models using ROC curves. *Psychological Review*, *99*(3), 518–35.

Ratcliff, R., & Starns, J. J. (2009). Modeling confidence and response time in recognition memory. *Psychological Review*, *116*(1), 59–83. doi:10.1037/a0014086

Ratcliff, R., & Starns, J. J. (2013). Modeling confidence judgments, response times, and multiple choices in decision making: recognition memory and motion discrimination. *Psychological Review*, *120*(3), 697–719. doi:10.1037/a0033152

Ratcliff, R., Thapar, A., & McKoon, G. (2004). A diffusion model analysis of the effects of aging on recognition memory. *Journal of Memory and Language*, *50*(4), 408–424. doi:10.1016/j.jml.2003.11.002

Roediger, H. L., & Butler, A. C. (2011). The critical role of retrieval practice in long-term retention. *Trends in Cognitive Sciences*, *15*(1), 20–7. doi:10.1016/j.tics.2010.09.003

Roediger, H. L., & Karpicke, J. D. (2006a). Test-enhanced learning: taking memory tests improves long-term retention. *Psychological Science*, *17*(3), 249–55. doi:10.1111/j.1467-9280.2006.01693.x

Roediger, H. L., & Karpicke, J. D. (2006b). The Power of Testing Memory . *Perspectives on Psychological Science*, *1*(3), 181–210. doi:10.1111/j.1745-6916.2006.00012.x

Roediger, H. L., & McDermott, K. B. (1995). Creating false memories: Remembering words not presented in lists. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *21*(4), 803–814. doi:10.1037/0278-7393.21.4.803

Rombouts, S. a, Barkhof, F., Witter, M. P., Machielsen, W. C., & Scheltens, P. (2001). Anterior medial temporal lobe activation during attempted retrieval of encoded visuospatial scenes: an event-related fMRI study. *NeuroImage*, *14*(1 Pt 1), 67–76. doi:10.1006/nimg.2001.0799

Rotello, C. M., & Macmillan, N. A. (2007). Response Bias in Recognition Memory. *The Psychology of Learning and Motivation*, *48*(07), 61–94. doi:10.1016/S0079-7421(07)48002-1

Rouleau, I., Imbault, H., Laframboise, M., & Bédard, M. A. (2001). Pattern of intrusions in verbal recall: comparison of Alzheimer's disease, Parkinson's disease, and frontal lobe dementia. *Brain and Cognition*, *46*(1-2), 244–249. doi:10.1016/S0278-2626(01)80076-2

Roy, M., Shohamy, D., & Wager, T. D. (2012). Ventromedial prefrontal-subcortical systems and the generation of affective meaning. *Trends in Cognitive Sciences*, *16*(3), 147–156. doi:10.1016/j.tics.2012.01.005

Rugg, M. D., & Vilberg, K. L. (2013). Brain networks underlying episodic memory retrieval. *Current Opinion in Neurobiology*, *23*(2), 255–60. doi:10.1016/j.conb.2012.11.005

Rugg, M. D., & Wilding, E. L. (2000). Retrieval processing and episodic memory. *Trends in Cognitive Sciences*, *4*(3), 108–115.

Rugg, M. D., & Yonelinas, A. P. (2003). Human recognition memory: a cognitive neuroscience perspective. *Trends in Cognitive Sciences*, *7*(7), 313–319. doi:10.1016/S1364-6613(03)00131-1

Rutishauser, U., Ye, S., Koroma, M., Tudusciuc, O., Ross, I. B., Chung, J. M., & Mamelak, A. N. (2015). Representation of retrieval confidence by single neurons in the human medial temporal lobe. *Nature Neuroscience*, *18*(7). doi:10.1038/nn.4041

Sagar, H. J., Sullivan, E. V, Cooper, J. A., & Jordan, N. (1991). Normal release from proactive interference in untreated patients with Parkinson's disease. *Neuropsychologia*. doi:10.1016/0028-3932(91)90075-J

Schacter, D. L. (1999). Suppressing False Recognition in Younger and Older Adults: The Distinctiveness Heuristic. *Journal of Memory and Language*, *40*(1), 1–24. doi:10.1006/jmla.1998.2611

Schacter, D. L., & Addis, D. R. (2007). The cognitive neuroscience of constructive memory: remembering the past and imagining the future. *Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences*, *362*(1481), 773–86. doi:10.1098/rstb.2007.2087

Schacter, D. L., Curran, T., Galluccio, L., Milberg, W. P., & Bates, J. F. (1996). False recognition and the right frontal lobe: a case study. *Neuropsychologia*, *34*(8), 793–808.

Schacter, D. L., Eich, J. E., & Tulvlng, E. (1978). Richard Semon's Theory of Memory. *Journal of Verbal Learning and Verbal Behaviour*, *17*, 721–743.

Schacter, D. L., & Wagner, A. D. (1999). Medial temporal lobe activations in fMRI and PET studies of episodic encoding and retrieval. *Hippocampus*, *9*(1), 7–24. doi:10.1002/(SICI)1098-1063(1999)9:1<7::AID-HIPO2>3.0.CO;2-K

Schnyer, D. M., Verfaellie, M., Alexander, M. P., LaFleche, G., Nicholls, L., & Kaszniak, A. W. (2004). A role for right medial prefontal cortex in accurate feeling-of-knowing judgements: evidence from patients with lesions to frontal cortex. *Neuropsychologia*, *42*(7), 957–66. doi:10.1016/j.neuropsychologia.2003.11.020

Schonberg, T., Daw, N. D., Joel, D., & O'Doherty, J. P. (2007). Reinforcement Learning Signals in the Human Striatum Distinguish Learners from Nonlearners during

Reward-Based Decision Making. *Journal of Neuroscience*, *27*(47), 12860–12867. doi:10.1523/JNEUROSCI.2496-07.2007

Schott, B. H., Sellner, D. B., Lauer, C.-J., Habib, R., Frey, J. U., Guderian, S., … Düzel, E. (2004). Activation of midbrain structures by associative novelty and the formation of explicit memory in humans. *Learning & Memory (Cold Spring Harbor, N.Y.)*, *11*(4), 383–7. doi:10.1101/lm.75004

Schultz, W. (1998). Predictive reward signal of dopamine neurons. *Journal of Neurophysiology*, *80*(1), 1–27.

Schultz, W., Dayan, P., & Montague, P. R. (1997). A neural substrate of prediction and reward. *Science (New York, N.Y.)*, *275*(5306), 1593–9.

Schultz, W., & Dickinson, A. (2000). Neuronal coding of prediction errors. *Annual Review of Neuroscience*, *23*, 473–500.

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*(3), e54324. doi:10.1371/journal.pone.0054324

Scimeca, J. M., & Badre, D. (2012). Striatal Contributions to Declarative Memory Retrieval. *Neuron*, *75*(3), 380–392. doi:10.1016/j.neuron.2012.07.014

Scoville, W., & Milner, B. (1957). Loss of recent memory after bilateral hippocampal lesions. *Journal of Neurology, Neurosurgery & …*, *20*, 11–21.

Seger, C. a, & Miller, E. K. (2010). Category learning in the brain. *Annual Review of Neuroscience*, *33*, 203–19. doi:10.1146/annurev.neuro.051508.135546

Shen, W., Flajolet, M., Greengard, P., & Surmeier, D. J. (2008). Dichotomous dopaminergic control of striatal synaptic plasticity. *Science (New York, N.Y.)*, *321*(5890), 848–851. doi:10.1126/science.1160575

Shiffrin, R., & Steyvers, M. (1997). A model for recognition memory: REM—retrieving effectively from memory. *Psychonomic Bulletin & Review*, *4*(2), 145–166.

Shohamy, D., & Adcock, R. A. (2010). Dopamine and adaptive memory. *Trends in Cognitive Sciences*, *14*(10). doi:10.1016/j.tics.2010.08.002

Shohamy, D., Myers, C. E., Grossman, S., Sage, J., Gluck, M. a, & Poldrack, R. a. (2004). Cortico-striatal contributions to feedback-based learning: converging data from neuroimaging and neuropsychology. *Brain : A Journal of Neurology*, *127*(Pt 4), 851–9. doi:10.1093/brain/awh100

Simpson, E. (1951). The interpretation of interaction in contingency tables. *Journal of the Royal Statistical Society. Series B (…*, *13*(2), 238–241.

Solomon, A. C., Stout, J. C., Johnson, S. a, Langbehn, D. R., Aylward, E. H., Brandt, J., … Paulsen, J. S. (2007). Verbal episodic memory declines prior to diagnosis in Huntington's disease. *Neuropsychologia*, *45*(8), 1767–76. doi:10.1016/j.neuropsychologia.2006.12.015

Spaniol, J., Davidson, P. S. R., Kim, A. S. N., 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*(8-9), 1765–79. doi:10.1016/j.neuropsychologia.2009.02.028

Squire, L. R. (1986). Mechanisms of memory. *Science*, *232*(4758), 1612–1619. doi:10.1126/science.3086978

Squire, L. R. (1992). Memory and the hippocampus: a synthesis from findings with rats, monkeys, and humans. *Psychological Review*, *99*(2), 195–231.

Squire, L. R., & Wixted, J. T. (2011). The cognitive neuroscience of human memory since H.M. *Annual Review of Neuroscience*, *34*, 259–88. doi:10.1146/annurev-neuro-061010-113720

Stadler, M. a, Roediger, H. L., & McDermott, K. B. (1999). Norms for word lists that create false memories. *Memory & Cognition*, *27*(3), 494–500.

Starns, J. J., Ratcliff, R., & McKoon, G. (2012). Evaluating the unequal-variance and dual-process explanations of zROC slopes with response time data and the diffusion model. *Cognitive Psychology*, *64*(1-2), 1–34. doi:10.1016/j.cogpsych.2011.10.002

Starns, J. J., Ratcliff, R., & White, C. N. (2012). Diffusion model drift rates can be influenced by decision processes: an analysis of the strength-based mirror effect. *Journal of Experimental Psychology. Learning, Memory, and Cognition*, *38*(5), 1137–51. doi:10.1037/a0028151

Starns, J. J., White, C. N., & Ratcliff, R. (2010). A direct test of the differentiation mechanism: REM, BCDMEM, and the strength-based mirror effect in recognition memory. *Journal of Memory and Language*, *63*(1), 18–34. doi:10.1016/j.jml.2010.03.004

Stelzel, C., Basten, U., Montag, C., Reuter, M., & Fiebach, C. J. (2010). Frontostriatal Involvement in Task Switching Depends on Genetic Differences in D2 Receptor Density. *Journal of Neuroscience*, *30*(42), 14205–14212. doi:10.1523/JNEUROSCI.1062-10.2010

Stewart, N., Brown, G. D. a, & Chater, N. (2005). Absolute identification by relative judgment. *Psychological Review*, *112*(4), 881–911. doi:10.1037/0033-295X.112.4.881

Stretch, V., & Wixted, J. (1998a). Decision rules for recognition memory confidence judgments. *Journal of Experimental Psychology. Learning, Memory, and Cognition*, *24*(6), 1397–1410.

Stretch, V., & Wixted, J. T. (1998b). On the difference between strength-based and frequency-based mirror effects in recognition memory. *Journal of Experimental Psychology. Learning, Memory, and Cognition*, *24*(6), 1379–96.

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. *Neuropsychology*, *8*(3), 355–373.

Sutton, R. (1988). Learning to predict by the methods of temporal differences. *Machine Learning*, *3*, 9–44.

Sutton, R., & Barto, A. G. (1998). *Reinforcement Learning: An Itroduction*. Cambridge, MA: MIT Press.

Talairach, J., & Tournoux, P. (1988). *Co-planar stereotaxic atlas of the human brain*. New York: Theime.

Taylor, A. E., Saint-Cyr, J. a, & Lang, A. E. (1990). Memory and learning in early Parkinson's disease: Evidence for a "Frontal Lobe Syndrome." *Brain and Cognition*, *13*, 211–232.

Thompson-Schill, S. L., D'Esposito, M., Aguirre, G. K., & Farah, M. J. (1997). Role of left inferior prefrontal cortex in retrieval of semantic knowledge: a reevaluation. *Proceedings of the National Academy of Sciences of the United States of America*, *94*(26), 14792–7.

Thompson-Schill, S. L., Swick, D., Farah, M. J., D'Esposito, M., Kan, I. P., & Knight, R. T. (1998). Verb generation in patients with focal frontal lesions: a neuropsychological test of neuroimaging findings. *Proceedings of the National Academy of Sciences of the United States of America*, *95*(26), 15855–60.

Tricomi, E. M., Delgado, M. R., & Fiez, J. a. (2004). Modulation of caudate activity by action contingency. *Neuron*, *41*(2), 281–92.

Tricomi, E. M., & Fiez, J. a. (2008). Feedback signals in the caudate reflect goal achievement on a declarative memory task. *NeuroImage*, *41*(3), 1154–67. doi:10.1016/j.neuroimage.2008.02.066

Tulving, E. (1972). Epsodic and semantic memory. In E. Tulving & W. Donaldson (Eds.), *Organization of memory* (pp. 382–403). New York: Academic Press.

Tzourio-Mazoyer, N., Landeau, B., Papathanassiou, D., Crivello, F., Etard, O., Delcroix, N., … Joliot, M. (2002). Automated anatomical labeling of activations in SPM using a macroscopic anatomical parcellation of the MNI MRI single-subject brain. *NeuroImage*, *15*(1), 273–89. doi:10.1006/nimg.2001.0978

Van Zandt, T. (2000). ROC curves and confidence judgements in recognition memory. *Journal of Experimental Psychology. Learning, Memory, and Cognition*, *26*(3), 582–600.

Verde, M. F., & Rotello, C. M. (2007). Memory strength and the decision process in recognition memory. *Memory & Cognition*, *35*(2), 254–62.

Vingerhoets, G., Vermeule, E., & Santens, P. (2005). Impaired intentional content learning but spared incidental retention of contextual information in non-demented patients with Parkinson's disease. *Neuropsychologia*, *43*, 675–81. doi:10.1016/j.neuropsychologia.2004.09.003

Voss, A., Rothermund, K., & Brandtstädter, J. (2008). Interpreting ambiguous stimuli: Separating perceptual and judgmental biases. *Journal of Experimental Social Psychology*, *44*(4), 1048–1056. doi:10.1016/j.jesp.2007.10.009

Wagner, A. D., Shannon, B. J., Kahn, I., & Buckner, R. L. (2005). Parietal lobe contributions to episodic memory retrieval. *Trends in Cognitive Sciences*, *9*(9), 445–53. doi:10.1016/j.tics.2005.07.001

Warren, D. E., Jones, S. H., Duff, M. C., & Tranel, D. (2014). False recall is reduced by damage to the ventromedial prefrontal cortex: implications for understanding the neural correlates of schematic memory. *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience*, *34*(22), 7677–82. doi:10.1523/JNEUROSCI.0119-14.2014

Weiermann, B., Stephan, M. a, Kaelin-Lang, A., & Meier, B. (2010). Is there a recognition memory deficit in Parkinson's disease? Evidence from estimates of recollection and familiarity. *The International Journal of Neuroscience*, *120*(3), 211–6. doi:10.3109/00207450903506510

Wheeler, M. A., Stuss, D. T., & Tulving, E. (1997). Toward a Theory of Episodic Memory : The Frontal Lobes and Autonoetic Consciousness. *Psychological Bulletin*, *121*(3), 331–354.

Wheeler, Stuss, & Tulving. (1995). Frontal Lobe Damage Produces Episodic Memory Impairment. *Journal of the International Neuropsychological Society : JINS*, *1*, 525–536.

White, C. N., & Poldrack, R. a. (2014). Decomposing bias in different types of simple decisions. *Journal of Experimental Psychology. Learning, Memory, and Cognition*, *40*(2), 385–98. doi:10.1037/a0034851

Whittington, C. J., Podd, J., & Kan, M. M. (2000). Recognition memory impairment in Parkinson's disease: Power and meta-analyses. *Neuropsychology*, *14*(2), 233–246. doi:10.1037//0894-4105.14.2.233

Whittlesea, B. W., & Price, J. R. (2001). Implicit/explicit memory versus analytic/nonanalytic processing: rethinking the mere exposure effect. *Memory & Cognition*, *29*(2), 234–46.

Wickens, T. D. (2002). *Elementary signal detection theory*. New York: Oxford University Press.

Wiecki, T. V, Sofer, I., & Frank, M. J. (2013). HDDM: Hierarchical Bayesian estimation of the Drift-Diffusion Model in Python. *Frontiers in Neuroinformatics*, *7*(August), 14. doi:10.3389/fninf.2013.00014

Wittmann, B. C., Bunzeck, N., Dolan, R. J., & Düzel, E. (2007). Anticipation of novelty recruits reward system and hippocampus while promoting recollection. *NeuroImage*, *38*(1), 194–202. doi:10.1016/j.neuroimage.2007.06.038

Wittmann, B. C., Daw, N. D., Seymour, B., & Dolan, R. J. (2008). Striatal activity underlies novelty-based choice in humans. *Neuron*, *58*(6), 967–73. doi:10.1016/j.neuron.2008.04.027

Wittmann, B. C., Schott, B. H., Guderian, S., Frey, J. U., Heinze, H.-J., & Düzel, E. (2005). Reward-related FMRI activation of dopaminergic midbrain is associated with enhanced hippocampus-dependent long-term memory formation. *Neuron*, *45*(3), 459–67. doi:10.1016/j.neuron.2005.01.010

Wixted, J. T. (2007). Dual-Process Theory and Signal-Detection Theory of Recognition Memory. *Psychological Review*, *114*(1), 152–176. doi:10.1037/0033-295X.114.1.152

Wixted, J. T., & Gaitan, S. C. (2002). Cognitive theories as reinforcement history surrogates: the case of likelihood ratio models of human recognition memory. *Animal Learning & Behavior*, *30*(4), 289–305.

Wolosin, S. M., Zeithamova, D., & Preston, A. R. (2012). Reward modulation of hippocampal subfield activation during successful associative encoding and retrieval. *Journal of Cognitive Neuroscience*, *24*(7), 1532–47. doi:10.1162/jocn_a_00237

Yonelinas, A. P. (2002). The Nature of Recollection and Familiarity: A Review of 30 Years of Research. *Journal of Memory and Language*, *46*(3), 441–517. doi:10.1006/jmla.2002.2864

Yonelinas, A. P., Otten, L. J., Shaw, K. N., & Rugg, M. D. (2005). Separating the Brain Regions Involved in Recollection and Familiarity in Recognition Memory. *Journal of Neuroscience*, *25*(11), 3002–3008. doi:10.1523/JNEUROSCI.5295-04.2005

Yonelinas, A. P., Regehr, G., & Jacoby, L. L. (1995). Incorporating response bias in a dual-process theory of memory. *Journal of Memory and Language*. doi:10.1006/jmla.1995.1036

Zajonc, R. B. (1968). Attitudinal Effects of Mere Exposure. *Journal of Personality and Social Psychology*, *9*(2), 1–27.

Zajonc, R. B. (2001). Mere Exposure: A Gateway to the Subliminal. *Current Directions in Psychological Science*, *10*(6), 224–228. doi:10.1111/1467-8721.00154

Zgaljardic, D. J., Borod, J. C., & Foldi, N. S. (2003). Disease : Relationship to Frontostriatal Circuitry. *Cognitive and Behavioral Neurology*, *16*(4).