MEMORY, 2017
https://doi.org/10.1080/09658211.2017.1354999

Routledge
Taylor & Francis Group

REVIEW ARTICLE

Check for updates

Unintended side effects of a spotless mind: theory and practice

Justin C. Hulbert, Zall Hirschstein, Clarence A. L. Brontë and Eleanor Broughton

Psychology Program, Bard College, Annandale-on-Hudson, NY, USA

ABSTRACT

Forgetting can be either a source of great frustration or one of great relief, depending on whether the memories in question are relevant to one's immediate goals. Adopting an appropriate strategy or memory mode can help achieve these goals. But do efforts to control memory engender unintended side effects? Presently, we expand on a theoretical perspective of memory control, wherein efforts to suppress episodic encoding or retrieval result in the systemic downregulation of the hippocampal memory system. We review evidence from multiple methodologies, highlighting a non-invasive means of inducing amnesia that casts a shadow over memory for unrelated events. By establishing the causes and consequences of the amnesic side effects associated with memory control, we argue it may be possible to harness hippocampal dynamics to promote more adaptive memory performance in the lab, clinic, and broader context of daily life.

ARTICLE HISTORY
Received 7 January 2017
Accepted 5 July 2017

KEYWORDS
Amnesia; cognitive control; inhibition; working memory; long-term memory; hippocampus

A case for dynamic memory

There is a certain appeal to the notion that memories are immutably stored somewhere in the brain. Names would never be forgotten, dates never confused, and details never lost. However appealing this idea may be, the evidence is lacking. Research has shown that memories are subject to modification and dependent upon context (Dudai, 2012; Hupbach, Hardt, Gomez, & Nadel, 2008). Consider the following example: when driving to work, remembering the customary route generally makes the commute easier. This dominant memory could, however, block recall of a lesser-used detour when required to navigate around unexpected road construction. To reach our goals, it is often necessary to navigate around distracting obstacles, both on the road and in memory.

When avoiding mnemonic distractors is impractical, inhibition may be recruited to decrease their accessibility and future competitiveness (Anderson, 2003; Hulbert & Norman, 2015). Here, inhibition refers to a reversible mechanism of executive control that acts to reduce the excitation state (and, in turn, the accessibility) of one or more memories (Anderson, 2007; but see MacLeod, Dodd, Sheard, Wilson, & Bibi, 2003, for a contrary opinion). As a remedy for interference, inhibition could be used to selectively target distracting memories while leaving others unaffected. To return to our example, this type of targeted inhibition could selectively dampen memory of the customary route to work, resulting in unfettered access to both the memory of the detour, as well as other non-target memories like the need to pick up a prescription at the pharmacy.

Such targeted control may work well under certain circumstances (reviewed by Storm & Levy, 2012), but when goals are more general – like entirely avoiding the retrieval of existing memories and/or the creation of new ones – it may be more adaptive to target the underlying memory system, rather than the unwanted memories themselves (Hulbert, Henson, & Anderson, 2016). Thus, as a result of trying to suppress memory, all memories supported by the affected region – even those unrelated to the target – may be rendered less accessible. Should the brain regions or networks supporting aspects of general memory functioning be disrupted, as in this form of systemic inhibition, one risks the unintended side effect of throwing out the proverbial baby with the no-longer-wanted bathwater. Returning to our driving example: if a passing blue '96 Chevy Impala threatens to dredge up memories of a past car accident, trying to clear the mind of that specific memory in order to stay focused on the road may "shut down memory lane" altogether.

This type of systemic inhibition theoretically is achievable via interneurons situated at the interface between prefrontal control regions and subcortical areas responsible for relatively localised functioning (Munakata et al., 2011; see *Active ingredients*, below). One subcortical structure of particular interest is the hippocampus, notable for its role in the encoding and retrieval of strong, contextualised memory representations (see Nadel & Hardt, 2011, for a review). While exogenous disruption of the hippocampus has long been known to compromise such mnemonic abilities, mounting evidence suggests that this region also may be endogenously modulated through prefrontal control

CONTACT Justin C. Hulbert jhulbert@bard.edu Psychology Program, Bard College, PO Box 5000, Annandale-on-Hudson, NY 12504, USA
© 2017 Informa UK Limited, trading as Taylor & Francis Group

pathways in attempts to reach mnemonic goals (reviewed by Anderson, Bunce, & Barbas, 2016; Anderson & Hanslmayr, 2014).

Exercising inhibitory control can be a lifesaver both literally and figuratively (for more on the adaptive benefits of forgetting, see Bjork, 1989; Bjork, Bjork, & MacLeod, 2006; Schooler & Hertwig, 2005; Storm, 2011), but the benefits must be qualified. For instance, tasks that depend on attending to previously irrelevant information may benefit from periods when inhibitory control is lacking (Amer, Campbell, & Hasher, 2016). Even when inhibition suits the current task, it has the potential to undercut future mnemonic goals once circumstances change (Hulbert & Norman, 2015; MacLeod & Hulbert, 2011; Schilling, Storm, & Anderson, 2014). Presently, we consider a theoretical perspective on memory control in which certain efforts to suppress episodic encoding or retrieval are achieved through the systemic attenuation of hippocampally dependent memory processes. After reviewing evidence that these circumstances give rise to side effects resembling temporary amnesia, we identify factors that may moderate such consequences. We argue that these side effects should be considered when attempting to optimise a dynamic memory system.

Treatment modalities

While the primary goal of any defence system is to block the intrusion of unwanted entities, certain defences may outlive their immediate usefulness and ensnare innocent bystanders. Consider the immune system. When working properly, it learns to identify and thwart external threats. Specific cues may precipitate a heightened immune response, whereas allergic reactions may require suppressive intervention. We argue that memory control can serve similar functions, engaging or disengaging neurocognitive systems that support memory depending on whether the situation demands encoding, retrieval, or suppression. And much like the immune system, the current modal state of the memory system could be expected to affect the way other mnemonic stimuli are processed around the same time.

According to this perspective on memory control, modulations in brain regions supporting memory should predict memory performance. In their review, Cohen et al. (2015) highlight that activations in the medial temporal lobe (MTL; particularly the hippocampus) during intentional and incidental encoding predict the later ability to remember those materials (e.g., Eichenbaum, Yonelinas, & Ranganath, 2007; Stark & Okado, 2003). Moreover, later remembering can be predicted by hippocampal activation well before the memoranda are originally presented (e.g., Park & Rugg, 2010).

Candidate cognitive mechanisms underlying such predictive effects include attention and the carryover of mnemonic states (Cohen et al., 2015). Much like attention selects the aspects of the environment that undergo additional processing, episodic memory tasks may lead to the adoption of a tonic state known as "retrieval mode" that biases individuals to interpret stimuli as episodic memory cues and/or deeply process retrieved information (for a review, see Rugg & Wilding, 2000). Attempting to remember which local gas station offers the lowest price may increase the likelihood that a roadside billboard hawking foot-long hoagies will serve as reminder that the leftovers intended for lunch are still sitting in the refrigerator. Moreover, the chance of mistaking the aforementioned billboard for a similar one may depend on whether unrelated memory cues leading up to the sign's appearance happened to be familiar.

Duncan, Sadanand, and Davachi (2012) exposed participants to sequences of objects that were new, old, or similar to old objects. When preceded by an old object, similar objects were more likely to be mistaken as old, compared to those preceded by new objects. They postulated that the hippocampal memory system can be biased toward a pattern completion mode – filling in incomplete representations from partial cues, as would benefit retrieval of old information – or a pattern separation mode – teasing apart similar memories, as would benefit new encoding. Accordingly, encountering an old object (like a sign seen previously) should bias the system in favour of a retrieval mode such that pattern completion would "fill in" inconsistencies between the next object in the sequence and a previously presented, similar object, risking their later confusion. Thus, a preceding memory decision can affect how the next stimulus is processed, with effects wearing off within seconds.

To recap: (1) hippocampal activity helps predict subsequent remembering in certain tasks; (2) mnemonic processing may be biased by preceding activities; (3) fluctuations in these domains take some non-zero amount of time. Might it also be possible to adopt a *suppression* mode that inhibits hippocampal processes, rather than a *retrieval/encoding mode* that engages them? We will next consider two cognitive tasks that lend support to this notion before considering their potential side effects. First, we examine a working memory task thought to engender strategies that indirectly affect long-term memory functionality before turning our attention to a task that is thought to more directly modulate the hippocampal memory system.

Memory depressors

Working memory tasks

The first type of long-term memory depressor emerges from what might be considered a somewhat unexpected domain: working memory. Long-term memory of the sort described above is commonly distinguished both functionally and neurally from working memory, a time- and capacity-limited buffer used for the active maintenance and manipulation of pertinent information (Baddeley, 2003). However, there are reasons to doubt early claims of this dissociation (reviewed by Ranganath & Blumenfeld, 2005).

Increased hippocampal engagement may be particularly conducive to maintaining multiple items in working memory (like keeping in mind four serially presented faces over a short delay), while reductions in hippocampal activity have been associated with lower working memory loads, such as the maintenance of a single item. Evidence for this relationship comes from both functional magnetic resonance imaging (fMRI) and intracranial electroencephalography measures (Axmacher et al., 2007). Even under low-load conditions, however, working memory for single items may be maintained in a hippocampally dependent manner (Axmacher, Elger, & Fell, 2009), simultaneously increasing the likelihood that they will be encoded in long-term memory (Axmacher, Elger, & Fell, 2009). These findings suggest that individuals may exert some control over which strategy they employ during working memory tasks and that hippocampal involvement interacts with long-term memory performance. To this point, Axmacher, Haupt, Cohen, Elger, & Fell (2009) discovered that maintaining multiple items in working memory (a process associated with hippocampal engagement) interfered with simultaneous long-term memory encoding of the faces (another hippocampally dependent task). In contrast, maintaining a single item (presumably without hippocampal involvement) was associated with improved recognition of the faces. Such observations suggest that certain modes of working memory maintenance may incidentally affect hippocampal functionality.

Further evidence for working memory-related hippocampal modulations (HMs) stems from the *n*-Back working memory task. This task requires participants to detect stimuli that were repeated *n* trials back in the presentation sequence. Take the following sequence of numerals, for example: 1-4-4-3-2-3. When $n = 0$, participants might be asked to simply press the button corresponding to the number currently on the screen. When $n = 1$, however, participants should wait to press anything until they get to the third numeral. Then they should press "4" because "4" happened to be presented *one* trial before. When $n = 2$, participants must refrain from pressing any buttons until the sixth trial when they should press "3" because "3" was presented *two* trials before. This 2-Back task requires participants to maintain the current numeral as well as the one immediately before it, as those numerals will become the relevant comparators on future trials. To limit interference from trials further back in the sequence, some have argued that it is advantageous to actively inhibit incorporation of these items into long-term memory (Jonides et al., 1997; Mullally & O'Mara, 2013). In the 0-Back condition such inhibition is thought to be largely unnecessary, as participants are merely responding to stimuli currently presented on the screen. Comparatively, the 2-Back task is associated with reductions in hippocampal activity and increases in frontal control regions, including the dorsolateral prefrontal cortex (e.g., Cousijn, Rijpkema, Qin, van Wingen, & Fernández, 2012).

If the observed HMs reflect efforts to control long-term memory encoding and limit distraction, one might expect a similar response in tasks that more explicitly involve such control. We next consider one such task.

The Think/No-Think paradigm

Anderson and Green's (2001) Think/No-Think (TNT) paradigm permits researchers to empirically investigate the consequences of repeated attempts to exert control over established long-term memories. Participants in a typical TNT task either retrieve (*Think* trials) or suppress (*No-Think* trials) well-practised memory associates when confronted with strong reminders – often drawing from word pairings learned to criterion (see Figure 1). Importantly, participants are instructed to fully attend the presented cues across conditions to curb differential processing. As such, they are trained never to look away from cues on the screen, regardless of condition.

This paradigm relies on the assumption that, even though demands on control may be acute, inhibition may leave residual effects (for alternative accounts of suppression-induced forgetting, see Raaijmakers & Jakab, 2013; Tomlinson, Huber, Rieth, & Davelaar, 2009). Provided that associates pose a threat of intruding into awareness and are controllable (Detre, Natarajan, Gershman, & Norman, 2013), No-Think associates should be rendered less accessible on a surprise final memory test relative to baseline items that were learned but neither retrieved nor suppressed in the intervening phase. Below-baseline impairment, termed "suppression-induced forgetting", is counter-intuitive in that repeatedly cueing a memory – rather than facilitating later recall of the target – can, with a simple change in instructions, encourage forgetting beyond what would be expected over a filled delay without any reminders. Despite wide-ranging individual differences (Levy & Anderson, 2008), non-clinical samples of young adults demonstrated an 8% average suppression-induced forgetting effect (Anderson & Huddleston, 2011).

Functional neuroimaging has helped identify a network of brain regions associated with memory suppression. Compared to retrieval, blood-oxygen-level dependent (BOLD) fMRI activity during No-Think trials is enhanced in many cognitive control regions associated with detecting and resolving forms of interference. These activations, reviewed by Anderson and Hanslmayr (2014), include the dorsolateral and ventrolateral prefrontal cortices (dLPFC and vLPFC, with a tendency to be right lateralised) and the anterior cingulate cortex (ACC). Such findings support the assertion that attempts to suppress memories involve more than merely disengaging brain regions associated with retrieval. The predictive relationship between dLPFC and suppression-induced forgetting (Anderson et al., 2004), in combination with neuropsychological data relating prefrontal lesions to deficits in memory inhibition (Conway & Fthenaki, 2003), further suggests that the dLPFC may play a key role in successful suppression of unwanted memories.

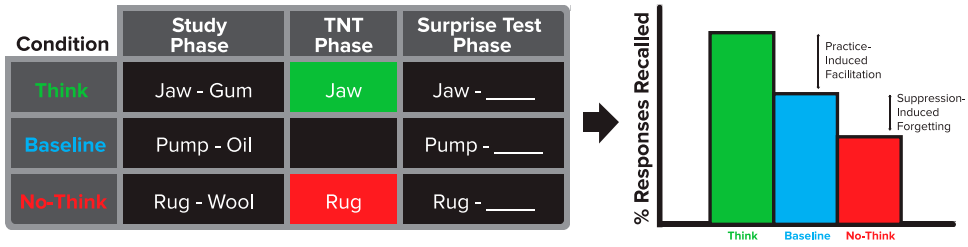

Figure 1. The TNT paradigm.

Notes: Participants learn word-pair associates to criterion in the initial study phase. In the following TNT phase, participants are presented with cues from the No-Think (now presented in red) and Think (now presented in green) conditions, which respectively indicate that participants are to suppress the learned associates or to instead retrieve them. In the surprise test phase that follows, participants' memory for all learned associates (including baseline items that were not cued during the TNT phase) is assessed. The right panel illustrates typical facilitated recall for the practised Think associates, along with a below-baseline impairment for No-Think items. This latter finding, thought to reflect the aftereffects of inhibition, is called suppression-induced forgetting.

Attempts to suppress retrieval are also associated with brain regions that are *less active* during No-Think compared to Think trials. Notably, robust reductions in hippocampal activity are regularly observed across a range of materials, including words, negatively valenced scenes, visual objects, and imagined future episodes (see Anderson & Hanslmayr, 2014, for a review; Benoit, Davies, & Anderson, 2016). Not only do these deactivations often predict suppression-induced forgetting (Benoit & Anderson, 2012; Depue, Curran, & Banich, 2007), effective connectivity analyses have revealed a top-down modulatory influence of the dLPFC on the hippocampus (Benoit & Anderson, 2012; Benoit, Hulbert, Huddleston, & Anderson, 2015; Benoit et al., 2016; Gagnepain, Henson, & Anderson, 2014). Furthermore, negative coupling between the dLPFC and the hippocampus often predicts suppression-induced forgetting on subsequent memory tests (Benoit & Anderson, 2012; Benoit et al., 2016) and reduced intrusion frequency (thinking of a No-Think associate) across the TNT phase (Benoit et al., 2015). Evidence of inhibition, apparent both during and after suppression attempts, suggests that participants can strategically modulate hippocampally dependent mnemonic processes to either access desirable memories or block unwanted retrieval.

Interpreting negative test results from memory depressors

Do TNT-induced HMs truly reflect an inhibited mnemonic state or simply the relative engagement of retrieval during Think trials? Ideally, a separate neuroimaging baseline could parse these possibilities. Accordingly, hippocampal deactivations have been noted relative to passive fixation (Depue et al., 2007; Levy & Anderson, 2012). However, during passive fixation, the hippocampus is free to engage in incidental encoding and/or retrieval, potentially inflating baseline measures and reducing the relative activation of more demanding comparison conditions (Stark & Squire, 2001). Hence, suppression-related deactivations might simply reflect a byproduct of restrictions placed on mind wandering and spontaneous memory processing. By this account, cognitive demands should decrease with practice, and consequently, mind wandering (along with hippocampal activity) should increase (Banich et al., 2009). But per cent signal change from fixation baseline appears to become more negative over suppression attempts, undermining a simple restrictive mind-wandering account (Depue et al., 2007).

Together, below-baseline hippocampal deactivations, negative coupling with vascularly distinct prefrontal regions, and correlations with both subjective (involuntary intrusions) and objective (suppression-induced forgetting) measures of control provide strong convergent evidence of a functional account of hippocampal downregulation during memory suppression. Still, the assumption that negative BOLD responses (NBRs) observed during control tasks necessarily reflect inhibited processing remains vulnerable to well-founded criticisms (Aron, 2007).

Side effects

In response, we argue that further traction is gained by examining the side effects that are predicted by a functional account of suppression-related hippocampal NBRs. Here, we return to our central question: do attempts to keep the mind tidy and free of mnemonic distractions result in unintended behavioural consequences? Below we discuss two approaches to determining whether healthy individuals attempting to exert inhibitory control over memories exhibit symptom profiles associated with hippocampal amnesia.

Hippocampal modulation (HM) paradigm

Recall that the TNT paradigm helps reveal the direct effects of attempting to not think of memory associates in the face of strong reminders. It becomes harder to recall the previously suppressed associates even after circumstances change and participants are incentivised to recall them (Anderson & Green, 2001). Some have suggested that attempts to banish unwanted thoughts can lead to their ironic rebound (Wegner, 2009) and negative clinical outcomes, including the persistence of post-traumatic stress symptoms in some individuals (Bomyea & Lang, 2016; Dalgleish, Hauer, & Kuyken, 2008; Ehlers, Mayou, & Bryant,

1998; Myers, 2010). Others have reported associations between positive real-world outcomes and better suppression abilities, including fewer distressing memory intrusions and less severe post-traumatic stress disorder symptoms (Catarino, Küpper, Werner-Seidler, Dalgleish, & Anderson, 2015; Fawcett et al., 2015; Hertel & Gerstle, 2003; Joormann, 2010; Küpper, Benoit, Dalgleish, & Anderson, 2014; Marzi, Regina, & Righi, 2014; Streb, Mecklinger, Anderson, Lass-Hennemann, & Michael, 2016). Assuming – at a minimum – that the ability to suppress distracting memories may be adaptive for reaching short-term goals, ancillary questions remain regarding potential mnemonic side effects.

In addition to suppressing retrieval of No-Think associates, functionally meaningful hippocampal deactivations should also impair other hippocampally dependent functions, including the formation of stable, contextualised memories for new events. Therefore, even after the system has regained functionality, unrelated events before or after suppression attempts should remain unavailable. It would be as if attempting to suppress the memory of an earlier car accident when confronted with an evocative '96 Chevy Impala also impairs the formation of a new memory for a parade float seen a block later – even if observed with rapt attention.

Consider another metaphor for memory suppression and its side effects: if you hammer away at a single part of a frozen pond containing an unwanted bit of detritus, the surrounding ice may crack and sink into the water. Established memories contained within those submerged shards are rendered less accessible, much as memory for direct targets of suppression are impaired in the TNT paradigm. Moreover, your ability to capture and solidify new memory impressions is disrupted until you stop hammering and the surface reforms. Efforts to maintain a spotless pond, like a spotless mind, are not without unintended consequences.

To test these predictions in the laboratory, we developed the HM paradigm (Hulbert et al., 2016). Participants were initially trained to criterion on a series of word pairs. Interpolated between a pseudo-randomised order of Think and No-Think cues in the subsequent main phase were unrelated, novel stimuli. Participants were told these materials were "distractor tasks" within a larger attention test. As such, they were instructed to complete the distractor tasks with due diligence before returning to the purported primary task of interest (TNT trials). In fact, these distractor tasks (e.g., a semantic judgment relating to a picture of a peacock standing in a parking lot) were designed to orient participants to memoranda that later would be probed on a surprise memory test (being asked to recall the central object, given a photo of the empty parking lot). The results could then be analysed based on the number of surrounding "doses" of suppression to which these "bystanders" had been exposed. Before participants were debriefed and dismissed, their memory for the learned word-pair associates was also tested, and a post-experiment questionnaire measuring adherence to instructions and strategy use was administered (Anderson & Huddleston, 2011).

Bounding every critical bystander during the main phase were blocks of odd/even parity judgments. Parity judgments are independently known to yield hippocampal NBRs relative to more passive fMRI baselines (Stark & Squire, 2001) but could not drive any observed memory effects in the HM paradigm because they were administered uniformly before and after all bystanders, regardless of condition. Rather, they were included to match the local shifts into and out of the incidental encoding periods. Without the buffers, it was reasoned, participants might dwell on thoughts about bystanders into No-Think trials as a way of diverting their thoughts from unwanted responses. This would be less advantageous during Think trials, which require focus on retrieval. Without buffers, bystanders presented before suppression events might receive a disproportionate amount of rehearsal time, selectively inflating recall and compromising critical comparisons.

Specifically, we sought to determine whether epochs of memory suppression would render individuals relatively less able to later remember contextualised information about bystander events. In contrast to the standard TNT suppression-induced forgetting effect, the forgetting here referred to a memory impairment for novel events. Namely, the "amnesic shadow" was calculated as the difference in subsequent memorability between bystanders surrounded by two epochs of suppression and those surrounded by zero suppression epochs. Figure 2 depicts the expected effect and an overview of the HM paradigm.

Across multiple experiments, we demonstrated that when unrelated bystanders were interpolated between epochs of memory suppression, subsequent cued recall and source recognition of those bystanders were impaired relative to items interpolated between Think trials (Hulbert et al., 2016). This effect generally became reliable for distractors presented later in the presentation schedule, after participants had ample practice (see section on *Practising suppression*). The amnesic shadow survived a 24-hour retention interval, generalised across bystander materials (including words and images), and persisted despite the presence of the low-level parity buffer task. Moreover, self-reported perseverative thoughts about the distractors into subsequent trials was low and comparable across conditions, suggesting that the amnesic shadow is not likely a direct result of differences in task-switching costs or rumination.

Notably, when compared to a non-episodic baseline task (phonological rehearsal of a visually presented nonsense words), No-Think trials still disrupted later cued recall of bystanders. Compared to the same baseline, surrounding bystanders with hippocampally dependent Think trials tended to improve bystander recall, though not significantly. This suggests that the amnesic shadow reflects more than just a relative benefit imparted to bystanders surrounded by Think items.

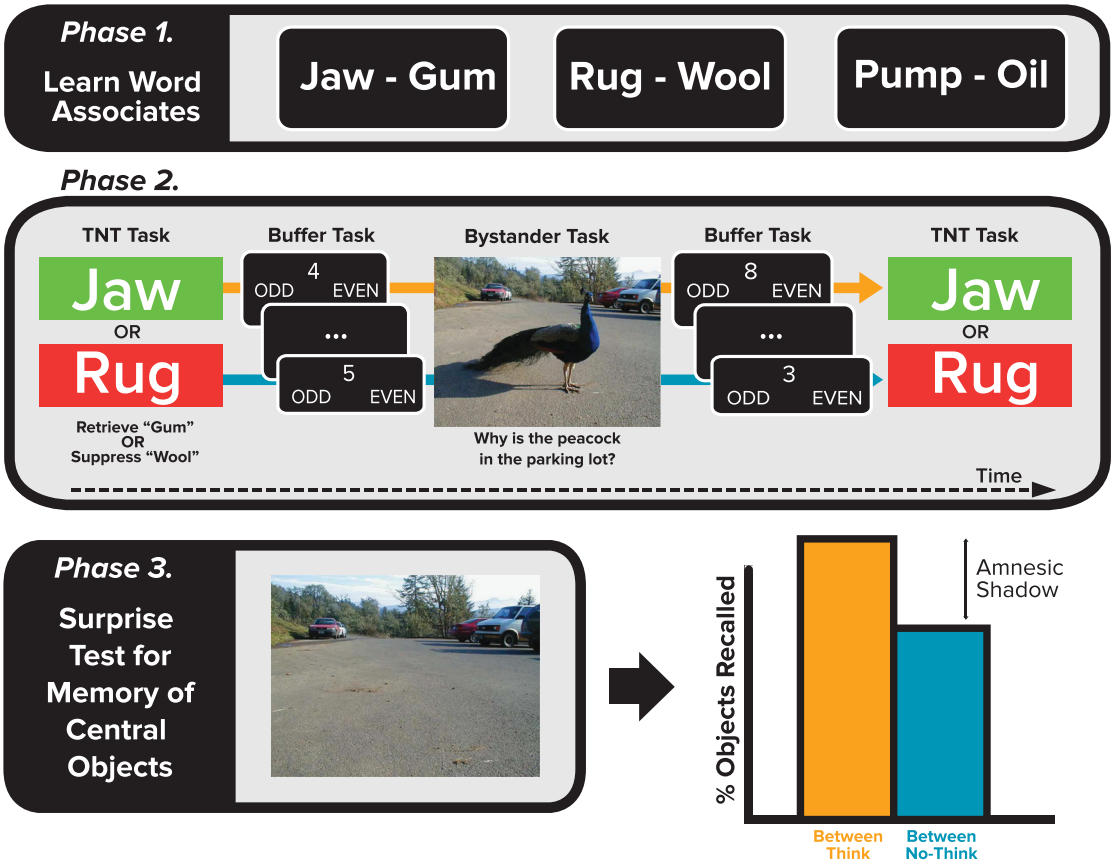

Figure 2. The HM paradigm.

Notes: As with the basic TNT paradigm, participants first learn word-pair associates to criterion. However, in addition to intermixed Think (green) and No-Think (red) cues, participants in Phase 2 occasionally encounter novel "bystander" materials (like an image of a peacock in a parking lot), for which they have to make a semantic judgment (like inventing an explanation for the appearance of the central object in the pictured location). A series of even/odd buffer judgments made before and after bystanders match their immediate task context across conditions. Of interest is whether having surrounded these unrelated materials with suppression (No-Think trials) impairs later memory, causing an "amnesic shadow" relative to those bystanders that had been surrounded by retrieval (Think trials). This is assessed in Phase 3 by asking participants to recall the associated object for each background scene, for example. Typical results are presented.

Difficulty did not seem to drive the amnesia either. The magnitude of the amnesic shadow was unrelated to the subjective disparity in difficulty across the surrounding trials. More convincingly, a strong difficulty manipulation that did not involve the intentional suppression of memories failed to produce a reliable amnesic shadow. Furthermore, substituting the unwanted associates with other memories (a strategy that makes demands on hippocampally dependent retrieval processes) left bystander memory unimpaired, in contrast to a strategy in which participants are asked to stop retrieval entirely (see *Alternative treatments*).

Additional experiments provided evidence that source recognition, but not necessarily item recognition, is subject to the amnesic shadow (Hulbert et al., 2016). In essence, participants were able to recognise bystander objects as familiar even if they had been sandwiched between two No-Think trials. They primarily had trouble recollecting the context in which those bystanders had been presented. The behavioural specificity of the amnesic shadow to strong, contextually bound memories deserves further study. Such results may be accommodated by certain models of the MTL ascribing recollection, contextualisation, and/or the processing of strong memories predominately to the hippocampus (Aggleton & Brown, 1999; Eichenbaum et al., 2007; Ranganath, 2010).

We observed signs of both anterograde and retrograde amnesia across experiments. This pattern of results would be expected if disruptions to hippocampally dependent memory functioning outlast individual No-Think trials (impairing the memorability of subsequent events and explaining the anterograde effect) and truncate continuing memory stabilisation processes after bystanders have already been presented (explaining the retrograde effect; Ben-Yakov & Dudai, 2011). Our understanding of why and when an anterograde or retrograde effect dominates will likely be illuminated through further investigation of the temporal boundary conditions of the HM effect (see *Expiration*).

Finally, an fMRI version of the HM paradigm revealed the expected deactivation of the hippocampus, along with a correlated pattern of activation in a right-lateralised frontal control region (Hulbert et al., 2016). Importantly, this neuroimaging study also demonstrated that the extent to

which the bilateral hippocampus was modulated by TNT trials predicts individual differences in the magnitude of the amnesic shadow.

Together, the above findings suggest that directly suppressing unwanted memories modulates hippocampal activity and that this modulation is functionally relevant for the ability to remember novel information.

2-Back to forget

Motivated attempts to directly suppress unwanted memories are not the only activity associated with HMs. If the 2-Back working memory task encourages an interference-reducing strategy by which long-term memory encoding is inhibited through systemic suppression, hippocampally dependent memory for bystander materials embedded between blocks of the *n*-Back task should be worse when $n = 0$ relative to when $n = 2$. Mullally and O'Mara (2013) conducted a series of experiments testing this logic. Participants in the first experiment studied 15 sequentially presented words, followed by a free recall test. This helped establish comparable memory across groups randomly assigned to the 0- and 2-Back conditions. After training, the groups completed a single block of their assigned working memory task. Once ready to advance (the transitions between *n*-Back task blocks and encoding/retrieval tasks were self-paced), participants studied and were tested on a second list of words. This process was repeated such that, eventually, three lists of words were presented and tested, each following a block of the *n*-Back task. The authors predicted that recall would be worse for lists that were presented and tested after the 2-Back blocks.

They also predicted that a recognition test coming after a final block of the *n*-Back task would reveal the selective nature of this memory impairment. To address this, participants were further subdivided into those taking an explicit recognition test and those taking an implicit version. The authors predicted that the explicit test – recognising studied words amongst lures – would reveal impaired memory for items presented after 2-Back blocks. The implicit test depended on perceptual priming for visually degraded words, feasible even without a functioning hippocampus (Gabrieli, 1998). Therefore, to the extent that the brain regions supporting implicit item recognition tend not to be disrupted by the 2-Back task, there should be no difference in implicit recognition across the 0- and 2-Back conditions.

As predicted, performing a 2-Back task before each experimental block of list learning reduced overall free recall performance relative to that of the 0-Back condition. The group difference was individually reliable for each non-baseline list and did not interact with word list. Moreover, there were no recall differences between the first and last words presented (discussed further under *Expiration*). These results suggest that the mnemonic shift induced by performing the 2-Back task had lasting effects consistent with the adoption of a sustained strategy. Moreover, the impairment appeared to be specific to a hippocampally dependent memory.

A follow-up experiment demonstrated the generalisability of these findings to learning associations between faces and names. Results were consistent with the first experiment. Notably, participants completed a visual attention task as a distractor between each round of face–name encoding and name retrieval. Findings from this attention task revealed no group differences, indicating that the *n*-Back manipulation largely spared visual attention. Reversing the general sequence of events revealed that a hippocampally dependent associative memory task impaired subsequent 2-Back working memory performance.

As with the HM paradigm, these results suggest that, when tasks separated in time drive individuals to adopt conflicting modes of the hippocampal memory system, the hippocampally dependent components of the secondary task are wont to suffer.

Potential interactions

To recap, down-regulating hippocampal activity may be akin to non-invasively (and reversibly) lesioning the hippocampus. The results from both memory depressor tasks also hint at potential tradeoffs associated with up-regulating the system: hippocampally dependent Think trials resulted in a trend toward the improved memorability of bystander events. Face–name encoding and recall similarly dampened 2-Back working memory performance. These effects presumably reflect a lingering mnemonic state that would be more conducive to performing hippocampally dependent memory tasks (like studying for an exam). We also presume that the presentation of novel relational information (bystanders) should engage hippocampal encoding processes, potentially limiting the effectiveness of suppression attempts that follow.

What other factors might influence the scale and scope of the unintended side effects of memory control? In this section we consider how practice, attention, strategy, and timing might be expected to interact with the amnesic shadow.

Practising suppression

An ounce of prevention: preparatory effects

Behavioural and electrophysiological evidence substantiate the potential value of preparatory cues in inhibiting unwanted memories (Hanslmayr, Leipold, & Bäuml, 2010; Hanslmayr, Leipold, Pastötter, & Bäuml, 2009). Participants who were warned of the upcoming TNT trial type demonstrated reliably greater suppression-induced forgetting than those who did not receive diagnostic precues. Most variants of the HM paradigm incorporated precues (Hulbert et al., 2016); however, the extent to which precues affect properties of the amnesic shadow in this paradigm has not yet been established.

A pound of cure: reactive control

Subjective reports of successful memory control increase in tandem with suppression attempts. This suggests that, while difficult at first, many individuals are able to effectively exert control over their memories after practice (Levy & Anderson, 2012). Critically, the rate of change in participants' perceived success is related to suppression-induced forgetting. Intrusions of unwanted TNT associates in this paradigm were assayed on a trial-by-trial basis using a three-point scale anchored to whether the associate never, briefly, or came to mind often. Levy and Anderson (2012) also demonstrated that intrusions of unwanted memory associates during No-Think trials triggered a far more pronounced downregulation of hippocampal activity than did suppression attempts unaccompanied by subjective intrusions. The extent of intrusion-related hippocampal downregulation predicted suppression-induced forgetting; no such correlation was observed with modulation during non-intrusion trials.

Not only were intrusion trials marked by deeper hippocampal deactivations, they were also more spatially diffuse, spanning anterior and posterior regions of the hippocampus and extending to the surrounding MTL cortex. These results indicate that the reactivation of unwanted memory traces may be critical for both suppression-induced forgetting and deactivations of the wider MTL (Anderson et al., 2016). Independent evidence suggests that moderate reactivations (in contrast to reactivations of associates that are either so weak they pose no threat of intruding or are too strong to control given the available resources) may be most vulnerable to lasting memory impairments (Detre et al., 2013). They may also induce the greatest mnemonic side effects.

An integrated approach

Depue et al. (2007) proposed a two-stage process of memory suppression in which practising memory suppression eventually leads to hippocampal downregulation. Early signs of below-baseline activity in visual cortex and the thalamus, in tandem with positive vLPFC activations during No-Think trials, characterised the first stage. With minimal practice (during the first six attempts to suppress emotional memories), hippocampal and amygdalar BOLD activity hovered above that of a fixation baseline, decreasing steadily over the course of the experiment until it finally became reliably negative in the fourth quartile. Butler and James (2010) reported a similar pattern of decreasing hippocampal activity across attempts to suppress neutral items. Depue et al.'s (2007) hippocampal deactivation, itself associated with reciprocal increases in the dLPFC (namely, right middle frontal gyrus), was taken as a sign of a qualitative shift in control pathways.

Depue et al. (2007) interpreted the first phase of control as a stopgap attempt to terminate the reinstatement of sensory percepts. The second stage, they argued, represents an attempt to directly suppress access to memory representations and associated emotional components supported by the hippocampus and amygdala, respectively – after sufficient practice, participants need not rely on sensory suppression anymore. Only at this later stage of the experiment was participants' ability to modulate hippocampal and right middle frontal gyrus activity predictive of suppression-induced forgetting. This relationship revealed itself in the last experiment quartile, when group-level hippocampal activity was at its nadir.

The above findings suggest that, given proper practice, individuals may learn to enhance control over unwanted memory intrusions. Any mnemonic side effects that arise on account of this modulation should also increase in magnitude as participants shift from the first to the second phase of this process. Accordingly, the predicted build-up in the amnesic shadow was observed in all of the behavioural experiments employing the HM paradigm with direct-suppression instructions and a cued-recall test (Hulbert et al., 2016).

Theoretically, more distributed suppression training with a diverse array of memory associates (and, perhaps, neurofeedback) could yield increasingly stable strategies for effectively controlling memories through systemic modulation. Beyond establishing positive transfer to situations requiring cognitive control, such studies could also identify negative transfer with respect to the establishment of stable, contextualised bystander memories.

Life experience

Suppression practice need not take place within a single laboratory experiment. If memory control is a skill honed with age and experience, then young children may be especially poor at regulating thoughts. This status is compounded by the relatively slow development of the prefrontal cortex. Over the years from 8 to 12, children become increasingly more adept at controlling their memories, as evidenced by a progressively more adult-like pattern of both suppression-induced forgetting (Paz-Alonso, Ghetti, Matlen, Anderson, & Bunge, 2009) and coupling between the lateral prefrontal cortex and hippocampus (Paz-Alonso, Bunge, Anderson, & Ghetti, 2013).

Older children have also accumulated more life experiences (and memories) than their younger counterparts. Some of those extra memories may be unpleasant in nature, providing more opportunities to practise their control abilities. Carried forward, this logic suggests that individuals who have experienced more unfortunate life events would have more reasons to try to prevent unwanted memory intrusions. Were they equipped and motivated to regularly exercise this control process over an extended period, such practice could help them hone successful general coping strategies, including suppression. While it would be unethical to expose individuals to trauma to test this prediction, longitudinal natural experiments may speak to such claims.

Limits in practice

We now consider some general limits to the effects of suppression practice in both the short- and long-term. For example, an individual who has become so successful at suppressing unwanted memories eventually may discover that reminders no longer threaten to incite unwanted retrieval. As such, the need for cognitive control may diminish along with the mnemonic side effects. This is one reason to expect a non-monotonic pattern in the magnitude of the amnesia effect for bystanders as a function of suppression practice.

Future work examining the trainability of memory control should account for some of the challenges encountered in the working memory domain (Redick et al., 2013). And given the association between working memory capacity and memory control (Aslan & Bäuml, 2011; Brewin & Smart, 2005), researchers may gain additional traction by constraining the populations from which they sample (e.g., by initially focusing on individuals with high working memory spans).

Differential diagnoses

Fatigue

Given that suppression-induced forgetting is typically augmented by increasing the number of suppression attempts (Anderson & Green, 2001), one might also expect that longer epochs of suppression would similarly yield greater forgetting. Nevertheless, Lee, Lee, and Tsai (2007) discovered that increasing the duration of the TNT trial from 3 to 5 s eliminated the suppression-induced forgetting effect. Whether this result owes to the longer duration of each individual TNT trial or the overall lengthening of the phase, these data hint at an upper limit to people's ability to sustain cognitive control continuously. In fact, instances of cognitive fatigue resulting from the prolonged exertion of self-control are widely reported (Hagger, Wood, Stiff, & Chatzisarantis, 2010; but see Hagger & Chatzisarantis, 2016). As discussed above, practice may lead to more efficient control and increased tolerance, but it may need to be distributed in manageable sessions, in part to limit fatigue.

Could fatigue explain the amnesic shadow itself? If so, performance might be expected to uniformly worsen across blocks of the procedure. However, evidence suggests that individuals generally get better at controlling intrusions over the duration of a memory suppression task, not worse (Levy & Anderson, 2012). And independent tasks performed between epochs of cognitive control and the bystanders, like Mullally and O'Mara's (2013) visual attention task and the parity buffers in the HM paradigm (Hulbert et al., 2016), revealed uniformly strong performance across blocks.

To test whether memory suppression (compared to retrieval) differentially sapped attentional resources available for bystanders and resulted in higher rates of forgetting, we swapped out the No-Think task for a "Think-Harder" task, in which participants were required to retrieve multiple items, compare and manipulate them, and regularly update memory (Hulbert et al., 2016). Despite clear differences in participants' perceived difficulty across the two tasks, no parallel difference was observed in subsequent recall for the bystanders. As such, these results cast doubt on difficulty as an explanation for the amnesic shadow.

Moreover, we observed steady growth in the amnesic shadow across experiment blocks in the HM paradigm. This growth was predicted a priori based on data indicating that hippocampal activity decreases over the course of a TNT experiment while frontal engagement increases (Depue et al., 2007). Together, these findings are consistent with the notion that cognitive control is ramping up, not down, in the later blocks of the TNT phase. In other words, we have reason to suspect that cognitive resources have not yet been depleted.

In contrast to the growth of the amnesic shadow in the HM paradigm, recall for materials presented after the 2-Back task were consistently impaired across blocks, indicative of a quickly adopted and sustained strategy (Mullally & O'Mara, 2013). Despite this difference (discussed further below), both control paradigms seem to produce mnemonic side effects that are unlikely to have been driven by fatigue.

Nonetheless, we do not assume that participants are indefatigable. While some variants of the HM paradigm described above lasted over two hours, presumably, there is some upper limit to one's willingness (or ability) to engage in a cognitively demanding set of tasks. This represents another reason to expect a non-monotonic pattern in the amnesic shadow over long-enough timelines.

Inattention

Perhaps participants differentially attended Think and No-Think trials, accounting for the amnesic shadow. Performance on the buffer task and the presence of retrograde amnesia undercut this diagnosis, as do additional pieces of evidence. All 279 participants were trained to keep their eyes and attention locked on the presented stimulus materials, including the bystanders and surrounding TNT cues (Hulbert et al., 2016). A post-experiment questionnaire indicated that they almost never endorsed a strategy of diverting their eyes away from the No-Think cues to avoid thinking of the associates. Individual differences in self-reported use of this strategy contributed to a near-zero correlation with the size of the amnesic shadow. Moreover, the overall effect remained highly reliable in a conservative follow-up analysis excluding all participants who reported any tendency to divert their eyes. Future work incorporating eye-tracking could further substantiate that participants attend the bystanders equitably across conditions, as existing data suggest.

Critically, engagement with the bystander tasks themselves did not appear to vary by condition in the HM

paradigm. Responses collected during bystander orienting tasks were largely comparable across conditions, both in terms of actual responses and reaction time, suggesting that attentional differences were not responsible for the observed amnesia (Hulbert et al., 2016).

Attention and intention to encode bystanders may interact. Although participants were instructed to pay attention to the bystanders in the HM paradigm (as all available evidence suggests), they were unaware that their memory for these items eventually would be tested. Had participants adopted an intentional encoding strategy, they may have had reason to shift resources away from the "primary" TNT tasks and, in turn, temper HMs. Without additional controls, altering the incidental nature of bystander encoding and/or the balance between the number of bystanders and TNT trials may lead to an encoding strategy that moderates the amnesic shadow. Such claims have yet to be tested.

Mullally and O'Mara's (2013) participants, in contrast, were aware that they should try to memorise the words and face–name associations between the working memory task, and the repeated encoding-test cycles further reinforced this emphasis. Their finding of an amnesic side effect suggests that, at least under certain conditions, intentional encoding remains vulnerable to something akin to the amnesic shadow.

Alternative treatments

Generic strategies

Is the mere intention to avoid an unwanted memory sufficient to produce an amnesic side effect? Evidence from the HM paradigm suggests that it is not. Participants using a suppression strategy that is reliant on hippocampally dependent mnemonic processes – namely one involving the retrieval of a substitute memory – showed no sign of an amnesic shadow for bystanders (Hulbert et al., 2016). These results stood in contrast to the results from their matched counterparts who were instead asked to stop retrieval entirely, a so-called direct-suppression strategy. Independent evidence suggests that even when suppression-induced forgetting across the two strategies is equated, HM is only observed for participants under direct-suppression instructions (Benoit & Anderson, 2012). Such a dissociation is consistent with the hypothesised difference between targeted and systemic suppression (Anderson & Huddleston, 2011; Munakata et al., 2011), and it accurately predicts that a strategy compatible with systemic suppression is more apt to disrupt the formation of stable, contextualised bystander memories than thought substitution (Hulbert et al., 2016).

Similarly, tasks that demand memory control while simultaneously drawing on hippocampal dependent processes – for example, selective retrieval in the retrieval practice paradigm (Kuhl, Dudukovic, Kahn, & Wagner, 2007; Wu, Peters, Rittner, Cleland, & Smith, 2014) or list-2 learning in the list-method directed forgetting paradigm (Manning et al., 2016) – are not expected to give rise to an amnesic shadow. Still, the wide range of strategies endorsed by participants in memory control experiments (Levy & Anderson, 2008) invites further exploration of potential interactions with the amnesic side effect.

Adherence

Not all variants of the TNT task require participants to employ a consistent strategy. It is not uncommon for uninstructed participants to try one suppression strategy and then switch to another if they find the first is ineffective, perhaps accounting for some of the observed practice effects. Even if participants adopt a consistent strategy, the TNT paradigm generally demands great flexibility toggling between suppression and retrieval strategies on a trial-by-trial basis, as the order of Think and No-Think trials is typically randomised.

Mullally and O'Mara's (2013) design involves far more predictability: it is blocked; it involves a set number of memoranda in each encoding/recall block; the duration of the *n*-Back task blocks is fixed; and the switch from one task block to the next is self-paced. The predictability of this experiment and the extra training on the *n*-Back task before experimental blocks may have encouraged the adoption of a relatively stable strategy, biased toward performance on the *n*-Back task (in their first two experiments).

Still, their third experiment demonstrated that extra rounds of a hippocampally engaging task could flip the bias. Although the authors found no sign that the magnitude of the effect changed across (or within) blocks, they demonstrated the potential to induce more rapid strategy shifts. And more rapid strategy shifts theoretically would yield shorter windows of the amnestic side effect. Conversely, certain alterations to the HM paradigm could widen the window (e.g., blocking the TNT cues by condition or altering the ratio of Think to No-Think trials). At present, we believe the results garnered from these two paradigms are compatible with the same underlying model of systemic modulation.

Warning labels

Expiration

How long do the amnestic effects persist? If the hippocampal downregulation is indeed under voluntary control, this question may be ill-posed. The modulation could be sustained indefinitely – or at least until fatigue sets in. Mullally and O'Mara's (2013) findings suggest that bystanders presented shortly after a 2-Back working memory task block are subject to the same impairment as bystanders presented later in the encoding list. These results are consistent with the adoption of a strategy conducive to sustained downregulation of the hippocampal memory system (one that might even last through the final memory test) – a prediction awaiting further neuroimaging support. In contrast, the HM paradigm's mix of Think, No-Think, and incidental encoding trials may induce more transient fluctuations.

Our work with the HM paradigm clearly demonstrates an amnesic shadow over bystanders presented at a temporal lag of around 5 s – the average duration of the parity buffer task (on either side of the bystander) across our behavioural experiments (Hulbert et al., 2016). The exact timing of the buffer task was not always fixed in these experiments, partly to minimise expectancy effects. While the limited range of durations prevented a detailed examination of the temporal dynamics in these experiments, suggestive results stemmed from the extended range of jittered buffer durations (up to 22 s) in the fMRI version of the task. Our analyses for this experiment were focused on the relationship between individual differences in HMs and the amnesic shadow. But in terms of the average magnitude of the amnesic shadow across the entire sample, we reported that the effect was numerically muted for bystanders subjected to lags greater than 5 s before and after them. Even so, removing bystanders presented amidst long lags did not change the nature of the main correlation analysis.

Establishing the period in which bystanders are most vulnerable to HM represents an important future direction. But if memory suppression also affects lingering hippocampal representations of memories that were originally (re-)activated further back in time, bystander items presented early in the experiment would be subjected to further amnestic agents (No-Think trials), regardless of the immediate hippocampal state in which they were introduced. As such, the observed amnesic shadow for items presented early would be diminished relative to those presented later and relatively spared from the introduction of such noise (reflecting a purer contrast based on the immediate TNT context). This represents another possible reason for the observed growth of the amnesic shadow in the HM paradigm across experiment blocks.

Prognosis

Much of the evidence reviewed above is compatible with the notion that both the 2-Back (working memory) task and the TNT (with direct-suppression) task can modulate the hippocampal memory system and, in turn, produce an amnesic shadow over temporally proximate bystander materials. There are some notable differences, however. The long-term memory control strategy in the direct-suppression task is made explicit, with adherence to the strategy assessed regularly throughout the HM paradigm (Hulbert et al., 2016). Less is known about the HMs sometimes observed across different working memory tasks, how they might relate to strategies aimed at controlling long-term memory encoding (and retrieval), and the extent to which they predict bystander forgetting. Whether or not participants in a given working memory task draw on hippocampally dependent processes appears to rely on factors beyond load (e.g., Axmacher, Elger, et al., 2009). Future work may help define these factors and establish whether working memory paradigms also are capable of producing the type of retrograde amnesic shadow regularly seen for bystanders presented before direct-suppression attempts (Hulbert et al., 2016). Such endeavours may also address the apparent differences in the immediacy and sustainability of the amnesic shadow across the *n*-Back and TNT paradigms.

One recent report marries elements of both the *n*-Back and TNT paradigms. Compared to a condition in which participants had to identify a designated vowel in an on-screen TNT cue (0-Back condition), participants who had to report a vowel from two trials before (2-Back) tended to exhibit enhanced – rather than impaired – recall for No-Think items on a final test designed to tap non-associative forgetting (Noreen & de Fockert, 2017). While suppression-induced forgetting remained apparent when tested with the original cue words, the lack of an impairment on the independent test might seem at odds with the HM account: to the extent that direct suppression and the 2-Back task individually are thought to drive down hippocampally dependent memory processing, their combination might be expected to produce more – not less – suppression-induced forgetting. We speculate that the 2-Back task in Noreen and de Fockert's paradigm was incompatible with the type of hippocampal downregulation strategy suggested by Mullally and O'Mara (2013) because participants were also forced to complete the 2-Back task during Think trials – demanding overt retrieval from long-term memory. Along with the costs of switching between vowel identification and direct suppression during No-Think trials, the emphasis on low-level cue features (vowels) may have focused inhibition on word form rather than the deeper representation of No-Think associates or the whole hippocampal memory system (Depue et al., 2007). As such, reminders of the deeper representation may have led to facilitation of the conceptual representation assessed on the independent final test. Paradigmatic variations with interpolated bystanders and neuroimaging may help to resolve the distinctive effects of the constituent tasks.

Practical implications

The mnemonic side effects discussed in this paper promise to reveal much about the hippocampal memory system and the mechanisms involved in motivated suppression. Such work stands to inform theories of impaired control over intrusive thoughts and behaviours in clinical populations, such as those with post-traumatic stress disorder, depression, attention-deficit/hyperactivity disorder, obsessive-compulsive disorder, and addiction (for further discussion, see Hertel, 2007; Nørby, 2017). Faced with the prospect of unwanted memory intrusions, individuals recovering from trauma may be especially motivated to regularly exclude certain memories from awareness via intentional suppression (Bomyea & Lang, 2016; Ehlers et al., 1998). Many of these individuals also exhibit memory impairments extending to neutral memories unrelated to the trauma

(Brewin, 2011; Guez et al., 2011, 2013). Systemic memory suppression may contribute to such deficits: as the theoretical perspective outlined in this paper would predict, hippocampally dependent memories often suffer disproportionately (Guez et al., 2011) and tend to resolve as memory intrusions decline (Guez et al., 2013).

Procedures like the HM paradigm afford tools to conduct controlled examinations into whether frequent, intentional acts of memory suppression have the potential to produce unintentional gaps in memory for unrelated materials. Practically, the strategy-specific nature of the amnesic side effect – predicted by theory and observed in the laboratory – may suggest more sustainable prescriptions for coping with unwanted memories in everyday life (Hulbert et al., 2016). For example, a thought-substitution strategy may achieve the same immediate effects of control (suppression of unwanted thoughts) while avoiding the unintended side effects and building new, more favourable associations.

Moreover, it may be possible to harness the amnestic effects of memory control for beneficial forgetting when directly suppressing unwanted memories proves too intractable or unpleasant. Previous attempts at developing a so-called cognitive vaccine hint at exciting practical applications. For instance, Holmes, James, Coode-Bate, and Deeprose (2009) discovered that playing Tetris in the wake of exposure to traumatic images resulted in fewer flashbacks to the event over the following week, though recognition was seemingly unaffected. They discuss their findings in terms of competition for visual resources; yet, the need to focus on game-related details in order to successfully play Tetris potentially entails the suppression of otherwise distracting internal thoughts and memories. If so, such findings may represent a particular instance of a more general amnestic side effect.

Memory depressors like the TNT paradigm, cognitive interventions to rewrite emotional memories (Schiller et al., 2010) or reduce their intrusions (Holmes, James, Kilford, & Deeprose, 2010), and the growing promise of real-time fMRI neurofeedback (deCharms, 2008) highlight the potential power of non-invasive techniques for effectively exerting control over internal memory representations. Not only might these techniques be employed to aid individuals in achieving their immediate goals, they also open the possibility of inducing reversible hippocampal lesions in otherwise healthy individuals for study (Hulbert et al., 2016; Mullally & O'Mara, 2013). Techniques such as these may supplement existing findings while simultaneously circumventing methodological and ethical issues associated with more invasive research approaches.

Active ingredients

More work is necessary to bridge levels of analysis and identify the precise neural underpinnings of memory control's targeted and concomitant effects. Here we offer some speculations based on the current literature.

An indirect excitatory pathway from the dLPFC to inhibitory neurons within the deep layers of the entorhinal cortex is known to be routed through the ACC (reviewed by Anderson et al., 2016). This pathway could have the potential to deprive the hippocampus of sensory input necessary for pattern completion. With information from the No-Think cues blocked at the gate, hippocampally dependent retrieval, encoding, and BOLD activity would be similarly attenuated during the suppression episode and potentially beyond (Anderson et al., 2016). Suppression-related BOLD activation in the ACC generally accords with this entorhinal gating hypothesis, though modulations of the entorhinal cortex presumably would be expected more often than reported (for an exception, see Levy & Anderson, 2012).

Anderson et al. (2016) put forward another hypothesis based on substantial bidirectional connections between the ACC and the thalamic reuniens nucleus. This circuit has gained recent attention for its major role in the transmission of information between the prefrontal cortex and the hippocampus via the reuniens (Ito, Zhang, Witter, Moser, & Moser, 2015) and in enhancing encoding specificity (Xu & Südhof, 2013). The existence of projections from the reuniens nucleus to GABAergic inhibitory interneurons in CA1 of the hippocampus also allows for Anderson et al.'s (2016) thalamo-HM hypothesis of memory control: rather than gating hippocampal inputs, this mechanism could lead to the direct, global inhibition of the hippocampus and its attendant mnemonic processes. As such, it could account for both the apparent interruption of pattern completion (blocking unwanted intrusions) and the encoding and stabilisation processes necessary to form contextualised memories of bystander events (Hulbert et al., 2016). This speculative account would be bolstered by high-resolution functional neuroimaging specifically focused on detecting activations in the relatively small nucleus reuniens of the thalamus during memory suppression. Additional insights could be gleaned from techniques allowing for the detection of hippocampal GABA in relation to control abilities.

Anderson et al. (2016) go on to suggest that input gating may afford proactive control over mnemonic awareness. Were this process to fail, an involuntary memory intrusion might then initiate reactive control through the reuniens pathway, resulting in the systemic downregulation of the hippocampal memory system. Not only would this explain the marked intrusion-related hippocampal deactivations in the literature (Levy & Anderson, 2012), it may also explain the amnesic shadow observed in the HM paradigm (Hulbert et al., 2016). Prior experience with memory suppression (in the laboratory or in real life) might even lead to a more expedient adoption of a strategy reliant on thalamo-HM. Similarly, interference experienced during the initial 2-Back training task could have led Mullally and O'Mara's (2013) participants to adopt a consistent strategy of systemic modulation prior to the experimental blocks.

Postscription

The case for dynamic memory presented in these pages posits a suppression mode, in which the hippocampal memory system can be downregulated to dampen retrieval (as in the TNT paradigm) or encoding (as arguably is the case in certain forms of the *n*-Back task) of long-term memories. As with a retrieval mode, assuming a mode of memory suppression seems to be associated with sustained neurobehavioural consequences that extend beyond the precipitating circumstances. Our review of the literature from these different domains led us to consider whether attempts to keep the mind tidy and free of mnemonic distractions may result in unintended side effects in the form of a reversible amnesia. Evidence from the HM paradigm (Hulbert et al., 2016) and modified *n*-Back tasks (Mullally & O'Mara, 2013) is consistent with the notion that strategies conducive to the suppression of unwanted memories induce functional modulations in the hippocampal memory system, affecting the formation and stabilisation of strong, contextually bound memories for unrelated events occurring in the temporal surround. We refer to this generalised effect as the amnesic shadow.

Among the factors thought to influence the scale and scope of the amnesic shadow, experience suppressing unwanted memories in the wake of traumatic life events may incidentally contribute to some of the generalised memory deficits reported in the clinical literature. Further examination of the neurocognitive mechanisms that support the amnesic shadow may reveal optimised strategies for controlling memories with fewer side effects. Just as inhibition's potential to limit unwanted memory intrusions may backfire, to the extent that it sustains harmful behaviours or the suppressed material later returns to relevance, direct suppression may represent a risky strategy before or after critical learning periods. As no one is immune from the need to keep mnemonic distractions at bay, continuing this line of investigation may bring us all closer to the realisation of a more flexible – if not entirely spotless – mind.

So, the next time you pass a billboard for an app designed to boost your memory or for a pill to help you forget, recall that memory dynamics are driven not only by external factors but also by internal goals. While popping a pill or practising a particular memory control strategy may bring about neurocognitive changes favourable to one's current goals, those activities may have far-reaching consequences once goals change. Our theoretical perspective prescribes a research trajectory aimed at isolating the factors that influence endogenous HMs and their associated side effects. Increasing our appreciation for – and understanding of – these and other dynamic memory processes may suggest ways to better optimise the memory system and avoid potholes while driving down memory lane.

Disclosure statement

No potential conflict of interest was reported by the authors.

ORCID

Justin C. Hulbert http://orcid.org/0000-0002-2889-0901
Zall Hirschstein http://orcid.org/0000-0003-2830-2476
Clarence A. L. Brontë http://orcid.org/0000-0002-2339-0350
Eleanor Broughton http://orcid.org/0000-0002-0148-2481

References

Aggleton, J. P., & Brown, M. W. (1999). Episodic memory, amnesia, and the hippocampal-anterior thalamic axis. *Behavioral and Brain Sciences*, *22*(3), 425–444. Discussion 444–489.

Amer, T., Campbell, K. L., & Hasher, L. (2016). Cognitive control as a double-edged sword. *Trends in Cognitive Sciences*, *20*(12), 905–915.

Anderson, M. C. (2003). Rethinking interference theory: Executive control and the mechanisms of forgetting. *Journal of Memory and Language*, *49*(4), 415–445.

Anderson, M. C. (2007). Inhibition: Manifestations in long-term memory. In H. L. Roediger, Y. Dudai, & S. M. Fitzpatrick (Eds.), *Science of memory: Concepts* (pp. 295–300). New York: Oxford University Press.

Anderson, M. C., Bunce, J. G., & Barbas, H. (2016). Prefrontal-hippocampal pathways underlying inhibitory control over memory. *Neurobiology of Learning and Memory*, *134*, 145–161.

Anderson, M. C., & Green, C. (2001). Suppressing unwanted memories by executive control. *Nature*, *410*, 131–134.

Anderson, M. C., & Hanslmayr, S. (2014). Neural mechanisms of motivated forgetting. *Trends in Cognitive Science*, *18*(6), 279–292.

Anderson, M. C., & Huddleston, E. (2011). Towards a cognitive and neurobiological model of motivated forgetting. In R. F. Belli (Ed.), *Nebraska symposium on motivation* (pp. 53–120). New York: Springer.

Anderson, M. C., Ochsner, K. N., Kuhl, B., Cooper, J., Robertson, E., Gabrieli, S. W., … Gabrieli, J. D. (2004). Neural systems underlying the suppression of unwanted memories. *Science*, *303*(5655), 232–235.

Aron, A. R. (2007). The neural basis of inhibition in cognitive control. *The Neuroscientist*, *13*(3), 214–228.

Aslan, A., & Bäuml, K. H. (2011). Individual differences in working memory capacity predict retrieval-induced forgetting. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, *37*(1), 264–269.

Axmacher, N., Elger, C. E., & Fell, J. (2009). Working memory-related hippocampal deactivation interferes with long-term memory formation. *Journal of Neuroscience*, *29*(4), 1052–1060.

Axmacher, N., Haupt, S., Cohen, M. X., Elger, C. E., & Fell, J. (2009). Interference of working memory load with long-term memory formation. *European Journal of Neuroscience*, *29*(7), 1501–1513.

Axmacher, N., Mormann, F., Fernández, G., Cohen, M. X., Elger, C. E., & Fell, J. (2007). Sustained neural activity patterns during working memory in the human medial temporal lobe. *Journal of Neuroscience*, *27*(29), 7807–7816.

Baddeley, A. (2003). Working memory: Looking back and looking forward. *Nature Reviews Neuroscience*, *4*, 829–839.

Banich, M. T., Mackiewicz, K. L., Depue, B. E., Whitmer, A. J., Miller, G. A., & Heller, W. (2009). Cognitive control mechanisms, emotion and memory: A neural perspective with implications for psychopathology. *Neuroscience & Biobehavioral Reviews*, *33*(5), 613–630.

Ben-Yakov, A., & Dudai, Y. (2011). Constructing realistic engrams: Poststimulus activity of hippocampus and dorsal striatum predicts subsequent episodic memory. *Journal of Neuroscience*, *31*(24), 9032–9042.

Benoit, R. G., & Anderson, M. C. (2012). Opposing mechanisms support the voluntary forgetting of unwanted memories. *Neuron*, *76*(2), 450–460.

Benoit, R. G., Davies, D. J., & Anderson, M. C. (2016). Reducing future fears by suppressing the brain mechanisms underlying episodic simulation. *Proceedings of the National Academy of Sciences*, *113* (52), E8492–E8501.

Benoit, R. G., Hulbert, J. C., Huddleston, E., & Anderson, M. C. (2015). Adaptive top-down suppression of hippocampal activity and the purging of intrusive memories from consciousness. *Journal of Cognitive Neuroscience*, *27*(1), 96–111.

Bjork, R. A. (1989). Retrieval inhibition as an adaptive mechanism in human memory. In H. L. Roediger & F. I. Craik (Eds.), *Varieties of memory and consciousness: Essays in honour of Endel Tulving* (pp. 309–330). Hillsdale, NJ: Lawrence Erlbaum.

Bjork, E. L., Bjork, R. A., & MacLeod, M. D. (2006). Types and consequences of forgetting: Intended and unintended. In L.-G. Nilsson & O. Nobuo (Eds.), *Memory and society: Psychological perspectives* (pp. 141–165). New York, NY: Psychology Press.

Bomyea, J., & Lang, A. J. (2016). Accounting for intrusive thoughts in PTSD: Contributions of cognitive control and deliberate regulation strategies. *Journal of Affective Disorders*, *192*, 184–190.

Brewin, C. R. (2011). The nature and significance of memory disturbance in posttraumatic stress disorder. *Annual Review of Clinical Psychology*, *7*(1), 203–227.

Brewin, C. R., & Smart, L. (2005). Working memory capacity and suppression of intrusive thoughts. *Journal of Behavior Therapy and Experimental Psychiatry*, *36*(1), 61–68.

Butler, A. J., & James, K. H. (2010). The neural correlates of attempting to suppress negative versus neutral memories. *Cognitive, Affective, & Behavioral Neuroscience*, *10*(2), 182–194.

Catarino, A., Küpper, C. S., Werner-Seidler, A., Dalgleish, T., & Anderson, M. C. (2015). Failing to forget: Inhibitory-control deficits compromise memory suppression in posttraumatic stress disorder. *Psychological Science*, *26*(5), 604–616.

Cohen, N., Pell, L., Edelson, M. G., Ben-Yakov, A., Pine, A., & Dudai, Y. (2015). Peri-encoding predictors of memory encoding and consolidation. *Neuroscience and Biobehavioral Reviews*, *50*, 128–142.

Conway, M. A., & Fthenaki, A. (2003). Disruption of inhibitory control of memory following lesions to the frontal and temporal lobes. *Cortex*, *39*(4–5), 667–689.

Cousijn, H., Rijpkema, M., Qin, S., van Wingen, G. A., & Fernández, G. (2012). Phasic deactivation of the medial temporal lobe enables working memory processing under stress. *NeuroImage*, *59*(2), 1161–1167.

Dalgleish, T., Hauer, B., & Kuyken, W. (2008). The mental regulation of autobiographical recollection in the aftermath of trauma. *Current Directions in Psychological Science*, *17*(4), 259–263.

deCharms, R. C. (2008). Applications of real-time fMRI. *Nature Reviews Neuroscience*, *9*(9), 720–729.

Depue, B. E., Curran, T., & Banich, M. T. (2007). Prefrontal regions orchestrate suppression of emotional memories via a two-phase process. *Science*, *317*(5835), 215–219.

Detre, G. J., Natarajan, A., Gershman, S. J., & Norman, K. A. (2013). Moderate levels of activation lead to forgetting in the think/no-think paradigm. *Neuropsychologia*, *51*(12), 1–88.

Dudai, Y. (2012). The restless engram: Consolidations never end. *Annual Review of Neuroscience*, *35*(1), 227–247.

Duncan, K., Sadanand, A., & Davachi, L. (2012). Memory's penumbra: Episodic memory decisions induce lingering mnemonic biases. *Science*, *337*(6093), 485–487.

Ehlers, A., Mayou, R. A., & Bryant, B. (1998). Psychological predictors of chronic posttraumatic stress disorder after motor vehicle accidents. *Journal of Abnormal Psychology*, *107*(3), 508–519.

Eichenbaum, H., Yonelinas, A. P., & Ranganath, C. (2007). The medial temporal lobe and recognition memory. *Annual Review of Neuroscience*, *30*(1), 123–152.

Fawcett, J. M., Benoit, R. G., Gagnepain, P., Salman, A., Bartholdy, S., Bradley, C., … Anderson, M. C. (2015). The origins of repetitive thought in rumination: Separating cognitive style from deficits in inhibitory control over memory. *Journal of Behavior Therapy and Experimental Psychiatry*, *47*, 1–8.

Gabrieli, J. D. E. (1998). Cognitive neuroscience of human memory. *Annual Review of Psychology*, *49*, 87–115.

Gagnepain, P., Henson, R. N., & Anderson, M. C. (2014). Suppressing unwanted memories reduces their unconscious influence via targeted cortical inhibition. *Proceedings of the National Academy of Sciences*, *111*(13), E1310–E1319.

Guez, J., Cohen, J., Naveh-Benjamin, M., Shiber, A., Yankovsky, Y., Saar, R., & Shalev, H. (2013). Associative memory impairment in acute stress disorder: Characteristics and time course. *Psychiatry Research*, *209*(3), 479–484.

Guez, J., Naveh-Benjamin, M., Yankovsky, Y., Cohen, J., Shiber, A., & Shalev, H. (2011). Traumatic stress is linked to a deficit in associative episodic memory. *Journal of Traumatic Stress*, *24*(3), 260–267.

Hagger, M. S., & Chatzisarantis, N. L. D. (2016). A multilab preregistered replication of the ego-depletion effect. *Perspectives on Psychological Science*, *11*(4), 546–573.

Hagger, M. S., Wood, C., Stiff, C., & Chatzisarantis, N. L. D. (2010). Ego depletion and the strength model of self-control: A meta-analysis. *Psychological Bulletin*, *136*(4), 495–525.

Hanslmayr, S., Leipold, P., & Bäuml, K.-H. (2010). Anticipation boosts forgetting of voluntarily suppressed memories. *Memory*, *18*, 252–257.

Hanslmayr, S., Leipold, P., Pastötter, B., & Bäuml, K.-H. (2009). Anticipatory signatures of voluntary memory suppression. *Journal of Neuroscience*, *29*(9), 2742–2747.

Hertel, P. T. (2007). Impairments in inhibition or cognitive control in psychological disorders. *Applied and Preventive Psychology*, *12*(3), 149–153.

Hertel, P. T., & Gerstle, M. (2003). Depressive deficits in forgetting. *Psychological Science*, *14*(6), 573–578.

Holmes, E. A., James, E. L., Coode-Bate, T., & Deeprose, C. (2009). Can playing the computer game "tetris" reduce the build-up of flashbacks for trauma? A proposal from cognitive science. *PLoS One*, *4* (1), e4153.

Holmes, E. A., James, E. L., Kilford, E. J., & Deeprose, C. (2010). Key steps in developing a cognitive vaccine against traumatic flashbacks: Visuospatial tetris versus verbal pub quiz. *PLoS One*, *5*(11), e13706.

Hulbert, J. C., Henson, R. N., & Anderson, M. C. (2016). Inducing amnesia through systemic suppression. *Nature Communications*, *7*, 1–9.

Hulbert, J. C., & Norman, K. A. (2015). Neural differentiation tracks improved recall of competing memories following interleaved study and retrieval practice. *Cerebral Cortex*, *25*(10), 3994–4008.

Hupbach, A., Hardt, O., Gomez, R., & Nadel, L. (2008). The dynamics of memory: Context-dependent updating. *Learning & Memory*, *15*(8), 574–579.

Ito, H. T., Zhang, S.-J., Witter, M. P., Moser, E. I., & Moser, M.-B. (2015). A prefrontal-thalamo-hippocampal circuit for goal-directed spatial navigation. *Nature*, *522*(7554), 50–55.

Jonides, J., Schumacher, E. H., Smith, E. E., Lauber, E. J., Awh, E., Minoshima, S., & Koeppe, R. A. (1997). Verbal working memory load affects regional brain activation as measured by PET. *Journal of Cognitive Neuroscience*, *9*(4), 462–475.

Joormann, J. (2010). Cognitive inhibition and emotion regulation in depression. *Current Directions in Psychological Science*, *19*(3), 161–166.

Küpper, C. S., Benoit, R. G., Dalgleish, T., & Anderson, M. C. (2014). Direct suppression as a mechanism for controlling unpleasant memories in daily life. *Journal of Experimental Psychology: General*, *143*(4), 1443–1449.

Kuhl, B. A., Dudukovic, N. M., Kahn, I., & Wagner, A. D. (2007). Decreased demands on cognitive control reveal the neural processing benefits of forgetting. *Nature Neuroscience*, *10*, 908–914.

Lee, Y.-S., Lee, H.-M., & Tsai, S.-H. (2007). Effects of post-cue interval on intentional forgetting. *British Journal of Psychology*, *98*, 257–272.

Levy, B. J., & Anderson, M. C. (2008). Individual differences in the suppression of unwanted memories: The executive deficit hypothesis. *Acta Psychologica*, *127*, 623–635.

Levy, B. J., & Anderson, M. C. (2012). Purging of memories from conscious awareness tracked in the human brain. *Journal of Neuroscience*, *32*(47), 16785–16794.

MacLeod, C. M., Dodd, M. D., Sheard, E. D., Wilson, D. E., & Bibi, U. (2003). In opposition to inhibition. In H. R. Brian (Ed.), *Psychology of learning and motivation: Advances in research and theory* (Vol. 43, pp. 163–214). San Diego, CA: Academic Press.

MacLeod, M. D., & Hulbert, J. C. (2011). Sleep, retrieval inhibition, and the resolving power of human memory. In A. S. Benjamin (Ed.), *Successful remembering and successful forgetting: Essays in honor of Robert A. Bjork* (pp. 133–152). New York, NY: Elsevier.

Manning, J. R., Hulbert, J. C., Williams, J., Piloto, L., Sahakyan, L., & Norman, K. A. (2016). A neural signature of contextually mediated intentional forgetting. *Psychonomic Bulletin & Review*, *23*(5), 1534–1542.

Marzi, T., Regina, A., & Righi, S. (2014). Emotions shape memory suppression in trait anxiety. *Frontiers in Psychology*, *4*(1001), 1–10.

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), e50814.

Munakata, Y., Herd, S. A., Chatham, C. H., Depue, B. E., Banich, M. T., & O'Reilly, R. C. (2011). A unified framework for inhibitory control. *Trends in Cognitive Science*, *15*(10), 453–459.

Myers, L. B. (2010). The importance of the repressive coping style: Findings from 30 years of research. *Anxiety, Stress, & Coping*, *23*(1), 3–17.

Nadel, L., & Hardt, O. (2011). Update on memory systems and processes. *Neuropsychopharmacology*, *36*(1), 251–273.

Noreen, S., & de Fockert, J. W. (2017). The role of cognitive load in intentional forgetting using the think/no-think task. *Experimental Psychology*, *64*(1), 14–26.

Nørby, S. (2017). Forgetting and emotion regulation in mental health, anxiety and depression. *Memory*. doi:10.1080/09658211.2017.1346130

Park, H., & Rugg, M. D. (2010). Prestimulus hippocampal activity predicts later recollection. *Hippocampus*, *20*(1), 24–28.

Paz-Alonso, P. M., Bunge, S. A., Anderson, M. C., & Ghetti, S. (2013). Strength of coupling within a mnemonic control network differentiates those who can and cannot suppress memory retrieval. *Journal of Neuroscience*, *33*(11), 5017–5026.

Paz-Alonso, P. M., Ghetti, S., Matlen, B. J., Anderson, M. C., & Bunge, S. A. (2009). Memory suppression is an active process that improves over childhood. *Frontiers in Human Neuroscience*, *3*(24), 1–6.

Raaijmakers, J. G. W., & Jakab, E. K. (2013). Rethinking inhibition theory: On the problematic status of the inhibition theory for forgetting. *Journal of Memory and Language*, *68*(2), 98–122.

Ranganath, C. (2010). A unified framework for the functional organization of the medial temporal lobes and the phenomenology of episodic memory. *Hippocampus*, *20*(11), 1263–1290.

Ranganath, C., & Blumenfeld, R. S. (2005). Doubts about double dissociations between short- and long-term memory. *Trends in Cognitive Sciences*, *9*(8), 374–380.

Redick, T. S., Shipstead, Z., Harrison, T. L., Hicks, K. L., Fried, D. E., Hambrick, D. Z., … Engle, R. W. (2013). No evidence of intelligence improvement after working memory training: A randomized, placebo-controlled study. *Journal of Experimental Psychology: General*, *142*(2), 359–379.

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

Schiller, D., Monfils, M.-H., Raio, C. M., Johnson, D. C., LeDoux, J. E., & Phelps, E. A. (2010). Preventing the return of fear in humans using reconsolidation update mechanisms. *Nature*, *463*, 49–53.

Schilling, C. J., Storm, B. C., & Anderson, M. C. (2014). Examining the costs and benefits of inhibition in memory retrieval. *Cognition*, *133*(2), 358–370.

Schooler, L. J., & Hertwig, R. (2005). How forgetting aids heuristic inference. *Psychological Review*, *112*(3), 610–628.

Stark, C. E. L., & Okado, Y. (2003). Making memories without trying: Medial temporal lobe activity associated with incidental memory formation during recognition. *The Journal of Neuroscience*, *23*(17), 6748–6753.

Stark, C. E. L., & Squire, L. R. (2001). When zero is not zero: The problem of ambiguous baseline conditions in fMRI. *Proceedings of the National Academy of Sciences*, *98*(22), 12760–12766.

Storm, B. C. (2011). The benefit of forgetting in thinking and remembering. *Current Directions in Psychological Science*, *20*(5), 291–295.

Storm, B. C., & Levy, B. J. (2012). A progress report on the inhibitory account of retrieval-induced forgetting. *Memory & Cognition*, *40*(6), 1–17.

Streb, M., Mecklinger, A., Anderson, M. C., Lass-Hennemann, J., & Michael, T. (2016). Memory control ability modulates intrusive memories after analogue trauma. *Journal of Affective Disorders*, *192*, 134–142.

Tomlinson, T. D., Huber, D. E., Rieth, C. A., & Davelaar, E. J. (2009). An interference account of cue-independent forgetting in the no-think paradigm. *Proceedings of the National Academy of Sciences*, *106*(37), 15588–15593.

Wegner, D. M. (2009). How to think, say, or do precisely the worst thing for any occasion. *Science*, *325*(5936), 48–50.

Wu, J. Q., Peters, G. J., Rittner, P., Cleland, T. A., & Smith, D. M. (2014). The hippocampus, medial prefrontal cortex, and selective memory retrieval: Evidence from a rodent model of the retrieval-induced forgetting effect. *Hippocampus*, *24*(9), 1070–1080.

Xu, W., & Südhof, T. C. (2013). A neural circuit for memory specificity and generalization. *Science*, *339*(6125), 1290–1295.