Check for updates

Received: 23 October 2020 | Revised: 4 February 2021 | Accepted: 6 February 2021

DOI: 10.1002/hipo.23317

RAPID COMMUNICATION

WILEY

Probabilistic value learning in medial temporal lobe amnesia

Daniela J. Palombo[1,2] | Virginie M. Patt[2,3] | Renee Hunsberger[2] | Mieke Verfaellie[2,3]

[1]Department of Psychology, University of British Columbia, Vancouver, British Columbia, Canada

[2]VA Boston Healthcare System, Boston, Massachusetts

[3]Department of Psychiatry, Boston University School of Medicine, Boston, Massachusetts

Correspondence
Daniela J. Palombo, Department of Psychology, University of British Columbia, 2136 West Mall, Vancouver, BC, V6T 1Z4, Canada.
Email: daniela.palombo@ubc.ca

Mieke Verfaellie, VA Boston Healthcare System, 150 S Huntington Ave, Boston, MA 02130.
Email: mieke.verfaellie@va.gov

Funding information
Canadian Foundation for Innovation; Natural Sciences and Engineering Research Council of Canada; Social Sciences and Humanities Research Council of Canada; US Department of Veterans Affairs

Abstract

A prevailing view in cognitive neuroscience suggests that different forms of learning are mediated by dissociable memory systems, with a mesolimbic (i.e., midbrain and basal ganglia) system supporting incremental trial-and-error reinforcement learning and a hippocampal-based system supporting episodic memory. Yet, growing evidence suggests that the hippocampus may also be important for trial-and-error learning, particularly value or reward-based learning. In the present report, we use a lesion-based neuropsychological approach to clarify hippocampal contributions to such learning. Six amnesic patients with medial temporal lobe damage and a group of healthy controls were administered a simple value-based learning task involving probabilistic trial-and-error acquisition of stimulus–response-outcome (reward or none) contingencies modeled after Li et al. (Proceedings of the National Academy of Sciences, 2011, 108(1), 55–60). As predicted, patients were significantly impaired on the task, demonstrating reduced learning of the contingencies. Our results provide further supportive evidence that the hippocampus' role in cognition extends beyond episodic memory tasks and call for further refinement of theoretical models of hippocampal functioning.

KEYWORDS
amnesia, hippocampus, medial temporal lobes, reinforcement learning, reward learning

1 | INTRODUCTION

As we navigate a complex world, many of our actions are driven by our desire to maximize reward or avoid punishment. But the environment is not fully predictable, and we often learn through trial and error (over repeated experiences) what actions are most likely to yield a more positive outcome, adjusting our behavior accordingly. When deciding whether or not to pack an umbrella for a walk on a dark and cloudy day, we estimate that it will likely rain.

The field of cognitive neuroscience has long been interested in the neural substrates of such probabilistic reward learning, with a wealth of evidence emphasizing the importance of the mesolimbic (i.e., midbrain and basal ganglia) system (e.g., Knowlton, Mangels, & Squire, 1996; Shohamy, Myers, Kalanithi, & Gluck, 2008). Yet, some neuroimaging evidence suggests that the medial temporal lobe (MTL), particularly the hippocampus, is also important for this form of learning (e.g., Dickerson, Li, & Delgado, 2011; Li, Delgado, & Phelps, 2011; Palombo, Hayes, Reid, & Verfaellie, 2019; Schonberg et al., 2010), with hippocampal recruitment (alongside mesolimbic regions) demonstrated during probabilistic learning, particularly when the task has a value or reward component. Adding to this evidence, other human work examining field potentials in the hippocampus shows maximal neural firing in this region when rewards are uncertain (compared to certain), suggesting the hippocampus may be important for computing a reward uncertainty signal (Vanni-Mercier, Mauguière, Isnard, & Dreher, 2009). Animal studies also implicate the hippocampus in probabilistic value learning, providing additional converging evidence (e.g., Lee, Ghim, Kim, Lee, & Jung, 2012; Seib, Espinueva, Floresco, & Snyder, 2020).

Hippocampal involvement in such learning is surprising when considering the classic memory systems view (Knowlton et al., 1996;

Published 2021. This article is a U.S. Government work and is in the public domain in the USA.

Squire, 1987), which posits that the hippocampus is primarily involved in episodic memory (memory for singular episodes) but not in incremental learning (and vice versa for the striatum). However, hippocampal involvement in probabilistic reward learning is less surprising from an anatomical standpoint as studies show that the hippocampus is interconnected with the nucleus accumbens and the ventral tegmental area, providing important functional input into this mesolimbic circuitry (e.g., Floresco, Todd, & Grace, 2001; Gasbarri, Packard, Campana, & Pacitti, 1994; Groenewegen, der Zee, te Kortschot, & Witter, 1987) and suggestive of functional synergy. These regions also show functional coupling at rest in humans as measured through fMRI (Kahn & Shohamy, 2013).

Neuropsychological studies also elucidate the functional significance of prior observations of hippocampal fMRI engagement in learning. To that end, in a recent study (Palombo et al., 2019) we administered a novel probabilistic value-based learning task to a group of amnesic patients with MTL lesions and controls. Here, participants learned whether players, differentiated by patterned jumpsuits, win money in a "game." Patients were impaired in acquiring the value-based contingencies of the task. Also relevant is an earlier study by Hopkins, Myers, Shohamy, Grossman, and Gluck (2004), that showed that amnesic patients could not learn ice cream preferences for Mr. Potato Head characters (with feedback on correct trials provided in the form of coins in a jar). Together, these findings provide evidence that the hippocampus is necessary for probabilistic value or reward-based learning in certain conditions.

Here, we try to further elucidate these conditions. In prior studies, the tasks used entailed complex cue or feedback structure, with heavy demands placed on integration. For example, in our prior study, we orthogonalized reward outcome and accuracy; in that task, the feedback provided to the participant emphasized the outcome for the player but was independent of whether the participant's choice was correct. This design leaves open the possibility that the hippocampus was necessary to resolve the ambiguity between player and participant outcome (i.e., integrate two sources of information, such as "I was incorrect but the player won"). In a similar vein, it is possible that in a prior study by Hopkins et al. (2004), the configural nature of the cues contributed to the impairment, given the role of the hippocampus in configural learning (Ballard, Wagner, & McClure, 2018; Duncan, Doll, Daw, & Shohamy, 2018).

An important question, then, is whether the hippocampus is also necessary for probabilistic reward learning when the task has simple figural demands and direct feedback structure. To answer this question, we designed a task that closely mirrors a paradigm developed by Li et al. (2011), wherein participants learn, probabilistically, the numeric value of a simple stimulus to attain (hypothetical) reward (also see Dickerson et al., 2011). We administered the task to MTL patients and healthy controls. Based on Li et al., who showed MTL activation in their paradigm, we predicted that patients with amnesia would fail to learn well the contingencies in this task, consistent with the hypothesis that the MTL is necessary for incremental value-based learning.

Six amnesic patients with MTL lesions (one female, five males) participated. Their neuropsychological profiles indicated severe memory impairment in the context of otherwise preserved cognition (Table 1). Etiologies of memory impairment included hypoxic–ischemic injury (n = 3), status epilepticus followed by left temporal lobectomy (n = 1), and stroke (n = 2). Lesions for five of the six patients are presented in Figure 1, as either MRI or CT images. Due to medical contraindications, P04 could not be scanned, and his MTL pathology was inferred based on etiology and neuropsychological profile. Of the

TABLE 1 Demographics and neuropsychological information for patients

				WAIS III		WMS III			Volume loss (%)	
Etiology		Age	Edu	VIQ	WMI	GM	VD	AD	Hippocampal	Subhippocampal
Patients										
P1	Hypoxic–ischemic	70	12	88	75	52	56	55	N/A	N/A
P2	Status epilepticus + left temporal lobectomy	56	16	93	94	49	53	52	63	60[a]
P3	Hypoxic–ischemic	63	14	106	115	59	72	52	22	–
P4	Hypoxic–ischemic	67	17	131	126	86	78	86	N/A	N/A
P5	Stroke	66	18	117	88	67	75	55	62	–
P6	Stroke	55	20	111	99	60	65	58	43	–
Mean		62.8	16.2	108	99.5	62.2	66.5	59.7		
Controls										
Mean		62.3	16.1	113						

Abbreviations: AD, auditory delayed; Age, age in years; Edu, education in years; GM, general memory; N/A, not available; WAIS-III, Wechsler adult intelligence scale-third edition; WMI, working memory index; WMS-III, Wechsler memory scale-third edition; VD, visual delayed; VIQ, verbal intelligence quotient.

[a]Volume loss in left anterior parahippocampal gyrus (i.e., entorhinal cortex, medial portion of the temporal pole, and the medial portion of perirhinal cortex; see Kan, Giovanello, Schnyer, Makris, & Verfaellie, 2007, for methodology).

Note that scores on the WAIS III and the WMS III are standardized scores (mean of 100; *SD* of 15).

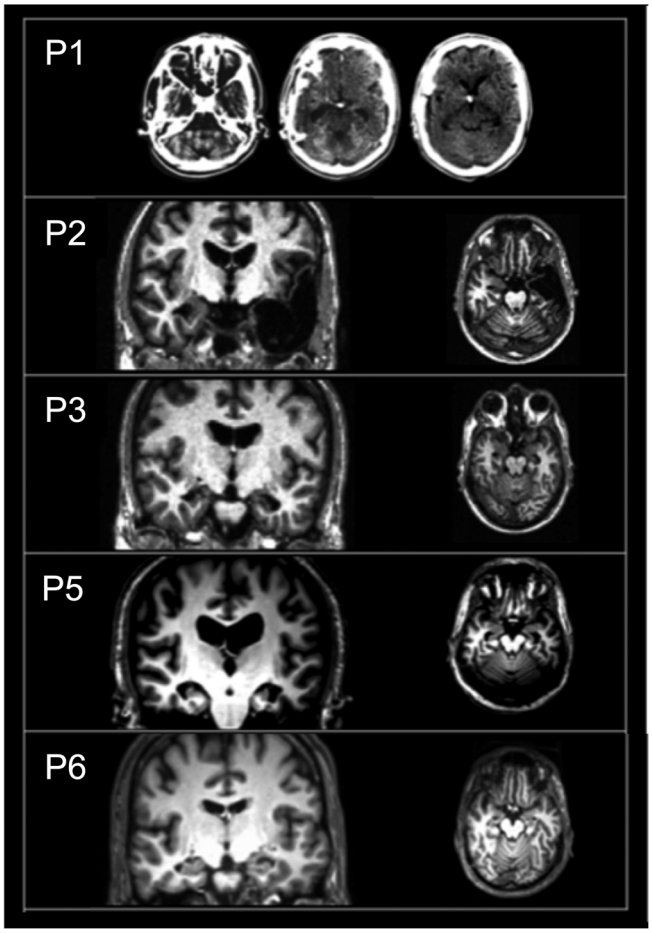

FIGURE 1 Structural MRI and CT scans depicting MTL lesions for five of the six participants with amnesia. The left side of the brain is displayed on the right side of the image. CT slices show lesion location for P1 in the axial plane. T1-weighted MRI images depict lesions for P2, P3, P5, and P6 in the coronal and axial plane

patients with available scans, two patients (P03, P05), had lesions restricted to the hippocampus, one patient's (P06) hippocampal lesion extended into the amygdala, whereas two patients had lesions including MTL cortices (P01, P02), with further extension into lateral temporal cortices in P02. As shown in Table 1, volumetric data for the hippocampus and MTL cortices was available for four of the six patients (P02, P03, P05, P06) using methodology reported elsewhere (see Kan et al., 2007).

Fourteen healthy control participants (four females, 10 males) matched to the patient group in age (M = 62.29, SD = 6.60), education (M = 16.07, SD = 2.59), and verbal IQ (M = 113.0, SD = 12.13) participated in the study. All participants provided informed consent in accordance with the Institutional Review Board at the VA Boston Healthcare System and were compensated for their time at a fixed hourly rate (participants were not compensated extra based on their performance).

The probabilistic learning paradigm, modelled after Li et al. (2011) required participants to learn whether each of four visual patterns usually hid a number larger or smaller than the number 10. Each of the patterns consisted of a nonverbalizable symbol presented in a 5 by 5 matrix (see Figure 2a). The number hidden beneath each visual pattern changed with each trial and could be any even number between 2 and 18, excluding 10. Two of the patterns were associated with numbers larger than 10 on 75% of the trials, whereas the other two were associated with numbers smaller than 10 on 75% of the trials.

On each trial, participants saw one of the visual patterns next to the number 10, and were instructed to choose which they thought was larger, the number hidden by the pattern or the number 10, via key press on a keyboard. To simplify the task, the number 10 was always displayed on the right side of the screen. Participants were given 4,000 ms to make their selection. If a participant failed to respond during this time, a screen displaying "Too late!" appeared. If a choice was made during the allotted time, the participant's selection was framed by a black square (3,000 ms). After this time, participants were shown the hidden number, and given feedback on the accuracy of their choice (for 1,500 ms). If a correct selection was made, a dollar bill appeared at the top of the screen, with the words "You win!" An incorrect selection was followed by a grey rectangle with the words "You lose!" (See Figure 2b). Trials were separated by a 2000 ms interstimulus interval.

Participants completed three blocks of 48 trials, with only a few seconds separating blocks for a brief instruction reminder. Within each block, each pattern appeared 12 times, and each of the potential hidden numbers appeared six times. The order of trials was pseudorandomized so that each pattern had one minority trial within each third of the block, and no pattern was presented more than twice in a row. Majority status of a visual pattern (i.e., whether typically associated with numbers larger or smaller than 10) was counterbalanced across participants.

Immediately after the three learning blocks, participants completed a test phase, using the same patterns. Here, participants saw two of the patterns side by side and were instructed to choose which they thought was associated with a number larger than 10, based on the game they had just played (i.e., the learning phase). Participants were given 4,000 ms to respond, and saw a screen displaying "Too late!" if they did not respond within the allotted time. Participants were not given feedback. The task was run using E-prime (Version 2.0) on a PC Lenovo Thinkpad laptop.

Patients on average responded on 97.11% of trials (SD = 3.76%) and control participants on average responded on 98.56% of the trials (SD = 1.77%). This suggests that participants in both groups had sufficient time to respond throughout the task.

For all analyses, accuracy (correct/[correct + incorrect trials]) was calculated based on the majority outcome status of the pattern (as in Palombo et al., 2019). Trials with no response were not included in the accuracy calculation. The mean accuracy for each group across the three learning blocks, as well as individual subject performance, is shown in Figure 3. A 2 (Group) × 3 (Block) mixed factorial ANOVA revealed a significant main effect of group, $F(1,18) = 7.87$, $p = .01$, $\eta_p^2 = 0.30$, and learning block, $F(2,36) = 14.88$, $p < .001$, $\eta_p^2 = 0.45$. The effect of block did not differ across groups, $F(2,36) = 1.93$,

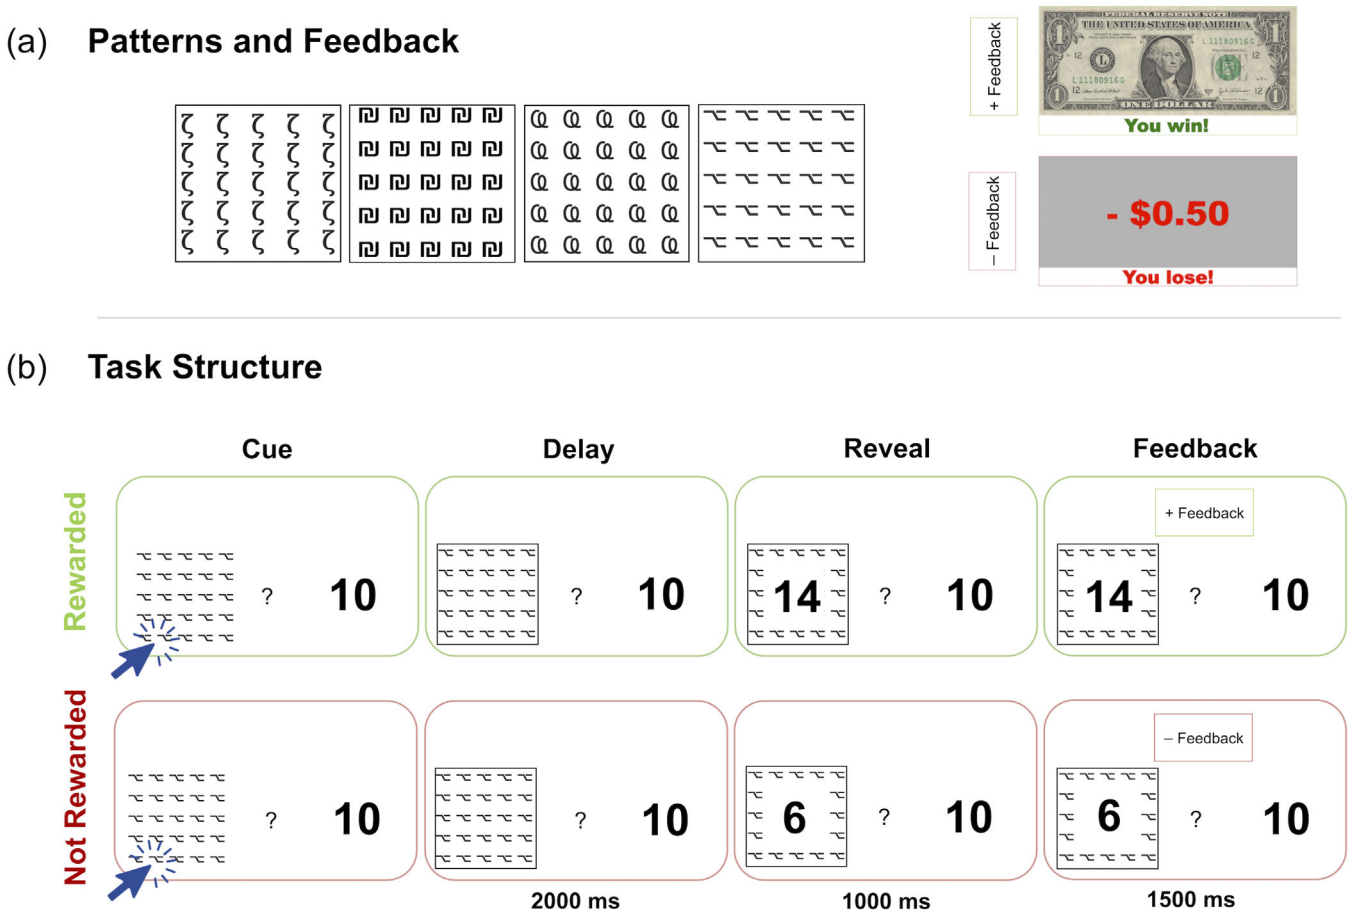

FIGURE 2 (a) The four stimulus patterns used in the learning paradigm. (b) Schematic of the learning paradigm with one stimulus pattern shown. As depicted in the upper row, the outcome was greater than 10 on 75% of trials (majority status trials); the lower row shows that the outcome was less than 10 on 25% of trials (minority status trials). The blue arrow is depicted for display purposes to indicate the participant's choice in this example [Color figure can be viewed at wileyonlinelibrary.com]

p = .16, η_p^2 = 0.10 (Patients, First Block: 50.89%, *SD* = 1.83%; Second Block: 52.89%, *SD* = 5.69%; Third Block: 63.26%, *SD* = 9.27%; Controls, First Block: 54.69%, *SD* = 10.06%; Second Block: 68.25%, *SD* = 12.33%; Third Block: 77.00%, *SD* = 14.23%). Independent samples t-tests showed that patients performed significantly worse on both the second, $t(18)$ = 2.89, p = .01, d = 1.60; and the third learning blocks, $t(18)$ = 2.16, p = .045, d = 1.14. This was not the case for the first learning block, $t(18)$ = 0.91, p = .38, d = 0.53.[1] Hence, the main effect of group was mainly driven by performance differences in the second and third blocks. We note that by the third block, both patients and controls performed significantly above chance (controls: $t(13)$ = 7.10, p < .001, d = 1.90; patients: $t(5)$ = 3.50, p = .02, d = 1.43). At the individual level, 11/14 controls (79%) performed above chance, whereas only 3/6 patients (50%) performed above chance, including one of the patients with a hippocampal-only lesion.

During the test phase, patients also performed significantly worse (*M* = 52.82%, *SD* = 21.33%) than controls (*M* = 83.33; *SD* = 23.57%), $t(18)$ = 2.72, p = .01, d = 1.36 (see Figure 3).[2] A 2 (Group) × 2 (Block) ANOVA comparing test phase performance relative to the third block of learning revealed a main effect of group, $F(1,18)$ = 9.04, p = .008, η_p^2 = 0.33, but no main effect of block, $F(1,18)$ = 0.15, p = .71, η_p^2 = 0.01, or group by block interaction, $F(1,18)$ = 2.43, p = .14, η_p^2 = 0.12. This suggests that there were no group differences in the retention of knowledge gained by the third learning block during the test phase. However, Figure 3 shows that whereas individual controls generally showed maintenance or nominally improved performance from the third block to test, all but one patient showed a nominal decrease.

An ancillary analysis examined whether the outcome status (greater or smaller than 10 on majority trials) associated with each pattern affected performance. There was not a significant main effect of pattern status, $F(1,18)$ = 0.40, p = .54, η_p^2 = 0.02, or a significant interaction between group and pattern status, $F(1,18)$ = 0.82, p = .38, η_p^2 = 0.04. Consistent with previous analyses, there was a significant main effect of group, $F(1,18)$ = 7.99, p = .011, η_p^2 = 0.31. These analyses suggest that neither performance of patients nor controls was affected by whether patterns were usually associated with values larger or smaller than 10.

To further clarify the role of the MTL in learning, we fitted a standard reinforcement learning (RL) model (Sutton & Barto, 2018) to the

trial-by-trial choice data. The expected reinforcing value "$V(s, a)$" of each possible action "a" (i.e., choice of the pattern or of the number 10) given presentation of a specific pattern "s" was initially set to 0.5 and subsequently updated for the selected action using a delta rule: $V(s, a) \leftarrow V(s, a) + \alpha(r - V(s, a))$, where "$r$" is the reinforcement outcome (equal to 1 if the participant's choice was reinforced and to 0 if it was not reinforced), $(r - V(s, a))$ represents the prediction error, and "α" is the updating parameter (i.e., the amount of influence given to the prediction error for updating expected values of options). The probability of choosing the pattern (or number 10) option was then calculated using a softmax rule:

$$P(a = \text{Pattern}/s) = e^{\beta V(s,a=\text{Pattern})} / \left[e^{\beta V(s,a=\text{Pattern})} + e^{\beta V(s,a=\text{Number 10})}\right]$$

where $P(a\ Pattern/s)$ is the probability of choosing the pattern option given presentation of a specific pattern "s", and "β" is the exploit-

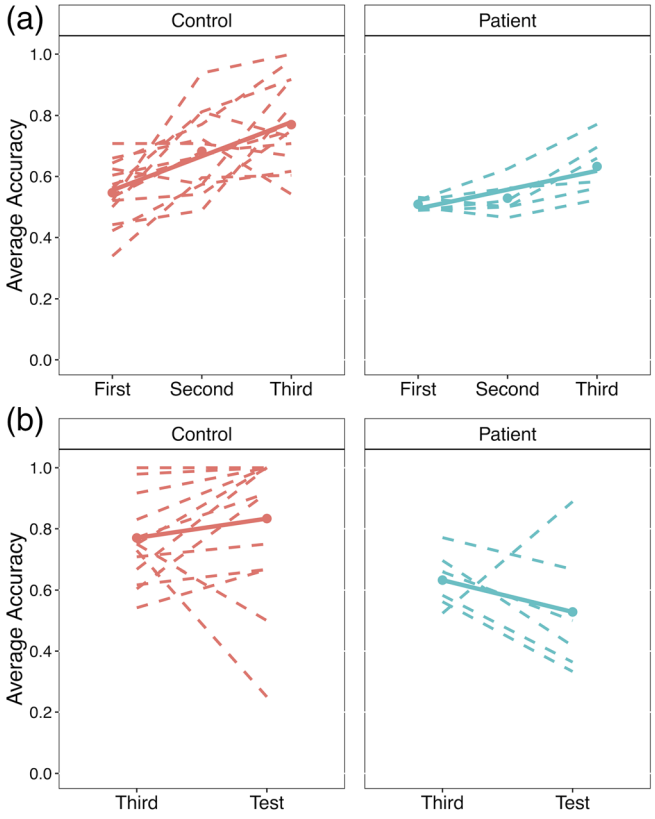

FIGURE 3 (a) Average accuracy based on majority status for all blocks. Each individual participant is shown with a dashed line. (b) Average accuracy based on majority status for the third learning block (left) and test phase (right) [Color figure can be viewed at wileyonlinelibrary.com]

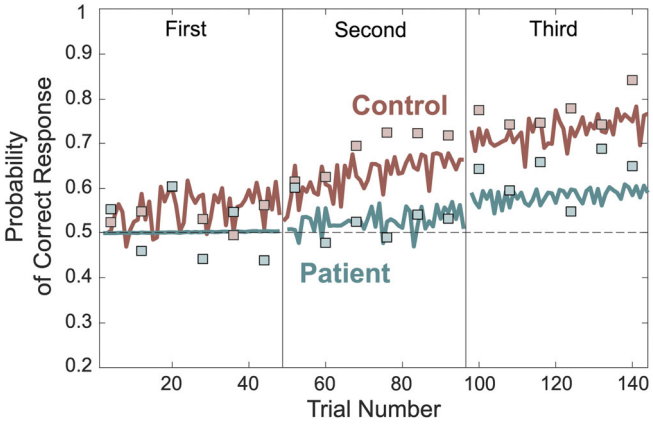

FIGURE 4 Trial-by-trial RL modeled probability of correct response calculated using best fit α and β parameters for each participant and each block, where "correct" refers to the majority status of the patterns. Curves represent group means. They indicate enhanced engagement of learning processes in controls compared to patients during the first block, but with probability of correct responses that are noisy and on average remain close to 0.5 for both groups. The greater accuracy in the control group compared to the patient group becomes more pronounced in the second and third blocks. Proportion of correct responses, averaged over bins of eight trials for each group are also presented. These data points were not directly fit by the model and are only included for illustrative purposes. Calculations of the R^2 coefficient of determination compared to a straight line at chance suggested that the model accounted for 45% (Block 1), 89% (Block 2), and 96% (Block 3) of additional variance in proportion of correct responses for control participants and for 8% (Block 1), 74% (Block 2), and 88% (Block 3) of additional variance for patients [Color figure can be viewed at wileyonlinelibrary.com]

TABLE 2 α and β values for patients and controls

	α			β		
	First	Second	Third	First	Second	Third
Patients						
Median[a]	0.02	0.05	0.06	0.2	0.2	1.4
Range[a]	0.02–0.03	0.03-0.59	0.03–0.12	0.20–0.60	0.20–2.40	0.20–2.80
Controls						
Median	0.05	0.08	0.12	1.1	2.1	2.3
Range	0.01–0.97	0.02–0.65	0.01–0.66	0.20–2.40	0.20–7.00	0.20–15.20

[a]As P6 was a statistical outlier here, for completeness, α and β are also reported without this patient: (Median[range]): α (first: 0.02 [0.02–0.03]; second: 0.4[0.03–0.07]; third: 0.05 [0.03–0.06]), β (first: 0.2 [0.20–0.60]; second: 0.02 [0.20–0.20]; third: 1.0 [0.20–2.80]).

explore parameter (i.e., larger β corresponds to more systematic choice of the option with greater expected reinforcing value, and smaller β corresponds to more random choice behavior). Parameters α and β were estimated for each participant and each block with Bayesian inference, using an implementation of the affine invariant ensemble Markov Chain Monte Carlo sampler of Goodman and Weare (2010) proposed by Foreman-Mackey, Hogg, Lang, and Goodman (2013) and computed in Matlab by Grinsted (2015). Flat priors were used, spanning the intervals (0–1) and (0–20) for α and β, respectively; and final estimates of parameters were selected as the maximum of the posterior distributions.

Nonparametric comparison (Mann–Whitney *U*) of estimated parameters revealed significant group differences during the first block for both α (U = 14, p = 0.02) and β (U = 15.5, p = .03; see Table 2). Similar patterns of smaller median α and β values in patients were found in the second and third blocks, albeit these did not reach significance for α (second U = 40.5, p = .90; third U = 27.5, p = .24) and were significant and marginally significant for β (second U = 11.5, p = .009, third U = 19.5, p = .06). These results indicate less updating and more random responding in patients compared to controls, a pattern that is evident by the first block (also see Figure 4).[3]

Overall, these results provide evidence that the MTL is necessary for value-based probabilistic learning and contributes to a growing body of literature suggesting that the role of the MTL in learning and cognition goes beyond memory for singular episodes. These results not only complement prior imaging work demonstrating MTL recruitment during a similar task (e.g., Li et al., 2011), they also extend this study by pointing to the necessity of the MTL for task performance. Critically, as the task used here had a simple feedback structure relative to our prior study, wherein the feedback to the participant and the value of the stimuli were orthogonal and likely required integration of the two (Palombo et al., 2019), the present findings rule out the possibility that the impairment in that study was simply due to this task demand.

How does the hippocampus support probabilistic learning? In our prior study (Palombo et al., 2019), we discussed the possibility that the hippocampus plays a more domain-general role in stimulus-outcome mappings via its computational role in pattern separation (Leutgeb, Leutgeb, Moser, & Moser, 2007). In keeping with this conceptualization, Ballard et al. (2018) recently showed that the hippocampus forms separable conjunctive representations, which, in turn, facilitate value-based learning via striatally mediated prediction errors. Such data suggest that the striatum and hippocampus play a synergistic role in cognitive processing. However, we deliberately avoided the use of conjunctive stimuli, opting instead for simple patterns as cues. A recent study by Seib et al. (2020) similarly showed impaired reward-based probabilistic learning using very simple cues in transgenic mice lacking adult hippocampal neurogenesis. Akin to our study, the demands on pattern separation in that study are not obvious.

Instead, the hippocampus may play a more fundamental role when learning entails value or reward information. Within this broader framework, one possibility is that the hippocampus is involved in learning the inherent value of a stimulus, namely in acquiring stimulus-value associations. A second albeit not mutually exclusive possibility is that the hippocampus is involved in reinforcement learning when reinforcement involves reward. The current data cannot adjudicate between these possibilities as our task involves value learning (i.e., whether a stimulus has a value greater or smaller than "10") and reward-based reinforcement. However, a prior fMRI study by Dickerson et al. (2011) showed hippocampus engagement in a similar value learning task (whether the value of a stimulus was greater or smaller than "5") where reinforcement did not entail reward. Such data favor the possibility that the hippocampus is involved in learning inherent stimulus value. Also consistent with this idea, in our own prior fMRI work (Palombo et al., 2019), where participants learned whether or not players won money, we showed that the hippocampus was more strongly engaged for correct versus incorrect trials, but not for rewarded versus nonrewarded trials. On the other hand, fMRI data from Delgado, Nystrom, Fissell, Noll, and Fiez (2000) support the possibility that the hippocampus is implicated when there is a reward component to reinforcement. In that study, MTL activation correlated with reward (i.e., winning after a choice) even though the task placed no demands on learning the value of a stimulus. Further research is needed to elucidate which aspect of our task is critical for hippocampal involvement.

There is now compelling evidence for synergistic communication between the hippocampus and basolimbic dopamine system. For instance in the domain of episodic memory, reward motivation boosts learning (Adcock, Thangavel, Whitfield-Gabrieli, Knutson, & Gabrieli, 2006; Madan, Fujiwara, Gerson, & Caplan, 2012; Murty & Adcock, 2014; Wolosin, Zeithamova, & Preston, 2012), largely, but not exclusively via VTA-hippocampal mechanisms (see Murty & Dickerson, 2017 for review). In a similar vein, in the context of the current task, the hippocampus may play a modulatory role in gating the phasic activity of dopaminergic neurons (e.g., Lodge & Grace, 2006), thus potentially scaling the efficacy of probabilistic reward learning.

In light of this possibility, it is interesting to note that the hippocampus does not appear necessary for probabilistic learning that does not entail value or reward information. Indeed, in a study of probabilistic learning by Foerde, Race, Verfaellie, and Shohamy (2013) amnesic patients learned mono-coloured butterfly flower associations as well as controls. Although the dissociation in amnesic performance between reinforcement learning with versus without a value or reward component is intriguing, given other task differences, it will be important to directly compare performance in amnesic patients in the same paradigm to illuminate this issue.

Alternatively, one might question whether our results might be due to the contributions of episodic memory to performance of control participants. This argument was advanced by Knowlton, Squire, and Gluck (1994) to explain their findings in the Weather Prediction Task in which participants must learn to classify multidimensional stimuli. Amnesic patients performed normally earlier in learning (i.e., the first 50 trials), but were impaired with further training (also see Knowlton et al., 1996), a pattern which was ascribed to the use of declarative strategies by controls as learning progressed.

Although our results analyzed by block are qualitatively similar, the RL modelling results provide evidence that even early in learning, amnesic patients are not updating as efficiently as controls. Albeit not definitive, these data do not compel an episodic memory explanation for our results.

In sum, we provide evidence for MTL involvement in simple probabilistic learning. Although the putative role of the hippocampus in probabilistic and value or reward-related learning remains to be clarified, our results, coupled with other recent study in this area, call into question the classic view of orthogonal memory systems, and demonstrate that the hippocampus supports broader aspects of learning than previously appreciated.

ACKNOWLEDGMENTS

Mieke Verfaellie was supported by a Senior Research Career Scientist Award from the Clinical Science Research and Development Service, US Department of Veterans Affairs. Daniela J. Palombo was supported by a Discovery Grant from NSERC, the John R. Evans Leaders Fund from the Canadian Foundation for Innovation, and an Insight Development Grant from SSHRC. The authors thank Alessandra Te for research assistance. The contents of this manuscript do not represent the view of the US Department of Veterans Affairs or the US Government.

ENDNOTES

[1] A similar pattern of results was observed with a non-parametric approach (first block, $U = 27.0$, $p = .24$; second block, $U = 10.0$, $p = .006$; third block, $U = 17.5$ $p = .04$).

[2] A similar pattern of results was observed with a non-parametric approach ($U = 13.5$, $p = .015$).

[3] As pointed out in previous work (Rutledge et al., 2009), α and β are not fully independent, and parameter comparison across groups can be difficult to interpret when both parameters are free. In the present work, we also tried RL models with fixed beta across groups (i.e., $\beta = 1$ and $\beta = 2$). These models resulted in slightly greater group differences in α, but without changing the pattern of significance of the results reported above.

DATA AVAILABILITY STATEMENT

The data that support the findings of this study are available on request from the corresponding author. The data are not publicly available due to privacy or ethical restrictions.

ORCID

Daniela J. Palombo https://orcid.org/0000-0001-8082-3522

REFERENCES

Adcock, R. A., Thangavel, A., Whitfield-Gabrieli, S., Knutson, B., & Gabrieli, J. D. (2006). Reward-motivated learning: Mesolimbic activation precedes memory formation. *Neuron*, *50*(3), 507–517. https://doi.org/10.1016/j.neuron.2006.03.036

Ballard, I., Wagner, A. D., & McClure, S. M. (2018). Hippocampal pattern separation supports reinforcement learning. *Nature Communications*, *10*, 1073. https://doi.org/10.1038/s41467-019-08998-1

Delgado, M. R., Nystrom, L. E., Fissell, C., Noll, D. C., & Fiez, J. A. (2000). Tracking the hemodynamic responses to reward and punishment in the striatum. *Journal of Neurophysiology*, *84*(6), 3072–3077. https://doi.org/10.1152/jn.2000.84.6.3072

Dickerson, K. C., Li, J., & Delgado, M. R. (2011). Parallel contributions of distinct human memory systems during probabilistic learning. *NeuroImage*, *55*(1), 266–276. https://doi.org/10.1016/j.neuroimage.2010.10.080

Duncan, K., Doll, B. B., Daw, N. D., & Shohamy, D. (2018). More than the sum of its parts: A role for the hippocampus in configural reinforcement learning. *Neuron*, *98*(3), 645-657.e6. doi:https://doi.org/10.1016/j.neuron.2018.03.042

Floresco, S. B., Todd, C. L., & Grace, A. A. (2001). Glutamatergic afferents from the hippocampus to the nucleus accumbens regulate activity of ventral tegmental area dopamine neurons. *Journal of Neuroscience*, *21*(13), 4915–4922. https://doi.org/10.1523/jneurosci.21-13-04915.2001

Foerde, K., Race, E., Verfaellie, M., & Shohamy, D. (2013). A role for the medial temporal lobe in feedback-driven learning: Evidence from amnesia. *Journal of Neuroscience*, *33*(13), 5698–5704. https://doi.org/10.1523/jneurosci.5217-12.2013

Foreman-Mackey, D., Hogg, D. W., Lang, D., & Goodman, J. (2013). emcee: The MCMC hammer. *Publications of the Astronomical Society of the Pacific*, *125*(925), 306–312. https://doi.org/10.1086/670067

Gasbarri, A., Packard, M. G., Campana, E., & Pacitti, C. (1994). Anterograde and retrograde tracing of projections from the ventral tegmental area to the hippocampal formation in the rat. *Brain Research Bulletin*, *33*(4), 445–452. https://doi.org/10.1016/0361-9230(94)90288-7

Goodman, J., & Weare, J. (2010). Ensemble samplers with affine invariance. *Communications in Applied Mathematics and Computational Science*, *5*(1), 65–80. https://doi.org/10.2140/camcos.2010.5.65

Grinsted, A. (2015). Grinsted/gwmcmc [Matlab]. https://github.com/grinsted/gwmcmc

Groenewegen, H. J., der Zee, E. V.-V., te Kortschot, A., & Witter, M. P. (1987). Organization of the projections from the subiculum to the ventral striatum in the rat. A study using anterograde transport of *Phaseolus vulgaris* leucoagglutinin. *Neuroscience*, *23*(1), 103–120. https://doi.org/10.1016/0306-4522(87)90275-2

Hopkins, R. O., Myers, C. E., Shohamy, D., Grossman, S., & Gluck, M. (2004). Impaired probabilistic category learning in hypoxic subjects with hippocampal damage. *Neuropsychologia*, *42*(4), 524–535. https://doi.org/10.1016/j.neuropsychologia.2003.09.005

Kahn, I., & Shohamy, D. (2013). Intrinsic connectivity between the hippocampus, nucleus accumbens, and ventral tegmental area in humans. *Hippocampus*, *23*(3), 187–192. https://doi.org/10.1002/hipo.22077

Kan, I. P., Giovanello, K. S., Schnyer, D. M., Makris, N., & Verfaellie, M. (2007). Role of the medial temporal lobes in relational memory: Neuropsychological evidence from a cued recognition paradigm. *Neuropsychologia*, *45*(11), 2589–2597. https://doi.org/10.1016/j.neuropsychologia.2007.03.006

Knowlton, B. J., Mangels, J. A., & Squire, L. R. (1996). A neostriatal habit learning system in humans. *Science*, *273*(5280), 1399–1402. https://doi.org/10.1126/science.273.5280.1399

Knowlton, B. J., Squire, L. R., & Gluck, M. A. (1994). Probabilistic classification learning in amnesia. *Learning & Memory*, *1*(2), 106–120. doi:https://doi.org/10.1101/lm.1.2.106

Lee, H., Ghim, J. W., Kim, H., Lee, D., & Jung, M. (2012). Hippocampal neural correlates for values of experienced events. *Journal of Neuroscience*, *32*(43), 15053–15065. https://doi.org/10.1523/jneurosci.2806-12.2012

Leutgeb, J. K., Leutgeb, S., Moser, M.-B., & Moser, E. I. (2007). Pattern separation in the dentate gyrus and CA3 of the hippocampus. *Science (New York, N.Y.)*, *315*(5814), 961–966. doi:https://doi.org/10.1126/science.1135801

Li, J., Delgado, M. R., & Phelps, E. A. (2011). How instructed knowledge modulates the neural systems of reward learning. *Proceedings of the National Academy of Sciences*, *108*(1), 55–60. https://doi.org/10.1073/pnas.1014938108

Lodge, D. J., & Grace, A. A. (2006). The hippocampus modulates dopamine neuro nresponsivity by regulating the intensity of phasic neuron activation. *Neuropsychopharmacology*, *31*(7), 1356–1361. https://doi.org/10.1038/sj.npp.1300963

Madan, C. R., Fujiwara, E., Gerson, B. C., & Caplan, J. B. (2012). High reward makes items easier to remember, but harder to bind to a new temporal context. *Frontiers in Integrative Neuroscience*, *6*, 61. https://doi.org/10.3389/fnint.2012.00061

Murty, V. P., & Adcock, R. A. (2014). Enriched encoding: Reward motivation organizes cortical networks for hippocampal detection of unexpected events. *Cerebral Cortex*, *24*(8), 2160–2168. https://doi.org/10.1093/cercor/bht063

Murty V. P., Dickerson K. C. (2017). Motivational influences on memory, *Advances in Motivation and Achievement* (pg. 203–227). Bingley, UK: Emerald Group Publishing Limited.

Palombo, D. J., Hayes, S. M., Reid, A. G., & Verfaellie, M. (2019). Hippocampal contributions to value-based learning: Converging evidence from fMRI and amnesia. *Cognitive, Affective, & Behavioral Neuroscience*, *19*(3), 523–536. https://doi.org/10.3758/s13415-018-00687-8

Rutledge, R. B., Lazzaro, S. C., Lau, B., Myers, C. E., Gluck, M. A., & Glimcher, P. W. (2009). Dopaminergic drugs modulate learning rates and perseveration in Parkinson's patients in a dynamic foraging task. *Journal of Neuroscience*, *29*(48), 15104–15114. https://doi.org/10.1523/jneurosci.3524-09.2009

Schonberg, T., O'Doherty, J. P., Joel, D., Inzelberg, R., Segev, Y., & Daw, N. D. (2010). Selective impairment of prediction error signaling in human dorsolateral but not ventral striatum in Parkinson's disease patients: Evidence from a model-based fMRI study. *NeuroImage*, *49*(1), 772–781. https://doi.org/10.1016/j.neuroimage.2009.08.011

Seib, D. R., Espinueva, D. F., Floresco, S. B., & Snyder, J. S. (2020). A role for neurogenesis in probabilistic reward learning. *Behavioral Neuroscience*, *134*(4), 283–295. https://doi.org/10.1037/bne0000370

Shohamy, D., Myers, C. E., Kalanithi, J., & Gluck, M. A. (2008). Basal ganglia and dopamine contributions to probabilistic category learning. *Neuroscience and Biobehavioral Reviews*, *32*(2), 219–236. https://doi.org/10.1016/j.neubiorev.2007.07.008

Squire, L. R. (1987). *Memory and brain*. New York, NY: Oxford University Press.

Sutton, R. S., & Barto, A. G. (2018). *Reinforcement learning: An introduction*. London, UK: MIT Press.

Vanni-Mercier, G., Mauguière, F., Isnard, J., & Dreher, J.-C. (2009). The hippocampus codes the uncertainty of cue–outcome associations: An intracranial electrophysiological study in humans. *Journal of Neuroscience*, *29*(16), 5287–5294. https://doi.org/10.1523/jneurosci.5298-08.2009

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), 1532–1547. https://doi.org/10.1162/jocn_a_00237

How to cite this article: Palombo DJ, Patt VM, Hunsberger R, Verfaellie M. Probabilistic value learning in medial temporal lobe amnesia. *Hippocampus*. 2021;31:461–468. https://doi.org/10.1002/hipo.23317