



Published in final edited form as:

Neuron. 2013 July 10; 79(1): 3–6. doi:10.1016/j.neuron.2013.06.033.

How did the chicken cross the road? With her striatal cholinergic interneurons, of course

Geoffrey Schoenbaum^{1,2,3}, Thomas A Stalnaker¹, and Yael Niv⁴

¹Behavioral Neuroscience Branch, NIDA/IRP/NIH/DHHS, Baltimore, MD

²Department of Anatomy and Neurobiology, University of Maryland, School of Medicine, 20 Penn St, HSF-2 S251, Baltimore, MD

³Solomon Snyder Department of Neuroscience, Johns Hopkins University, Baltimore, MD

⁴Department of Psychology & Princeton Neuroscience Institute, Princeton University, Princeton, NJ

Abstract

Recognizing when the world changes is fundamental for normal learning. Here, Bradfield and colleagues show that cholinergic interneurons in dorsomedial striatum are critical to the process whereby new states of the world are appropriately registered and retrieved during associative learning.

Recognizing when the world has changed—and when it has not—is a fundamental yet much ignored component of associative learning. Imagine relocating to Sydney, Australia. While much there might be familiar, one prominent difference is of life-or-death import: the cars come from the right. If you don't learn to look right-left-right before crossing, your visit might be quite short. On the other hand, since you plan to venture to proper-side-of-the-road-driving countries periodically, it would behoove you to also maintain your previous left-right-left behavior, applying that when appropriate. Optimally, rather than over-writing your original strategy for crossing the street, upon experiencing the strange driving habits in your new home town, you would form a new 'state' of 'I am in Sydney' and learn new mappings from actions to goals ('policies' in the jargon of reinforcement learning, 'action-outcome associations' in terms of learning theory) relevant to that state. Linking these learned policies to the new state would, conveniently, protect the old policies linked to the old state from being overwritten, so that behavior could be modified quickly if the old state were to re-appear.

As this example illustrates, appropriate recognition of when to form new states to which to attach information is vital to adaptive behavior. In this issue of *Neuron*, Bradfield and colleagues (Bradfield et al., 2013) use a series of complex yet highly controlled behavioral manipulations to show that input from a part of the thalamus, the parafascicular nucleus, onto cholinergic interneurons in the posterior compartment of the dorsomedial striatum (pDMS), is critical to the appropriate creation of new states during learning. We use 'state'

© 2013 Elsevier Inc. All rights reserved.

*Address correspondence to Geoffrey Schoenbaum (geoffrey.schoenbaum@nih.gov) and Yael Niv (yael@princeton.edu).

Publisher's Disclaimer: This is a PDF file of an unedited manuscript that has been accepted for publication. As a service to our customers we are providing this early version of the manuscript. The manuscript will undergo copyediting, typesetting, and review of the resulting proof before it is published in its final citable form. Please note that during the production process errors may be discovered which could affect the content, and all legal disclaimers that apply to the journal pertain.

here to refer to a high order representation of the environment in which actions are being chosen—a notion that encompasses the animal learning theory terms of ‘context’, ‘discriminative stimulus’ and ‘occasion setter’ as well as the statistical learning theory term ‘latent cause’ (Gershman and Niv, 2010), but is different from common usage of the term in reinforcement learning.

In the first phase of training, Bradfield et al. taught rats to associate two levers with two different, but equally valued, rewards (pellets or sucrose). Satiating the rats on one of the two outcomes (a so-called ‘devaluation test’) selectively reduced responding on the lever leading to that outcome and not on the lever leading to the other outcome. Notably, this was the case both for intact rats and for rats in which cholinergic signaling in the pDMS, an area previously shown to be necessary for goal-directed behavior (Yin et al., 2005), was disrupted via several different manipulations (Figure 1; left). This intact initial learning, and sensitivity to devaluation, demonstrated that behavior was indeed goal-directed (i.e., guided by an expectation of the specific outcome and its unique attributes (Dickinson and Balleine, 1994)) and that intact cholinergic activity in the pDMS is *not* necessary for this fundamental learning process.

But what happens if the world changes? This was tested in a second phase of training, in which rats faced three challenges (Figure 1; top), each designed to test how changes in the associative structure of the environment would be incorporated into the earlier learning. The first involved contingency degradation – the outcome associated with one of the levers was presented for free, meaning that rats no longer needed to work to receive that reward. The second, reversal learning, involved switching the outcomes associated with each lever, followed by another devaluation test of the effects of satiety on responding. The third, extinction, involved removal of all outcomes for actions, followed by a ‘reinstatement test’ in which one of the outcomes was delivered to test whether it could reinstate pressing on the lever most recently associated with that outcome. In each challenge, the critical question was whether rats would appropriately create new states in which to represent the new environmental contingencies. If so, each challenge should selectively affect responding on the lever most recently associated with the degraded, devalued or reinstated outcome. Any non-specific effects on both levers would suggest that something had gone awry.

The results were amazingly clear-cut: in each case, intact rats exhibited selective effects on subsequent testing (Figure 1; middle), whereas rats in which cholinergic signaling in the pDMS was disrupted showed intermediate or non-specific effects on testing (Figure 1; bottom). Importantly, this same behavioral pattern was induced by bilateral lesions of the parafascicular nucleus, crossed lesions of the parafascicular nucleus and the pDMS, or a pharmacological manipulation that disrupted cholinergic signaling in the pDMS only during the learning phases. This exhaustive characterization of the phenomenon shows both that it is reliable, and that it depends on cholinergic signaling *at the time of learning*, with the latter explicitly confirmed using an immunohistochemistry tool specific to cholinergic interneurons that was recently developed by the Balleine lab (Bertran-Gonzalez et al., 2012). In addition, these results were shown to be specific to cholinergic disruption in the pDMS as they were not reproduced by manipulations of cholinergic function in the anterior portion of the dorsomedial striatum.

What does all this mean?

While these results are extremely elegant in their consistency and convergence, they are not entirely straightforward to interpret theoretically. What exactly has gone wrong with the state generation process due to the cholinergic manipulations? Here the comprehensive set of metaphoric hoops through which the rats were made to jump becomes key to narrowing

down the options, highlighting the utility of using the incisive behavioral manipulations that animal learning theorists have spent decades developing.

To understand what went wrong, it is useful to first review what aspects of learning were not disrupted by cholinergic manipulations: in addition to intact goal-directed learning, the comprehensive battery of tests shows that new state formation was not completely abolished. This is evident in the test following the third challenge, extinction training, in which exposure to one of the outcomes led to reinstatement of responding. Reinstatement indicates that the extinction did not simply overwrite and erase previous associations between actions and outcomes (Gershman et al., 2010), but rather reward omission caused rats in both groups to create a new state (Figure 1; state 4). However, reinstatement in cholinergically-impaired rats was far from normal: these rats reinstated *both* actions (Figure 1; right).

Retrieval deficit

One possible explanation for this pattern of results is that upon reinstatement the rats erroneously retrieved two states – the most recent, post-reversal state (state 3 in Figure 1), in which the right lever was mapped to sucrose and the left to pellets, and the state from initial training (state 1 in Figure 1), in which the lever to reward mapping was the reverse. This may, in fact, sound familiar to world travelers: a foolproof policy for safe street-crossing in some countries is to look left-right-left-right repeatedly, that is, to act upon *both* pre-travel and in-travel states.

Such a retrieval deficit could also explain the lack of specificity of the post-reversal devaluation test, in which cholinergically-impaired rats devalued *both* actions rather than only the one associated with the satiated outcome (Figure 1; third column). Finally, it can also explain the intermediate level of responding in the contingency degradation test (Figure 1; second column) by assuming that the new state (state 2, in which not pressing was associated with the outcome) was retrieved together with the old state (state 1).

The deficit in reinstatement was observed even when cholinergic function was disrupted only during learning, yet this does not rule out a retrieval deficit, as retrieval of the appropriate states is also necessary during learning. That is, in order to learn, on every trial, the rat must retrieve and update the association within the relevant state. If multiple states were retrieved and updated during learning, the rat would show a non-selective response in the reinstatement test even though normal cholinergic function had been restored. Importantly, under this interpretation, new state formation is intact, however, retrieval of appropriate states is disrupted or at least less selective.

State creation deficit

A second possible explanation is that the rats with disrupted cholinergic function might have been able to form a new state in extinction but not in the other challenges. Why would this happen? To answer this, it is useful to ask how does the brain know that a new state should be formed in the first place. One impetus for state creation is significant differences between the current situation and past experience (Gershman et al., 2010). According to this idea, prediction errors— differences between what is expected (driving is on the right of the road, mass transportation is called ‘subway’, etc.) and what is currently experienced (cars are on the left, the underground train is ‘the city circle’)—drive state formation. Importantly, these prediction errors include both errors in predicted value (the city circle is not cheap), and errors in predicted identity (would you expect ‘the city circle’ to indicate an underground train system?). The former are typically termed reward prediction errors (though we use ‘value’, as changes in rewarding events can also induce identity prediction errors), and

Bradfield et al. refer to the latter as ‘state prediction errors’, though we prefer ‘identity’ as any sort of error could lead to recognition of state change.

Bradfield et al.’s first two manipulations—contingency degradation and reversal learning—involved only identity prediction errors, since the underlying value of the reward associated with lever pressing did not change. However, the last manipulation introduced value prediction errors since the reward was entirely omitted. If cholinergic transmission in the striatum is important for detecting, representing, or learning from identity prediction errors, one would expect to see no new state formation in the first two manipulations due to the cholinergic manipulation, but intact state formation during extinction learning. Thus, like a retrieval deficit, a selective effect on the formation of new states following identity prediction errors would also produce the observed pattern of results (Figure 1 bottom).

What’s Ach got to do with it?

Though relatively little is known about the function of cholinergic striatal interneurons, what we know so far relates nicely to these two interpretations. For example, one can easily imagine a key role for striatal acetylcholine in retrieval: Cholinergic interneurons are inhibitory, tonically active, and innervate (and receive input from) a large number of medium spiny neurons (Zhou et al., 2002). This places this local modulatory system in a prime position to provide network-wide inhibition, promoting retrieval of only the relevant state at each point in time (Apicella, 2007). By reducing cholinergic tone, Bradfield et al. could have thus caused rats to retrieve multiple states during decision making and learning, thereby supporting the first interpretation above. On the other hand, cholinergic interneurons also respond to important events with phasic changes in firing that are notably unrelated to value prediction errors (Morris et al., 2004). Do these responses relate instead to identity prediction errors? This has yet to be tested, and would support the second interpretation.

However, even without complete understanding of the striatal circuitry and its reliance on acetylcholine, the powerful toolkit provided by traditional animal learning theory could be used to test and differentiate the above two hypotheses. One key experiment would be to train rats to associate the two levers with rewards of decidedly different magnitude and then put them through Bradfield et al.’s series of tests. If the deficit depended on the need to learn from identity prediction errors, behavior should now be impervious to cholinergic interventions in the pDMS, since all three manipulations would involve value as well as identity prediction errors. If, on the other hand, the problem was one of retrieval, then the rats’ responding should still reflect the erroneous association of both levers with both outcomes. Of course, single unit recordings would still be useful for understanding the relationship between either of these roles and the precise firing patterns of the neurons, as well as the dynamics of learning in the striatal network that gives rise to these functions (and associated deficits). However, it is always inspiring to see well-controlled behavioral designs reveal underlying neural processes even absent electrodes.

References

- Apicella P. Leading tonically active neurons of the striatum from reward detection to context recognition. *Trends in Neurosciences*. 2007; 30:299–306. [PubMed: 17420057]
- Bertran-Gonzalez J, Chieng B, Laurent V, Valient E, Balleine BW. Striatal cholinergic interneurons display activity-related phosphorylation of ribosomal protein S6. *PLoS ONE*. 2012; 7:e53195. [PubMed: 23285266]
- Bradfield LA, Bertran-Gonzalez J, Chieng B, Balleine BW. The thalamo-striatal pathway and cholinergic control of goal-directed action: interlacing new and existing action-outcome associations in the striatum. *Neuron*. 2013

- Dickinson A, Balleine BW. Motivational control of goal-directed action. *Animal Learning and Behavior*. 1994; 22:1–18.
- Gershman SJ, Blei DM, Niv Y. Time, context and extinction. *Psychological Review*. 2010; 117:197–209. [PubMed: 20063968]
- Gershman SJ, Niv Y. Learning latent structure: carving nature at its joints. *Current Opinion in Neurobiology*. 2010; 20:251–256. [PubMed: 20227271]
- Morris G, Arkadir D, Nevet A, Vaadia E, Bergman H. Coincident but distinct messages of midbrain dopamine and striatal tonically active neurons. *Neuron*. 2004; 43:133–143. [PubMed: 15233923]
- Yin HH, Ostlund SB, Knowlton BJ, Balleine BW. The role of the dorsomedial striatum in instrumental conditioning. *European Journal of Neuroscience*. 2005; 22:513–523. [PubMed: 16045504]
- Zhou F, Wilson C, Dani J. Cholinergic interneuron characteristics and nicotinic properties in the striatum. *Journal of Neurobiology*. 2002; 53:590–605. [PubMed: 12436423]

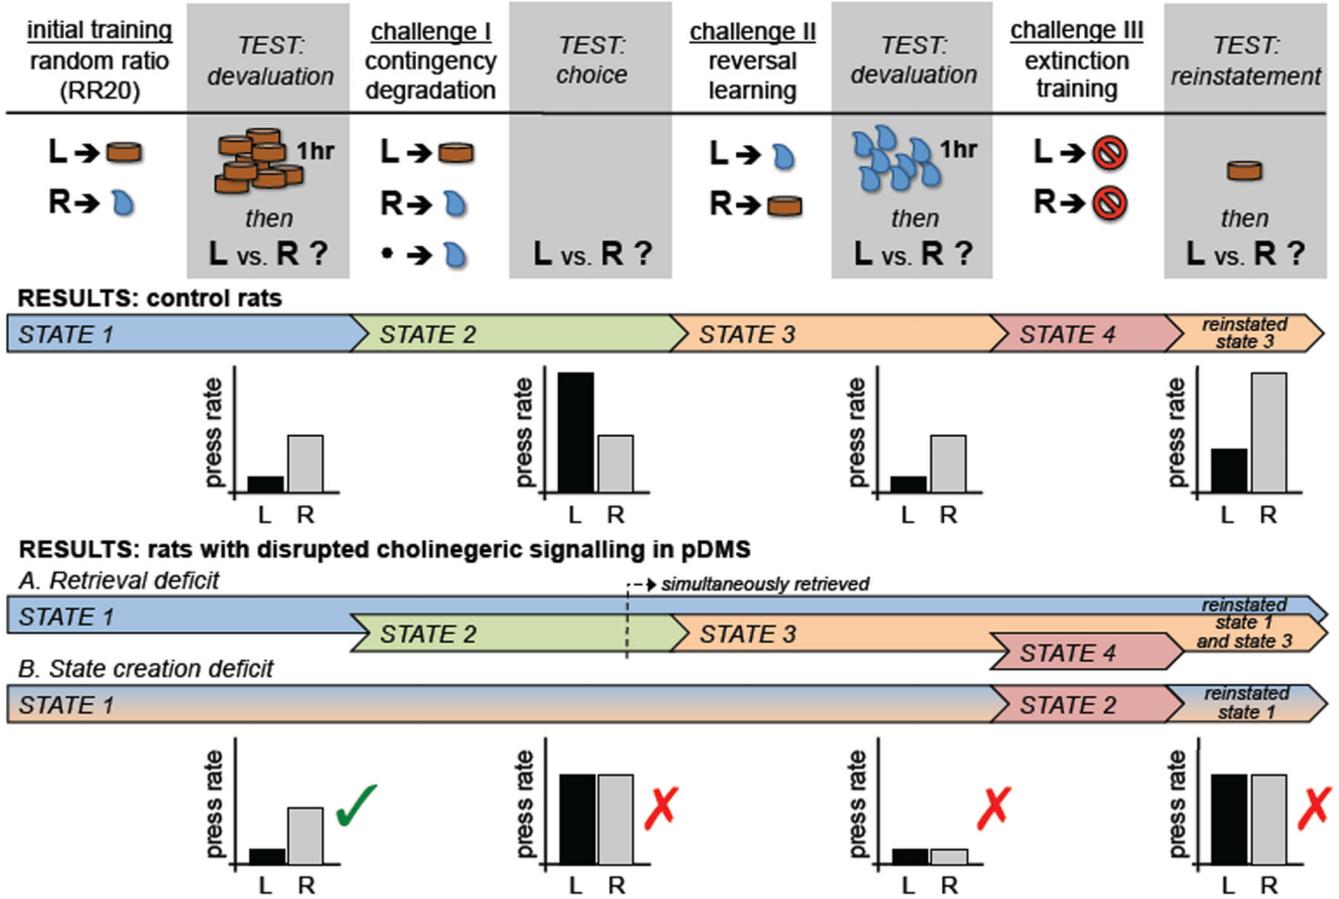


Figure 1. Illustration of Bradfield et al.’s (2013) experimental design, results & interpretation. Rewards for the two actions were a food pellet or sucrose solution (**Top**). All tests were conducted without rewards, retraining was administered after tests as necessary, • denotes no action, and all conditions were counterbalanced (see Bradfield et al., 2013). Selective test responding only on the lever that had previously been mapped to the non-devalued/non-degraded/reinstated outcome in the control group suggests that each change in task contingencies was encoded by the rats as a new state (**Middle**). In contrast, disruption of cholinergic activity in pDMS resulted in nonspecific degradation, devaluation and reinstatement of both actions, but only after initial learning (**Bottom**). This could be ascribed to (A) a retrieval deficit that causes multiple states to be retrieved throughout, or (B) a deficit in creating new states when the identity but not the value of outcomes is changed, such that training in challenges I and II is combined with the initial training.