

# Amnesia in healthy people via hippocampal inhibition: A new forgetting mechanism

Quarterly Journal of Experimental Psychology  
2024, Vol. 77(1) 1–13  
© Experimental Psychology Society 2023  
Article reuse guidelines:  
sagepub.com/journals-permissions  
DOI: 10.1177/17470218231202728  
qjep.sagepub.com



Michael C Anderson<sup>1,2</sup>  and S Subbulakshmi<sup>1,3</sup>

## Abstract

Structural damage to the hippocampus gives rise to a severe memory deficit for personal experiences known as organic amnesia. Remarkably, such structural damage may not be the only way of creating amnesia; windows of amnesia can also arise when people deliberately disengage from memory via a process known as retrieval suppression. In this review, we discuss how retrieval suppression induces systemic inhibition of the hippocampus, creating “amnesic shadow” intervals in people’s memory for their personal experiences. When new memories are encoded or older memories are reactivated during this amnesic shadow, these memories are disrupted, and such disruption even arises when older memories are subliminally cued. Evidence suggests that the systemic inhibition of the hippocampus during retrieval suppression that gives rise to the amnesic shadow may be mediated by engagement of hippocampal GABAergic inhibitory interneurons. Similar amnesic shadow effects are observed during working memory tasks like the n-back, which also induce notable hippocampal downregulation. We discuss our recent proposal that cognitive operations that require the disengagement of memory retrieval, such as retrieval suppression, are capable of *mnemonic process inhibition* (the inhibition of mnemonic processes such as encoding, consolidation, and retrieval and not simply individual memories). We suggest that people engage mnemonic process inhibition whenever they shift attention from internal processes to demanding perceptual-motor tasks that may otherwise be disrupted by distraction from our inner world. This hitherto unstudied model of inhibition is a missing step in understanding what happens when attentional shifts occur between internally and externally oriented processes to facilitate goal-directed behaviour. This process constitutes an important novel mechanism underlying the forgetting of life events.

## Keywords

Memory control; forgetting; amnesia; retrieval suppression; inhibition

Received: 16 May 2023; revised: 01 October 2023; accepted: 26 June 2023

Organic amnesia refers to a severe deficit in memory arising from damage to a brain structure involved in memory, such as the hippocampus (Cipolotti & Bird, 2006; Huppert & Piercy, 1979; Levin et al., 1983; Parkin, 1984; Shimamura, 1992). Perhaps the most famous case of organic amnesia is patient HM (Henry Molaison), who, as a young man, developed intractable epileptic seizures. HM’s hippocampi, thought to be the source of the electrical instability that generated these seizures, were removed in an experimental neurosurgery intended to ameliorate his symptoms. Removal of the hippocampi, however, rendered HM permanently amnesic; he did not encode another episodic memory again for the remaining six decades of his life (Huppert & Piercy, 1979; Scoville & Milner, 1957; Squire, 2009) (Figure 1). HM lost the ability to form new memories experienced after the surgery (anterograde

amnesia) and lost many memories stored in the years immediately prior to the surgery (retrograde amnesia). The discovery of such profound forgetting arising from such a focal form of brain damage was a watershed moment in the history of memory research. Over the years, this discovery has inspired a tremendous volume of research on


<sup>1</sup>MRC Cognition and Brain Sciences Unit, University of Cambridge, Cambridge, UK

<sup>2</sup>Behavioural and Clinical Neuroscience Institute, University of Cambridge, Cambridge, UK

<sup>3</sup>Department of Psychology, Stanford University, Stanford, CA, USA

## Corresponding author:

Michael C Anderson, MRC Cognition and Brain Sciences Unit, University of Cambridge, 15 Chaucer Road, Cambridge CB2 1TN, UK.  
Email: michael.anderson@mrc-cbu.cam.ac.uk



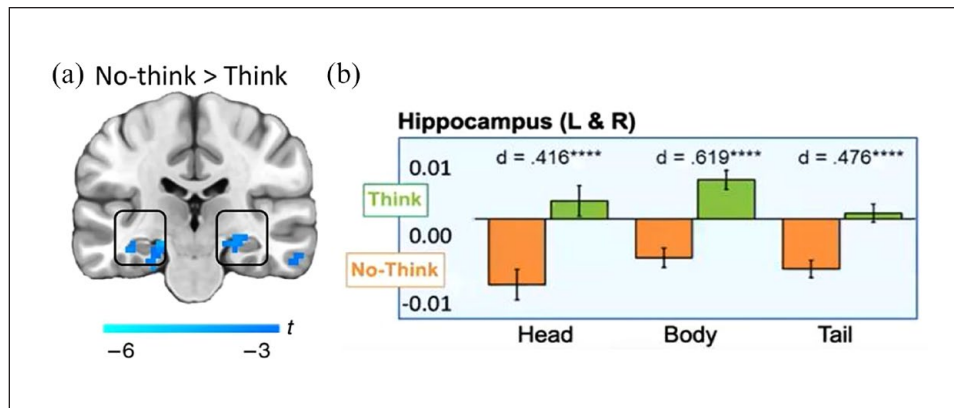
**Figure 1.** An MRI image of HM's brain, illustrating the absence of both of his hippocampi (darkened areas within white boxes). Source. Image reproduced from Corkin et al. (1997).

the role of the hippocampus in encoding, consolidating, and retrieving long-term memories, and on how hippocampal damage causes organic amnesia. This review, however, will not address amnesia arising from permanent structural damage to the hippocampus; rather, it will focus on amnesia arising in otherwise healthy individuals, from the normal operation of cognitive processes. Specifically, we propose that cognitive operations exist that have as their outcome the suppression of hippocampal activity and the disruption of its functional state. This disruption can induce amnesia much like that suffered by HM, but occurring in briefer windows of time.

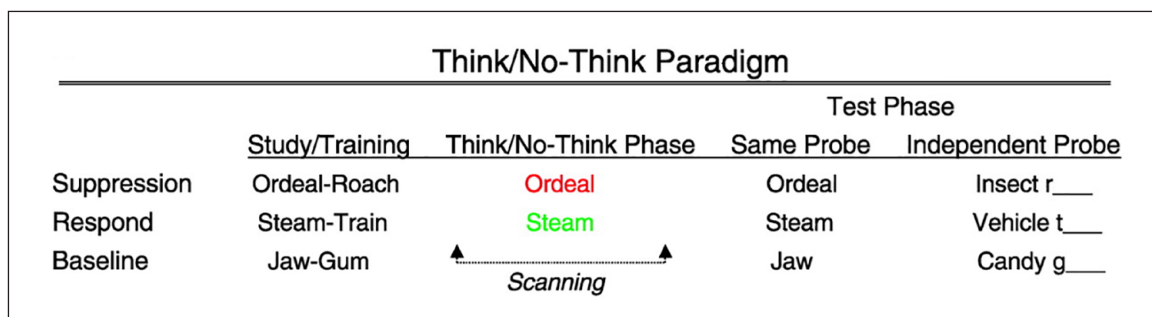
Although many cognitive processes might affect hippocampal activity, our work focuses on retrieval suppression. Retrieval suppression refers to the process of stopping retrieval of a long-term memory in the face of a reminder to that memory (Levy & Anderson, 2002). Although reminders often have the capacity to automatically elicit retrieval of an associated memory, people can willfully stop the memory from entering awareness even when focusing full attention on the reminder. Retrieval suppression is operationalised in the laboratory by the Think/No-Think (TNT) task (Anderson & Green, 2001). In this task, participants are presented with reminders to memories that they have encoded previously in the laboratory; on some trials they are asked to retrieve the associated memory (hereinafter, Think trials), and on other trials, to stop themselves from retrieving the memory (hereinafter, No-Think trials). A large body of evidence concerns the mnemonic after-effects of engaging in retrieval suppression (Anderson & Hanslmayr, 2014; Anderson & Huddleston, 2012; Marsh & Anderson, 2020). It has been found, for example, that repeated retrieval suppression makes it harder to recall the suppressed memories, leading to subsequent forgetting of these memories. This phenomenon is known as *suppression-induced forgetting*.

Of particular importance for the present review, however, is that suppressing retrieval not only disrupts memory for the suppressed content, but also downregulates hippocampal activity to achieve that result (Figure 2a). To illustrate retrieval–suppression's impact on hippocampal activity, consider an analysis of functional imaging data from 10 TNT studies (Figure 2b). As Figure 2b illustrates, whereas retrieving a memory upregulates hippocampal activity (in green), suppressing retrieval downregulates it (in orange), and this pattern arises throughout the anterior–posterior axis of the hippocampus. Thus, intentionally stopping episodic retrieval appears to suppress hippocampal function in a manner analogous to how intentionally stopping an action suppresses motor cortical activity (Apšvalka et al., 2022; Schmitz et al., 2017). In this review, we focus on a striking and counterintuitive question raised by this apparent downregulation: if retrieval stopping interrupts retrieval by suppressing hippocampal activity, might hippocampal function be globally disrupted, rendering it unable to perform its usual range of functions? Might such disruption mimic the characteristics of organic amnesia?

Why might retrieval–suppression induce a generalised amnesia? The answer lies in a possibility that one might not immediately consider when thinking about how retrieval stopping works. If one encounters a reminder to an unpleasant memory and one tries to push the unwelcome memory out of awareness, the intention is to clear awareness of that specific memory. One might imagine—and many researchers in experimental psychology have tacitly assumed—that an inhibitory process engaged for this purpose might act on that specific memory. But the foregoing brain imaging data revealing a hippocampus-wide downregulation suggests an alternative—that the suppression of a specific memory is achieved by a mechanism targeted at hippocampally mediated retrieval processes more broadly. Thus, while a person may intend to exclude a specific memory, the implementation of that intention may involve suppression acting systemically. Precedents for such a global inhibition exist. For example, in the stop signal task, stopping a motor action given a stop signal globally affects motor cortex, not the specific effector one intends to stop (J.R. Wessel & Aron, 2017). If retrieval suppression globally suppresses hippocampal activity, any process requiring the hippocampus ought to be interrupted. For example, if hippocampal suppression lingers after retrieval stopping has ended, any new experiences encountered afterwards ought to be less effectively encoded, as perceptual inputs arriving to the hippocampus may not result in trace formation. Anterograde amnesia for that event would result. On the contrary, hippocampally mediated memories encoded prior to suppression may require ongoing hippocampal processing to be stabilised and consolidated; suppression may interrupt this process, causing retrograde amnesia. Thus, memories for these temporally adjacent events would be lost, even though



**Figure 2.** (a) Hippocampal downregulation during No-Think trials (from Schmitz et al., 2017,  $N=24$ ). (b) Region of interest analysis of 10 TNT studies illustrating hippocampal downregulation during No-Think trials. Plotted is average activation in hand-traced hippocampi of 330 participants in the Think and No-Think conditions in the head, body, and tail of the hippocampus averaged over the left (L) and right (R) hippocampi.



**Figure 3.** The structure of the Think/No-Think procedure, using word pairs. Source. From Anderson et al. (2004).

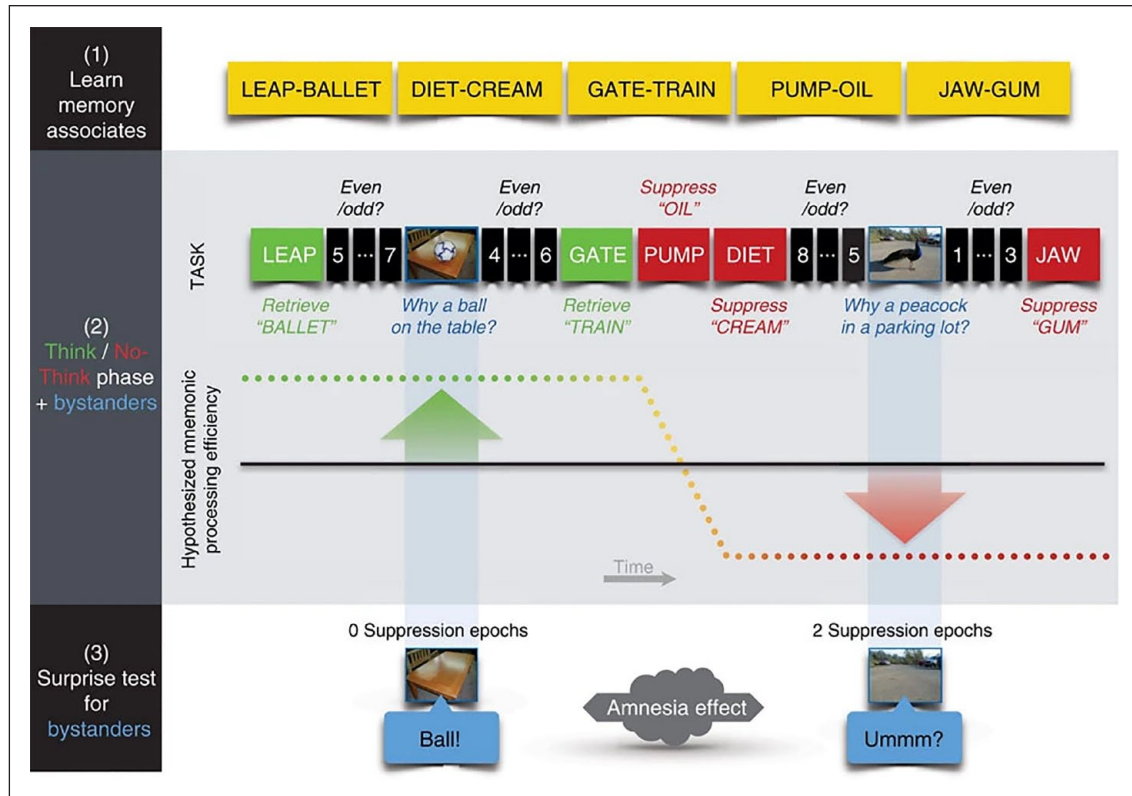
the retrieval suppression itself was not targeted towards them, and even if they are entirely unrelated to the content being suppressed. Thus, retrieval suppression may create a transient “virtual lesion” of the hippocampus that undermines its role in the formation or stabilisation of memories (Hulbert et al., 2016).

In this article, we review the work we have done to test the predictions of this systemic suppression hypothesis. We focus specifically on the idea that retrieval suppression may induce an *amnesic shadow* for hippocampally dependent memories that are encoded or reactivated in the temporal surround of suppression. First, we describe how we adapted to TNT task to test these predictions. We then review two series of studies that explore the existence and characteristics of the amnesic shadow for both newly encoded events are older memories that one is merely reminded of in the vicinity of suppression. We close by arguing that although we have discovered the amnesic shadow in the context of intentional retrieval suppression, it may be induced by a much wider range of conditions. Critically, we argue that this process constitutes an entirely

novel mechanism of forgetting from long-term memory that could be far more pervasive than researchers realise.

### Testing the existence of the amnesic shadow

We tested the existence of the amnesic shadow by adapting the TNT task. A typical TNT task consists of three phases (Anderson et al., 2004; Anderson & Green, 2001): a Study phase, a TNT phase, and a Recall phase (Figure 3). During the Study phase, participants learn several dozen paired associates. Generally, these paired associates could be word pairs, word–picture pairs, or picture–picture pairs; the suppression-induced forgetting effect has been observed with a variety of stimuli (for a review, see Anderson & Hulbert, 2021; Marsh & Anderson, 2020). After studying all the associated pairs, participants are tested for their knowledge of the pairs to ensure that they meet our learning standard before the TNT phase (typically, a minimum of 50% correct). Generally, all the experiments exploring the amnesic shadow used simple neutral



**Figure 4.** The modified Think/No-Think task for studying the amnesic shadow. Source. From Hulbert et al. (2016).

pairs of words as stimuli for the main TNT portion of the adapted amnesic shadow task.

After learning the pairs, participants enter the critical TNT phase, wherein they are required to control the retrieval process. During this task, participants perform trials in which they receive a cue from a previously studied pair and are asked to either retrieve the memory associated to it (if the cue appears in green), or to suppress retrieval of the associate (if the cue appears in red). Participants are instructed to focus on the reminders fully for both types of trials. Critically, during red (No-Think) trials, participants are instructed to stop themselves from retrieving the associate at all (not even for a second) and are asked to do this without generating distracting thoughts. For a given participant, particular items are either always suppressed or always retrieved, so are consistently cued with red or green. Targets are retrieved or suppressed multiple times (typically 12 times), to maximise the manipulation's effect. The TNT task ends with a Recall phase, wherein participants are tested on their memory for all items. During Recall, either the studied reminders, or novel "independent" probes are presented as cues (see Figure 3). These independent cues are related to the targets, but not to the learnt reminders (e.g., Insect R\_\_ to cue recall of "roach, studied with Ordeal).

A subset of the pairs is learned during the study phase, but not cued during the TNT phase. These items are

included to estimate how well people could recall pairs on the final test, given that they neither suppressed, nor retrieved the target associate in the interim, and so constitute the baseline used to determine whether suppressing items in the No-Think condition impairs memory. No-Think items are typically remembered more poorly than are the baseline items even though No-Think reminders are presented repeatedly during the TNT phase, whereas the baseline reminders are not. This below-baseline forgetting reflects the typical suppression-induced forgetting pattern (Anderson & Green, 2001; Anderson & Hulbert, 2021; Mamat & Anderson, 2023).

Although most studies using the TNT task focus on suppression-induced forgetting, our concern in the present studies is different. We are interested in whether people's efforts to suppress retrieval induce an adverse hippocampal state that disrupts memory for events in the temporal vicinity of suppression. To test for this possibility, we adapted the TNT phase. In the modified task, we added "bystander" events into the TNT phase (Hulbert et al., 2016). "Bystander" events are incidentally encoded stimuli inserted in between Think and No-Think trials, selected to be as unrelated to the TNT items as possible, in both form and content. For example, whereas our TNT stimuli were word pairs, bystanders were photographs of objects in natural settings selected to be unrelated to any pair (e.g., a peacock in a parking lot, Figure 4). When participants



encountered these pictures, they needed to covertly explain how the central object came to be in its surroundings and to rate on a scale of 1–4, how difficult generating that explanation was. Some bystander items were both preceded by and followed by Think trials, whereas others were nestled between two No-Think trials. The remainder of the bystanders appeared between a Think and a No-Think trial (Figure 4).

The purpose of inserting bystander scenes between Think and No-Think trials was to determine whether people's ability to form a memory trace of a novel event (the bystander) was influenced by suppressing No-Think items in the preceding or following trials. Would having just suppressed retrieval of a word pair impair new encoding? Would a scene that one has just encoded be affected by suppressing retrieval of a word pair after it? One challenge to answering these questions, however, is to rule out alternative explanations of memory deficits for bystanders in terms of divided attention or distraction. For example, perhaps people find trying to suppress retrieval on No-Think trials particularly hard, leaving them distracted after doing it, making them less able to attend to and encode a new scene. Perhaps this distraction is greater after a No-Think than after a Think trial, which might cause a difference in memory for bystander items on later tests. Such effects would not reflect amnesia due to suppressed hippocampal function, but simply distraction during encoding. To control for distraction, participants never encoded bystander scenes adjacently to Think or No-Think trials. Rather, participants performed an odd-even judgement task in which they had to judge as quickly as possible whether each in a series of presented numbers was odd or even (Figure 4). Participants always performed this judgement task both just before and just after encoding the bystander events. The judgement task serves as a buffer that precedes the encoding of bystander events for 5–10 s, refocusing participants attention onto a common task, prior to encoding. In doing so, the odd/even task controls for differential task-switching costs (Monsell, 2003; Wylie & Allport, 2000) across Think and No-Think trials. This demanding task also acted as a buffer for 5–10 s after bystander encoding, so that any lingering thoughts about bystanders were cleared away before the following TNT trials (Figure 4).

In the typical TNT paradigm, the TNT phase would be followed by a final recall test for the Baseline, Think, and No-Think items, as described above. However, because our concern was not with suppression-induced forgetting effects, we replaced this final test for the studied pairs with a test of the bystanders themselves. In some of the current experiments, this involved trials that presented the scene in which the object had appeared on its own and asking participants to recall the object from that scene (Figure 5). In other studies, we tested recognition memory and source memory for presented content. According to the amnesic shadow hypothesis, the bystander events that were encoded

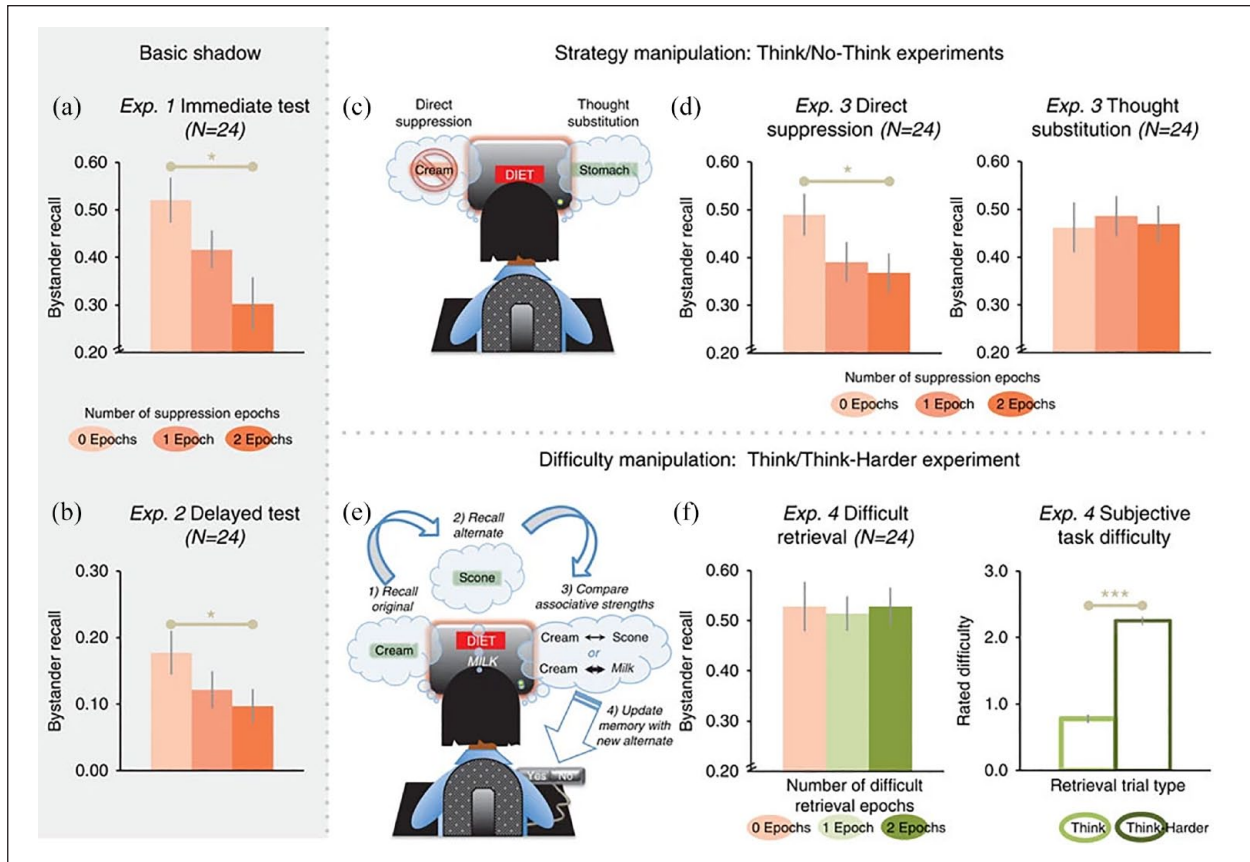
in the temporal surround of suppression trials would be affected by the suppression mechanism. Hence, they would be forgotten more easily compared with the bystander events which were encoded in the temporal surround of retrieval trials. More specifically, we predicted that bystander recall would vary systematically according to the “number of doses” of suppression (0, 1, or 2) surrounding it. Thus, bystanders surrounded on both sides by suppression (2 doses) would be recalled the worst, whereas bystanders surrounded on no sides by suppression (0 doses—i.e., item surrounded by Think trials) would be recalled the best, with performance in the 1 dose condition in between those two (i.e., items preceded by suppression and followed by retrieval, or vice versa).

Because the focus of this adapted TNT paradigm is very different and seeks to measure the impact of putative hippocampal modulation on episodic memory function, we refer to this new procedure as the hippocampal modulation or *HM paradigm*, partially in honour of amnesic patient HM (Hulbert et al., 2016). We note that the strategy and structure of this paradigm can be adapted to study the impact of any other task (aside from the TNT task) believed to affect hippocampal function, and so the HM paradigm represents a category of procedures addressing this question. So, for example, Think and No-Think trials could be replaced with a different pairing of tasks thought to modulate hippocampal activity. Illustrations of how one might do this will arise in later sections of this article.

## The amnesic shadow: basic findings

The results of our first amnesic shadow experiment with this design are reported in Figure 5a. As this figure illustrates, recall performance in the 0-epoch (i.e., dose) condition, in which the bystander was surrounded by 0 No-Think trials, was good (around 52%), especially given that encoding (a) was incidental, (b) occurred during a complex task, and (c) occurred after a sizable delay (20–30 min). In contrast, surrounding the bystander by 2 doses of suppression virtually cut recall performance in half, compared with the 0-dose condition. Recall for the 1-dose condition fell in between these two extremes. Thus, final recall performance varied systematically according to the number of doses of suppression in the temporal vicinity of the bystander, consistent with the predicted existence of the amnesic shadow.

The sizable drop in recall performance arising from surrounding suppression trials cannot be readily explained by known forgetting mechanisms. All standard factors which could potentially cause variation in retention were held constant between the three “dose” conditions, including (a) retention interval, until the final test, (b) amount of time and attention devoted to the judgement encoding task (there were no differences in judgements or judgement times), and (c) the potential for proactive or retroactive



**Figure 5.** Main results of the amnesic shadow study. (a) Immediate cued-recall accuracy for bystanders as a function of the number of adjacent suppression epochs. The difference between the left- (peach colour) and right-most (dark orange) bars reveals an amnesic shadow (F-test). (b) Amnesic shadow observed after a 24-h delay. (c) Experiment 3's two No-Think strategies: direct suppression and thought substitution. (d) Direct suppression, not thought substitution, caused a shadow. (e) Experiment 4 replaced No-Think trials with a difficult "Think Harder" task. (f) No shadow was observed in Experiment 4 (left subpanel), despite a significant difficulty disparity across the Think and Think Harder conditions. Source. From Hulbert et al. (2016).

interference. These findings are consistent, however, with our amnesic shadow hypothesis derived from independent imaging data showing clear evidence of hippocampal suppression in the retrieval suppression task. We also confirmed that the amnesic shadow is an enduring effect, as one might expect if suppression had interrupted encoding or stabilisation processes. In a second experiment, we increased the retention interval from 5 min to 24 hr. As in the first experiment, recall varied according to the number of doses of suppression (Figure 5b).

So far, we have argued that stopping the retrieval process is essential to inducing an amnesic shadow, given the linkage of this process to hippocampal downregulation. To establish that retrieval suppression is truly necessary, however, requires that we show that other methods of controlling unwanted memories that do not involve retrieval suppression do not induce a similar effect. One such memory control process is thought substitution. Thought substitution refers to a method of memory control wherein, when confronted with an unwelcome reminder, people control

awareness of an unwanted memory by retrieving a substitute or alternate memory to stop the target memory's entrance into awareness. Both thought substitution and retrieval suppression methods induce forgetting of the controlled content (Benoit & Anderson, 2012; Bergström et al., 2009; Hertel & Calcaterra, 2005); however, hippocampal downregulation occurs during retrieval suppression, but not during thought substitution (Benoit & Anderson, 2012; see Anderson & Hulbert, 2021 for a review of other dissociations between these memory control methods). Indeed, thought substitution during No-Think trials requires hippocampal engagement to support participants' effort to retrieve the substitute memory, making hippocampal suppression counterproductive.

To test the specificity of the amnesic shadow to retrieval suppression, Experiment 3 juxtaposed the impact of retrieval suppression and thought substitution on the recall of bystander items in two groups (Figure 5c). Direct suppression participants performed the same task as described for Experiments 1 and 2; thought substitution participants,

in contrast, performed the same set of tasks but with modified instructions concerning how to control memory during the No-Think trials. Instead of focusing on the reminder and stopping the retrieval of the target, participants were asked to retrieve an alternate memory association that they had previously studied along with the reminder to prevent target retrieval. Notably, both groups were closely matched; both were trained in direct suppression and thought substitution methods in a practice session on filler items, prior to the main TNT task; and both groups were told about the same thought substitutes for each No-Think cue. However, for the main TNT task, the direct suppression group was told to control awareness exclusively by suppressing retrieval, whereas the thought substitution group was asked to retrieve the preassigned thought substitute for each No-Think item, to keep the target out of awareness. This experiment provides a strong test of our hypothesis; participants in both groups shared the goal of controlling awareness of the target. The groups, however, used opposite means of achieving that goal: whereas the suppression group disengaged retrieval, the thought substitution group engaged it, but directed retrieval at a replacement memory. Of these techniques, only suppression involves hippocampal downregulation. Interestingly, this study revealed that the amnesic shadow was indeed selective to retrieval suppression. The thought substitution group showed no memory deficits for bystander events (Figure 5d). Thus, the amnesic shadow does not stem from active forgetting in general but arises with a “shutdown” of retrieval processes.

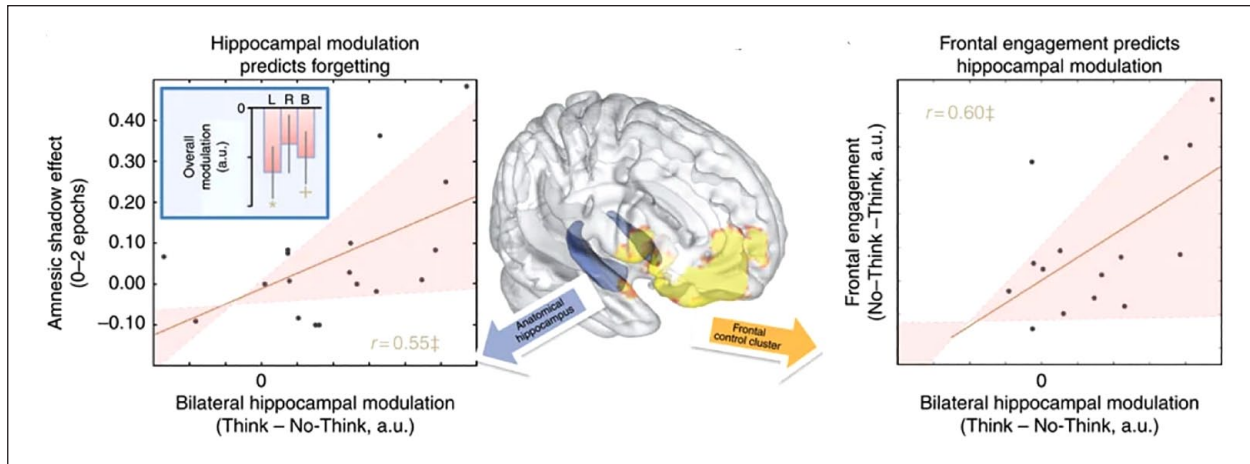
Although the foregoing findings suggest that the amnesic shadow is specific to retrieval suppression, there is an alternative possibility worth considering: it is not retrieval suppression per se that causes the amnesic shadow but some other attribute of retrieval suppression that goes along with it. For example, it could be that trying to exclude an associate from awareness by retrieval suppression when confronted with a reminder is unusually difficult, and perhaps more difficult than retrieving a thought substitute associate instead. More difficult tasks are generally more distracting and so may tend to disrupt hippocampal function more. This difficulty hypothesis could explain why an amnesic shadow occurs for retrieval suppression and not for thought substitution. If true, then a comparably challenging task that does not involve retrieval suppression should also produce an amnesic shadow, just as retrieval suppression does. We investigated this possibility by creating the “Think/Think Harder” task (hereinafter, the TTH task) (Hulbert et al., 2016) (Figure 5e).

The TTH task consisted of two types of trials, as does the TNT. The first trial type was the same as the Think trials in the “TNT” task; participants were simply cued to retrieve and maintain the associated target, given its cue during these trials. In place of the No-Think trials, however, the TTH procedure included “Think Harder” trials.

On Think Harder trials, the reminder of a pair would appear in red, as on No-Think trials, but participants were instructed instead to retrieve two associates, both previously learned with the reminder, the target, and an “alternate.” In addition, a novel item appeared on screen, and participants had to compare the strength of relationships between the associates themselves and between the associates and the novel item. For example, if the reminder was “diet,” the two associates might be “cream” and “scone.” Participants had to retrieve “cream” and “scone” when “diet” appeared. Furthermore, participants then had to compare the novel item “milk” presented on the screen to both “cream” and “scone,” and decide whether “cream-milk,” or “cream-scone,” were more strongly related to one another. In this example, people would normally judge “cream-milk” to be more strongly related than cream-scone. If so, then participants had to replace “scone” with “milk” and then, from the next Think Harder trial onwards, start retrieving “cream” and “milk” when presented with “diet.” This challenging task involved multiple retrievals, comparing strengths of associations, and updating one’s memory for the associates which go with the reminder.

The results of this experiment were clear. Surrounding bystander scenes by exceedingly difficult tasks that did not involve retrieval suppression did not cause an amnesic shadow effect. Recall performance for the bystander items did not vary as a function of the number of surrounding “Think Harder” trials. Nevertheless, participants rated the “Think Harder” trials as being exceedingly difficult, and more difficult than “Think” trials (Figure 5f). This finding suggests that task difficulty or distraction cannot explain the amnesic shadow effect and that disengaging the retrieval process is essential to observe the effect.

Although the foregoing findings suggest that stopping the retrieval process is necessary to induce an amnesic shadow, they do not directly demonstrate that the shadow derives from hippocampal downregulation. To establish this linkage, we conducted a functional magnetic resonance imaging (fMRI) study in which participants performed the TNT task inside the scanner (Hulbert et al., 2016). We sought to quantify the modulation of activation in the hippocampus by retrieval suppression and link this modulation to the size of amnesic shadow. If the amnesic shadow is a behavioural reflection of disrupted hippocampal function, then the more hippocampal downregulation occurs during No-Think trials, the greater should be the shadow effect in a subsequent test for the bystander events. The results strongly confirmed this prediction (Figure 6, left). The results also confirmed that stronger hippocampal downregulation was linked to greater engagement of the prefrontal cortex during the No-Think trials (Figure 6, right). Taken together, these findings indicate that retrieval suppression not only downregulates hippocampal activity but also induces a degraded functional state in this structure. This degraded state creates windows of amnesia



**Figure 6.** Neuroimaging results of the amnesic shadow study. Left: Across participants, bilateral hippocampal modulation (higher scores indicate greater modulation) predicted the observed amnesic shadow effect after the scanning session was complete. Right: The greater the prefrontal engagement during No-Think trials, the greater the observed hippocampal modulation. Source: Image modified from Hulbert et al. (2016).

resembling organic amnesia but induced by prefrontal regulation of hippocampal function.

Taken together, the preceding findings suggest that the severe memory deficits first identified in HM, arising from structural damage to the hippocampus can be mimicked in otherwise healthy individuals when those people engage inhibitory control to suppress retrieval, downregulating hippocampal activity systemically. We know that suppression is systemic, rather than targeted at individual traces, because it affects not only the thing that people seek to suppress, but any recent hippocampally dependent memory, irrespective of its relatedness to the suppressed content. Moreover, evidence suggests that this hippocampal downregulation reflects active inhibition at the neural level. For example, with magnetic resonance spectroscopy (MRS) we studied whether individual variation in the concentration of gamma-aminobutyric acid (GABA) within the hippocampus might be related to suppression of this structure during retrieval stopping (Schmitz et al., 2017). GABA is the chief inhibitory neurotransmitter in the brain, and the sole inhibitory neurotransmitter in the hippocampus, and is present in inhibitory interneurons throughout this structure (as well as arising extracellularly). We found that higher concentrations of GABA in the hippocampus predicted (a) larger suppression-induced forgetting, (b) more robust hippocampal downregulation during retrieval suppression, and (c) stronger modulation by the prefrontal cortex of the hippocampus for participants higher in GABA. Thus, higher GABA concentrations in the hippocampus enabled the prefrontal cortex to exert inhibitory control over the hippocampus. GABA concentrations within the visual cortex, in contrast, bore no relationship to these outcomes, demonstrating that the effects are region-specific. These results show that GABAergic interneurons in the hippocampus are important for mediating

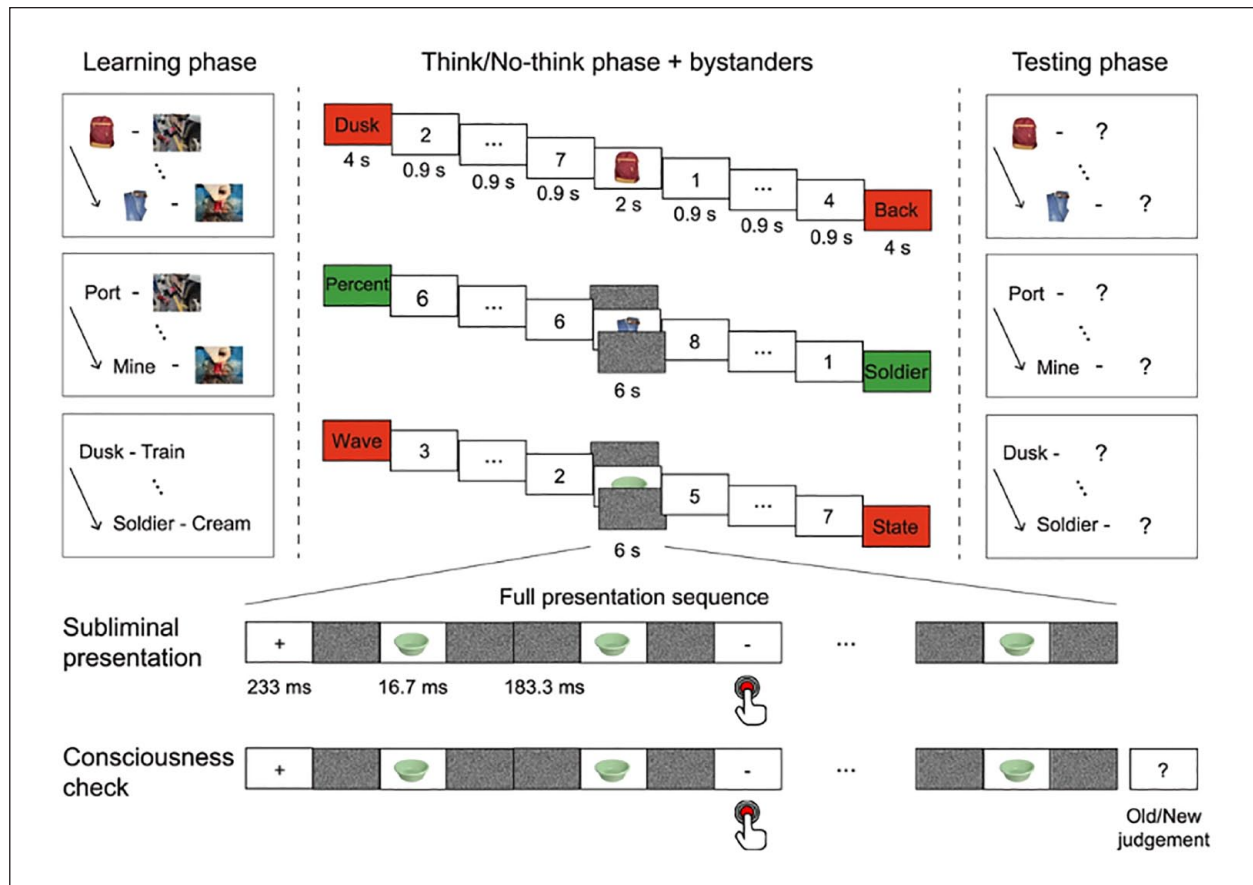
hippocampal downregulation and creating the dysfunctional state that produces the amnesic shadow.

### The amnesic shadow for older events

Based on the discussion thus far, one might wonder how pervasive the amnesic shadow could be in producing forgetting in daily life. Its impact is presumably limited to encoding or stabilisation of event representations acquired in the immediate vicinity of people's efforts to intentionally suppress a memory. Arguably, most of the memories we form are not vulnerable to these circumstances. For example, what about the vast store of memories we have acquired in our past, prior to an effort to suppress retrieval? Have those memories entirely escaped the destructive impacts of hippocampal modulation? If so, modulation of hippocampal activity may not be a major contributor to much of the forgetting we experience.

Recently, however, several findings suggest that the amnesic shadow has broader impact than this characterization would suggest. Two studies examined whether older memories, acquired prior to efforts to suppress, might also be vulnerable to disruption if they are simply cued in the temporal vicinity of suppression (Zhu et al., 2022; Zhu & Wang, 2021). For example, Zhu et al. (2022) modified the TNT task by asking participants to encode bystander memories before the TNT training was even begun (Figure 7). Bystanders were composed of a set of scenes, each one trained to be associated to two cues: an object and a word. Only once these bystanders had been learned did participants then learn the word pairs that would be used for the TNT task. During the TNT phase, instead of bystander encoding trials in between Think and No-Think trials, participants were presented with the object cue of a bystander pair and simply asked whether they recognised the cue





**Figure 7.** The Think/No-think task modified to study amnesic shadow effects for previously encoded bystanders. Bystanders were encoded before the TNT task began and were cued between Think and No-Think trials. Half of the encoded items were subliminally cued to see whether subliminal reactivation of the memories in the temporal surround of a suppression trial also induced amnesia for these items.

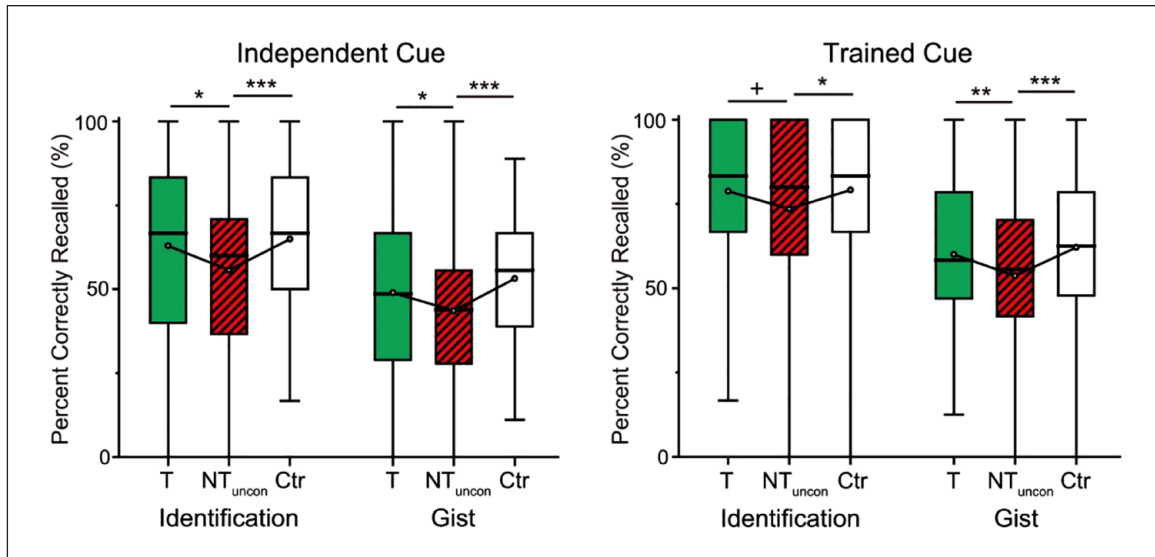
Source. Image reproduced from Zhu et al. (2022).

itself (without any need to recall the associated scene). As in previous amnesic shadow studies, the cuing of these scenes was preceded and succeeded by an odd/even buffer task to match task-set switching effects across conditions. Importantly, half of the bystander object cues were presented subliminally: they appeared preceded and followed by white noise masks so that they could not be consciously perceived. Prior work suggests that even this imperceptible cue presentation may subliminally activate the associated memory in the hippocampus (Degonda et al., 2005; Duss et al., 2011, 2014; Henke et al., 2003; Reber & Henke, 2011; Reber et al. 2012; Wuethrich et al., 2018). Thus, bystander scenes, whether cued consciously or subliminally may, by virtue of hippocampal reactivation, be rendered vulnerable to the amnesic shadow created by the surrounding retrieval suppression trials.

One advantage that this cuing design has over the one used in the original amnesic shadow studies (that focused on novel encoding) is that it allows for the inclusion of baseline pairs that enable the separate quantification of the effects of surrounding bystanders by Think trials or

No-Think trials. So, for example, in addition to encoding pairs that would be cued between No-Think or between Think trials, participants encoded a third set of bystander items that were never cued during the TNT phase at all. On the final recall test after the TNT phase, how would bystanders cued in between No-Think trials fare, in comparison to performance on these baseline items? Importantly, participants showed forgetting for such bystanders, relative to Baseline pairs, indicating a clear amnesic shadow effect (Figure 8). Strikingly, the forgetting effect arose even for the scenes which were only subliminally reactivated and not consciously recalled, even when the scenes were probed by the second set of (word) cues that never appeared during the TNT phase (Figure 8). These findings illustrate that the memory deficits induced by hippocampal suppression were not cue-dependent. In contrast, scenes cued in between Think trials showed none of these forgetting effects (Figure 8).

The foregoing studies demonstrate that the amnesic shadow is not limited to recently encoded memories but can affect any memory that happens to be thought about in



**Figure 8.** Results from Zhu et al. (2022). Previously encoded bystanders were recalled less when they were cued in the temporal surround of No-Think trials; forgetting was also found in bystanders cued subliminally. In this experiment (Experiment 2), all bystanders were subliminally cued.  
 Source. From Zhu et al. (2022).

the temporal vicinity of suppression. And such reactivation need not be deliberate or elaborate: even memories cued subliminally were vulnerable to disruption. It is presently unknown how extended the “temporal vicinity” around suppression could be that is vulnerable to the amnesic shadow effect. All current studies have inserted 5–10s of buffer tasks around bystander encoding to control for task-switching effects, indicating that suppression can affect memories for at least the surrounding 10s in either direction. This does not mean, however, that the effect does not extend well beyond this, as this has not been examined. But we now know that, irrespective of the width of this window, older memories can be affected.


### Mnemonic process inhibition in other tasks

Although the preceding work expands the range of memories that could be affected by mnemonic process inhibition, one might wonder whether retrieval suppression happens often enough in daily life for it to be a major factor underlying episodic forgetting. If someone does not suppress retrieval very often, does that mean that hippocampal suppression rarely leads them to forget? An important possibility, however, is that the suppression of mnemonic functions required by intentional retrieval stopping could also arise in other task contexts to support optimal performance. For example, working memory updating tasks, such as the n-back task have been shown to robustly engage prefrontal control regions, and also downregulate the hippocampus. As a memory updating task, the n-back task (especially more demanding versions, such as the

2-back task) places heavy demands on the suppression of irrelevant memory contents (on every trial some portion of the content retained in memory must be expelled), even though this suppression may not be the main focus of the task.

Building on the observation of hippocampal suppression, Mullally and O’Mara (2013) tested whether performing an n-back working memory task led to deficits in the encoding of new items into memory in its immediate aftermath. Participants were asked to perform the n-back task repeatedly in blocks; each block lasted for 1 min. Between blocks, participants were asked to memorise a collection of face-name associations and then were immediately tested on their retention of these pairs. This cycle of n-back/face-name encoding was repeated for 4 runs. Participants showed improved recall for the face-name associations over the runs, as they were encoding the same set repeatedly. Of main interest, however, was how well participants could recall these associations as a function of whether they had been assigned to perform a 2-back version of the n-back task or a 0-back version (different groups). Prior work establishes robust disengagement of hippocampal activity in the 2-back condition, but not in the 0-back condition, suggesting that embedding face-name encoding between blocks of the 2-back task may lead hippocampally dependent episodic encoding of face-name pairs to suffer. Consistent with this prediction, participants in the 2-back condition showed significantly worse recall of the pairs than did participants in the 0-back condition (Figure 9).

Given the evidence for hippocampal downregulation in the n-back task, Mullally and O’Mara’s findings may



**Figure 9.** Results from Mullally and O’Mara (2013). Face-name pairs were recalled more poorly when they were encoded between 2-back blocks, compared with when the pairs were encoded between blocks of the 0-back task. Source. From Mullally and O’Mara (2013).

reflect another instance of an amnesic shadow. If so, these results lend credence to the possibility that the amnesic shadow arises even when there is no intentional suppression of retrieval, as long as the hippocampus is functionally disengaged and mnemonic processing, inhibited. If amnesic shadows arise in other tasks, when might they happen and under what conditions? Anderson and Hulbert (2021) proposed a broad hypothesis about the conditions that give rise to hippocampal suppression: hippocampal activity may be suppressed when attention is oriented away from memory and towards the external world. Here, external attention refers to attention directed to external perceptions such as attention to visual, auditory, or somatosensory stimuli; in contrast, internal attention refers to attention to processes or contents that are not currently being perceived, such as memories, ideas, and feelings. When focusing externally, especially during demanding tasks, it may be distracting and maladaptive to dwell too much on one’s internal thoughts. Hence, it may be useful to disengage from internal processes. Anderson & Hubert hypothesised that retrieval suppression is a special case when a person voluntarily and intentionally disengages from memory. However, other processes may cause hippocampal disengagement, even when one is not consciously suppressing retrieval.

Strikingly, a large body of evidence already supports the idea that focusing attention externally suppresses brain structures critical to memory, including the hippocampus. For instance, it has long been known that the default mode network is suppressed when attention is reoriented from internal cues to external tasks (Anticevic et al., 2012; Binder et al., 1999; Buckner et al., 2008; Buckner & DiNicola, 2019; Fakhraei et al., 2021; Harrison et al., 2011; Mazoyer et al., 2001; McKiernan et al., 2006;

Raichle, 2015; Shulman et al., 1997; Singh & Fawcett, 2008; Smallwood et al., 2013, 2021). The default mode network (Shulman et al., 1997) refers to a collection of brain regions activated when people are not performing a particular task but are resting and typically engaged in mind-wandering. These regions are strongly functionally connected to each other during a “resting state” scan (Andrews-Hanna et al., 2010; Fransson & Marrelec, 2008). Importantly, many default network regions are part of the network involved in episodic and autobiographical memory (D’Argembeau et al., 2005; Diana et al., 2007; Schacter et al., 2007; Spreng & Grady, 2010), consistent with the possibility that when people are at rest (and mind-wandering), their cognition focuses on content originating from memory. The fact that switching attention to a difficult perceptual or motor task suppresses these regions is therefore consistent with the possibility that such tasks suppress mnemonic processing.

If the foregoing interpretation is correct, it suggests a substantial role of the amnesic shadow in everyday forgetting. This argument suggests that any time that people focus intently on the external world to achieve some task, mnemonic process inhibition would occur. Such regular bouts of mnemonic suppression may induce an amnesic shadow for all memories recently activated within the hippocampus. Given that much of our daily life involves attention to the external world, recurring inhibition of the hippocampus’s functional state may constitute a ubiquitous factor promoting the gradual erosion of experience from memory.

## Concluding remarks

Organic amnesia is the most profound form of forgetting ever demonstrated. Here we have argued that structural damage to the hippocampus is not the only way in which one can become amnesic. We argued that there exist cognitive operations that have, as their outcome, the suppression of the hippocampus’s functional state, creating reversible dysfunction of the hippocampus’s that mimics organic amnesia. We developed the case for the existence of an amnesic shadow as a result of retrieval suppression in particular detail, ruling out a great variety of alternative interpretations that could explain the data without positing hippocampal suppression. The amnesia that results is durable and affects not only encoding and stabilisation of recent memories, but also older memories that simply have been reactivated by cues in the vicinity of suppression. Data suggest that these amnesic effects arise from the global suppression of mnemonic processes driven by GABAergic inhibitory processes that suppress hippocampal activity in response to prefrontal control input.

Importantly, we argued that this form of amnesia may not be specific to the act of intentional retrieval suppression.

Cases have already been reported of amnesic-shadow-like effects induced by tasks such as the n-back working memory task, difficult versions of which are known to suppress hippocampal activity. We argued that any task that requires orientation of attention to external stimuli and away from memory, especially if challenging in nature, may suppress hippocampal activity and induce an amnesic shadow effect. If this hypothesis is correct, then systemic suppression of hippocampal activity may be a ubiquitous mechanism of forgetting from episodic memory. As far as we are aware, no such mechanism has been proposed previously, and so mnemonic process inhibition may be an entirely novel and important source of forgetting. If this ultimately proves to be the case, then mnemonic process inhibition and the resulting amnesic shadow may be yet another scientific discovery that ultimately can be attributed to the tragic loss of memory suffered by HM.

### Declaration of conflicting interests

The author(s) declared no potential conflicts of interest with respect to the research, authorship, and/or publication of this article.

### Funding

The author(s) received no financial support for the research, authorship, and/or publication of this article: Preparation for this manuscript was supported by Medical Research Council MC-A060-5PR00 to MCA.

### ORCID iD

Michael C Anderson  <https://orcid.org/0000-0001-9505-9299>

### References

- Anderson, M. C., & Green, C. (2001). Suppressing unwanted memories by executive control. *Nature*, *410*(6826), 366–369.
- Anderson, M. C., & Hanslmayr, S. (2014). Neural mechanisms of motivated forgetting. *Trends in Cognitive Sciences*, *18*(6), 279–292.
- Anderson, M. C., & Huddleston, E. (2012). Towards a cognitive and neurobiological model of motivated forgetting. In R. Belli (Ed.), *True and false recovered memories: Toward a reconciliation of the debate* (pp. 53–120). Springer.
- Anderson, M. C., & Hulbert, J. C. (2021). Active forgetting: Adaptation of memory by prefrontal control. *Annual Review of Psychology*, *72*, 1–36.
- Anderson, M. C., Ochsner, K. N., Kuhl, B., Cooper, J., Robertson, E., Gabrieli, S. W., Glover, G. H., & Gabrieli, J. D. (2004). Neural systems underlying the suppression of unwanted memories. *Science*, *303*(5655), 232–235.
- Andrews-Hanna, J. R., Reidler, J. S., Sepulcre, J., Poulin, R., & Buckner, R. L. (2010). Functional-anatomic fractionation of the brain's default network. *Neuron*, *65*(4), 550–562.
- Anticevic, A., Cole, M. W., Murray, J. D., Corlett, P. R., Wang, X. J., & Krystal, J. H. (2012). The role of default network deactivation in cognition and disease. *Trends in Cognitive Sciences*, *16*(12), 584–592.
- Apšvalka, D., & Anderson, M. C. (in preparation). *Memory control along the hippocampal longitudinal axis*.
- Apšvalka, D., Ferreira, C. S., Schmitz, T. W., Rowe, J. B., & Anderson, M. C. (2022). Dynamic targeting enables domain-general inhibitory control over action and thought by the prefrontal cortex. *Nature Communications*, *13*(1), 274.
- Benoit, R. G., & Anderson, M. C. (2012). Opposing mechanisms support the voluntary forgetting of unwanted memories. *Neuron*, *76*(2), 450–460.
- Bergström, Z. M., de Fockert, J. W., & Richardson-Klavehn, A. (2009). ERP and behavioural evidence for direct suppression of unwanted memories. *NeuroImage*, *48*(4), 726–737.
- Binder, J. R., Frost, J. A., Hammeke, T. A., Bellgowan, P. S. F., Rao, S. M., & Cox, R. W. (1999). Conceptual Processing during the Conscious Resting State: A Functional MRI Study. *Journal of Cognitive Neuroscience*, *11*(1), 80–93.
- Buckner, R. L., Andrews-Hanna, J. R., & Schacter, D. L. (2008). The brain's default network: Anatomy, function, and relevance to disease. *Annals of the New York Academy of Sciences*, *1124*(1), 1–38.
- Buckner, R. L., & DiNicola, L. M. (2019). The brain's default network: Updated anatomy, physiology and evolving insights. *Nature Reviews Neuroscience*, *20*(10), 593–608.
- Cipolotti, L., & Bird, C. M. (2006). Amnesia and the hippocampus. *Current Opinion in Neurology*, *19*(6), 593–598.
- Corkin, S., Amaral, D. G., González, R. G., Johnson, K. A., & Hyman, B. T. (1997). HM's medial temporal lobe lesion: Findings from magnetic resonance imaging. *Journal of Neuroscience*, *17*(10), 3964–3979.
- D'Argembeau, A., Collette, F., Van der Linden, M., Laureys, S., Del Fiore, G., Degueldre, C., Luxen, A., & Salmon, E. (2005). Self-referential reflective activity and its relationship with rest: A PET study. *NeuroImage*, *25*(2), 616–624.
- Degonda, N., Mondadori, C. R., Bosshardt, S., Schmidt, C. F., Boesiger, P., Nitsch, R. M., Hock, C., & Henke, K. (2005). Implicit associative learning engages the hippocampus and interacts with explicit associative learning. *Neuron*, *46*(3), 505–520.
- Diana, R. A., Yonelinas, A. P., & Ranganath, C. (2007). Imaging recollection and familiarity in the medial temporal lobe: A three-component model. *Trends in Cognitive Sciences*, *11*(9), 379–386.
- Duss, S. B., Oggier, S., Reber, T. P., & Henke, K. (2011). Formation of semantic associations between subliminally presented face-word pairs. *Consciousness and Cognition*, *20*(3), 928–935.
- Duss, S. B., Reber, T. P., Hänggi, J., Schwab, S., Wiest, R., Müri, R. M., Brugger, P., Gutbrod, K., & Henke, K. (2014). Unconscious relational encoding depends on hippocampus. *Brain*, *137*(12), 3355–3370.
- Fakhraei, L., Francoeur, M., Balasubramani, P. P., Tang, T., Hulyalkar, S., Buscher, N., Mishra, J., & Ramanathan, D. S. (2021). Electrophysiological correlates of rodent default-mode network suppression revealed by large-scale local field potential recordings. *Cerebral Cortex Communications*, *2*(2), tgab034.
- Fransson, P., & Marrelec, G. (2008). The precuneus/posterior cingulate cortex plays a pivotal role in the default mode network: Evidence from a partial correlation network analysis. *NeuroImage*, *42*(3), 1178–1184.



- Harrison, B. J., Pujol, J., Contreras-Rodríguez, O., Soriano-Mas, C., López-Solà, M., Deus, J., Ortiz, H., Blanco-Hinojo, L., Alonso, P., Hernández-Ribas, R., Cardoner, N., & Menchón, J. M. (2011). Task-induced deactivation from rest extends beyond the default mode brain network. *PLOS ONE*, *6*(7), Article e22964.
- Henke, K., Mondadori, C. R., Treyer, V., Nitsch, R. M., Buck, A., & Hock, C. (2003). Nonconscious formation and reactivation of semantic associations by way of the medial temporal lobe. *Neuropsychologia*, *41*(8), 863–876.
- Hertel, P. T., & Calcaterra, G. (2005). Intentional forgetting benefits from thought substitution. *Psychonomic Bulletin & Review*, *12*, 484–489.
- Hulbert, J. C., Henson, R. N., & Anderson, M. C. (2016). Inducing amnesia through systemic suppression. *Nature Communications*, *7*(1), 11003.
- Huppert, F. A., & Piercy, M. (1979). Normal and abnormal forgetting in organic amnesia: Effect of locus of lesion. *Cortex*, *15*(3), 385–390.
- Levin, H. S., Peters, B. H., & Hulkonen, D. A. (1983). Early concepts of anterograde and retrograde amnesia. *Cortex: A Journal Devoted to the Study of the Nervous System and Behavior*, *19*, 427–440.
- Levy, B. J., & Anderson, M. C. (2002). Inhibitory processes and the control of memory retrieval. *Trends in Cognitive Sciences*, *6*(7), 299–305.
- Mamat, Z., & Anderson, M. C. (2023). Improving mental health by training the suppression of unwanted thoughts. *Science Advances*, *9*(38), eadh5292.
- Marsh, L., & Anderson, M. C. (2020). Inhibition as a cause of forgetting. In M. J. Kahana & A. D. Wagner (Eds.), *The Oxford handbook of human memory* (Vol. 1). Oxford University Press.
- Mazoyer, B., Zago, L., Mellet, E., Bricogne, S., Etard, O., Houdé, O., Houde, O., Crivello, F., Joilliot, M., Petit, L., & Tzourio-Mazoyer, N. (2001). Cortical networks for working memory and executive functions sustain the conscious resting state in man. *Brain Research Bulletin*, *54*(3), 287–298.
- McKiernan, K. A., D'Angelo, B. R., Kaufman, J. N., & Binder, J. R. (2006). Interrupting the “stream of consciousness”: An fMRI investigation. *NeuroImage*, *29*(4), 1185–1191.
- Monsell, S. (2003). Task switching. *Trends in Cognitive Sciences*, *7*(3), 134–140.
- Mullally, S. L., & O'Mara, S. M. (2013). Suppressing the encoding of new information in memory: A behavioral study derived from principles of hippocampal function. *PLOS ONE*, *8*(1), Article e50814.
- Parkin, A. J. (1984). Amnesic syndrome: A lesion-specific disorder? *Cortex*, *20*(4), 479–508.
- Raichle, M. E. (2015). The brain's default mode network. *Annual Review of Neuroscience*, *38*, 433–447.
- Reber, T. P., & Henke, K. (2011). Rapid formation and flexible expression of memories of subliminal word pairs. *Frontiers in Psychology*, *2*, Article 343.
- Reber, T. P., Luechinger, R., Boesiger, P., & Henke, K. (2012). Unconscious relational inference recruits the hippocampus. *Journal of Neuroscience*, *32*(18), 6138–6148.
- Schacter, D. L., Addis, D. R., & Buckner, R. L. (2007). Remembering the past to imagine the future: The prospective brain. *Nature Reviews Neuroscience*, *8*(9), 657–661.
- Schmitz, T. W., Correia, M. M., Ferreira, C. S., Prescott, A. P., & Anderson, M. C. (2017). Hippocampal GABA enables inhibitory control over unwanted thoughts. *Nature Communications*, *8*(1), 1311.
- Scoville, W. B., & Milner, B. (1957). Loss of recent memory after bilateral hippocampal lesions. *Journal of Neurology, Neurosurgery, and Psychiatry*, *20*(1), 11.
- Shimamura, A. P. (1992). Organic amnesia. In L. R. Squire (Ed.), *Encyclopedia of learning and memory* (pp. 30–34). Macmillan.
- Shulman, G. L., Corbetta, M., Buckner, R. L., Fiez, J. A., Miezin, F. M., Raichle, M. E., & Petersen, S. E. (1997). Common blood flow changes across visual tasks: I. Increases in subcortical structures and cerebellum but not in nonvisual cortex. *Journal of Cognitive Neuroscience*, *9*(5), 624–647.
- Singh, K. D., & Fawcett, I. P. (2008). Transient and linearly graded deactivation of the human default-mode network by a visual detection task. *NeuroImage*, *41*(1), 100–112.
- Smallwood, J., Bernhardt, B. C., Leech, R., Bzdok, D., Jefferies, E., & Margulies, D. S. (2021). The default mode network in cognition: A topographical perspective. *Nature Reviews Neuroscience*, *22*(8), 503–513.
- Smallwood, J., Tipper, C., Brown, K., Baird, B., Engen, H., Michaels, J. R., Grafton, S., & Schooler, J. W. (2013). Escaping the here and now: Evidence for a role of the default mode network in perceptually decoupled thought. *NeuroImage*, *69*, 120–125.
- Spreng, R. N., & Grady, C. L. (2010). Patterns of brain activity supporting autobiographical memory, prospection, and theory of mind, and their relationship to the default mode network. *Journal of Cognitive Neuroscience*, *22*(6), 1112–1123.
- Squire, L. R. (2009). The legacy of patient HM for neuroscience. *Neuron*, *61*(1), 6–9.
- Wessel, J. R., & Aron, A. R. (2017). On the globality of motor suppression: Unexpected events and their influence on behavior and cognition. *Neuron*, *93*(2), 259–280.
- Wuethrich, S., Hannula, D. E., Mast, F. W., & Henke, K. (2018). Subliminal encoding and flexible retrieval of objects in scenes. *Hippocampus*, *28*(9), 633–643.
- Wylie, G., & Allport, A. (2000). Task switching and the measurement of “switch costs.” *Psychological Research*, *63*, 212–233.
- Zhu, Z., Anderson, M. C., & Wang, Y. (2022). Inducing forgetting of unwanted memories through subliminal reactivation. *Nature Communications*, *13*(1), 6496.
- Zhu, Z., & Wang, Y. (2021). Forgetting unrelated episodic memories through suppression-induced amnesia. *Journal of Experimental Psychology: General*, *150*(3), 401–413.