Is external memory memory? Biological memory and extended mind ☆

Kourken Michaelian *

Philosophy Department, Bilkent University, Turkey

ABSTRACT

Clark and Chalmers (1998) claim that an external resource satisfying the following criteria counts as a memory: (1) the agent has constant access to the resource; (2) the information in the resource is directly available; (3) retrieved information is automatically endorsed; (4) information is stored as a consequence of past endorsement. Research on forgetting and metamemory shows that most of these criteria are not satisfied by biological memory, so they are inadequate. More psychologically realistic criteria generate a similar classification of standard putative external memories, but the criteria still do not capture the function of memory. An adequate account of memory function, compatible with its evolution and its roles in prospection and imagination, suggests that external memory performs a function not performed by biological memory systems. External memory is thus not memory. This has implications for: extended mind theorizing, ecological validity of memory research, the causal theory of memory.

© 2012 Elsevier Inc. All rights reserved.

1. Criteria for memory

1.1. Clark and Chalmers' criteria for memory

The core case discussed by Clark and Chalmers in their original argument for the extended mind hypothesis is that of Otto, a fictional Alzheimer's patient:

Otto suffers from Alzheimer's disease, and like many Alzheimer's patients, he relies on information in the environment to help structure his life. Otto caries a notebook around with him everywhere he goes. When he learns new information, he writes it down. When he needs some old information, he looks it up. For Otto, his notebook plays the role usually played by biological memory. Today, Otto hears about the exhibition at the Museum of Modern Art, and decides to go see it. He consults the notebook, which says that the museum is on 53rd Street, so he walks to 53rd Street and goes into the museum.

Clearly, Otto walked to 53rd Street because he wanted to go to the museum and he believed the museum was on 53rd Street. And … it seems reasonable to say that Otto believed the museum was on 53rd Street even before consulting his notebook. [...] The information in the notebook functions just like information concerning an ordinary non-occurrent belief; it just happens that this information lies beyond the skin.

[...] Otto is constantly using his notebook as a matter of course. It is central to his actions in all sorts of contexts, in the way that an ordinary memory is central in an ordinary life. The same information might come up again and again, perhaps

being slightly modified on occasion, before retreating into the recesses of his artificial memory. Clark and Chalmers (1998, pp. 12–13)

My focus here is on the claim that the notebook constitutes a sort of *external memory* for Otto.[1]

My aim is not to argue against the extended mind hypothesis but rather to deal specifically with the concept of external memory: the question is not whether external memory really involves storage of mental contents (even dispositional beliefs) but rather whether external memory is really *memory*. That external memory is memory presumably implies that external memory stores mental contents, which would establish the extended mind hypothesis, but there are alternative routes to this conclusion. According to the parity principle, endorsed by extended mind theorists, what matters is the function of a resource, not its location relative to the agent's body or brain: an external resource counts as cognitive when it performs a function that would be counted as cognitive, were it performed by an internal resource (see Sutton (2010) for discussion). Even if external memory turns out not to be a type of memory, i.e., not to perform the function that is actually performed by biological memory, it clearly performs a function that we would count as cognitive if it were performed by an internal resource. Granted the parity principle, this is enough to establish the extended mind hypothesis.

Nevertheless, while my argument is compatible with the extended mind hypothesis, it does have implications for theorizing in the extended mind framework; I discuss these in Section 3. Additionally, it has implications for the ecological validity of work on memory using standard laboratory paradigms and for the traditional causal theory of memory in philosophy, also discussed in Section 3. Though I do not explore them here, the conclusion that external memory is not memory also has broader implications: what becomes, e.g., of Donald's anthropological account if his contention that external memory "is the *exact* internal analog of internal, or biological memory" (Donald, 1991, p. 309) turns out to be false?

One might worry that the claim that external memory is not memory is paradoxical, but any appearance of paradox here is misleading. My claim is, roughly, that external memory is so unlike memory that it does not belong to the same natural kind, that there cannot be a theory memory covering both biological memory and external memory and assigning them both the same role. This is no more paradoxical than the claim that fool's gold is not gold.

Given that I argue that the conclusion that external memory is not memory does not tell for or against the extended mind hypothesis, my argument should be distinguished from other discussions of differences between internal and external memory, in which such differences are taken to bear directly on the extended mind hypothesis. In such discussions, which focus on the search for a "mark of the cognitive" (Adams & Aizawa, 2001, 2010), it has sometimes been suggested that differences between internal and external memory mean that external memory is not a type of memory, but the suggestion has not in general been supported by a *systematic* discussion of the nature of internal memory. Rupert (2004) comes closest to a systematic approach (and he does explicitly maintain that "external memory" is a different explanatory kind than internal memory), but his focus differs significantly from mine. Rupert is concerned to attack the hypothesis of extended cognition itself (supporting instead the hypothesis of embedded cognition) and targets the claim that external memory is memory primarily in order to undermine the hypothesis. Thus, while he does discuss interference effects and other consequences of construction, his argument does not focus on the broader constructive character of memory; nor does he discuss forgetting or metamemory in any detail. As far as the question whether external memory is memory is concerned, Rupert's argument is largely compatible with mine here, though I do not take either his or my argument to significantly undermine the extended mind hypothesis.

The Otto case is intuitively plausible as a case of external memory. In order to explain why this is so, Clark and Chalmers point out that it satisfies four criteria (Clark & Chalmers, 1998, p. 17):

1. the agent has constant access to the resource;
2. the information in the resource is directly available to him without difficulty;
3. he automatically endorses information retrieved from the resource;
4. information is stored in the resource as a consequence of past endorsement by the agent.[2]

They do not, however, provide any argument for the adequacy of these features as criteria for memory. Obviously, we cannot evaluate the criteria in terms of their adequacy with respect to external memory, since we have little independent purchase on the category. But clearly the criteria should get things right with respect to internal, biological memory, since this category is well-understood (even if there is no uncontroversial characterization of biological memory in general available (Michaelian, 2011))—if the criteria are not satisfied by biological memory systems, they are inadequate.

---

[1] I focus on external memory in the form of artifacts, setting aside cases in which the memory system of one agent apparently serves as an external memory for another agent (a possibility mentioned by Clark and Chalmers); while, in such a case, the resource in question is obviously a memory, it is less obvious whether it is a memory *for* the relevant agent. Neither will I deal with the related phenomenon of group memory, which raises additional complications (since here the coupling goes in multiple directions) (Theiner, 2009, Theiner, Allen, & Goldstone, 2010; Sutton, Harris, Keil, & Barnier, 2010).

[2] They hedge their bets with respect to criterion 4 ("perhaps one can acquire beliefs through subliminal perception, or through memory tampering?" (Clark & Chalmers, 1998, p. 17)), but it fits naturally with their overall picture of memory, and presumably they take ordinary cases of memory not to involve such processes. See also (Clark, 2008, p. 80).
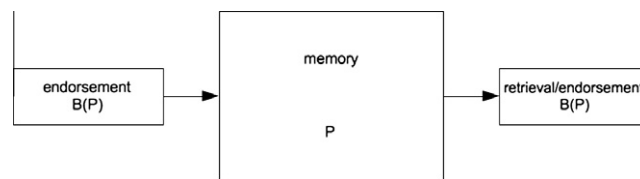
Fig. 1. The simple picture of memory: memory as a container for beliefs.

1.2. Problems for Clark and Chalmers' criteria

While Clark and Chalmers do not say why they take them to be relevant, these features do correspond to a simple, intuitively plausible picture of the operation of memory (Fig. 1). According to this picture, records are placed in memory when the subject endorses or accepts (believes) them. Records are discrete, stable items—they remain unchanged while in memory, and they are retrieved from memory unchanged. And records are endorsed or accepted (believed) upon retrieval, so that we can say that memory stores dispositional beliefs. (Clark and Chalmers are explicit on the point that what is stored in Otto's notebook are dispositional beliefs.) This picture of remembering as a purely preservative process is common among philosophers, though there are exceptions (e.g., Matthen, 2010); Burge (1993), e.g., discusses a purely preservative form of memory, and, while Bernecker's version of the causal theory of memory (Bernecker, 2008, 2010) allows that certain transformations of content are compatible with remembering, it acknowledges only very limited transformations, forbidding transformations which add new content.

Given this simple picture, on which memory is in effect a sort of container for beliefs, it is clearly possible in principle for an external resource such as a notebook to serve as a memory for an agent, to play the same role for him as is played by his internal memory (bracketing certain differences—e.g., speed of encoding and retrieval). The problem is that the simple picture is false on all counts[3]:

Storage is not a consequence of endorsement. Memory does not store only endorsed information – storage is rather determined primarily by a form of relevance, with the consequence that non-believed but relevant records are stored. The specific determinants of encoding are likely to be complex. Depth of processing (Craik, 2002) clearly plays an important role, with the type of processing performed at encoding affecting later memory performance (e.g., semantic processing produces better subsequent recognition of words than does phonological processing). Recent research suggests that survival processing also play a role (Nairne, Thompson, & Pandeirada, 2007; Nairne & Pandeirada, 2008, 2008), with information processed for survival-relevance enjoying better long-term retention. It is likely that additional factors play a role. I use "relevance" here as shorthand for this collection of factors.

Records are not discrete, stable items. Memory records are not discrete, stable items but rather are malleable (susceptible to transformation and recombination) during encoding, the consolidation process that precedes stable storage, and again during the reconsolidation process that follow reactivation (retrieval); nor are records retrieved unchanged, since retrieval is a reconstructive process (Dudai, 2004, 2006; Hassabis & Maguire, 2007, 2009; Koriat, Goldsmith, & Pansky, 2000; Loftus, 2005; Matthen, 2010; McClelland, 1995, 2011; Schacter & Addis, 2007; Schacter, Norman, & Koutstaal, 1998; Shanton, 2011; Sutton, 2003; Vosgerau, 2010).

Endorsement of retrieved records is not automatic. Since records are not stored as a consequence of endorsement by the agent, and due to the constructive character of encoding, consolidation, and retrieval, records are not endorsed automatically upon retrieval—metamemory processes rather intervene to determine endorsement/rejection of retrieved records, e.g., on the basis of source monitoring judgements (judgements about the probable origins of records given their level of detail, relation to other information, etc.) or fluency processing (use of fluency as a cue for accuracy) (Arango-Muñoz, 2011; Dunlosky & Metcalfe, 2008; Hertwig, Herzog, Schooler, & Reimer, 2008; Koriat, 2002; Metcalfe & Dunlosky, 2008; Mitchell & Johnson, 2000, 2009; Nelson & Narens, 1994, 1990; Proust, 2010).

Moreover, *if memory is a container, it is a rather leaky one.* The costs of searching a large memory store and the need for an acceptable level of speed in retrieval mean that forgetting (determined by retrieval history) is a necessity for a functional memory system (Anderson, 1990, 1991, 2000; Bjork, 1989; Bjork & Bjork, 1988; Kraemer & Golding, 1997).

I expand on these points below, but we are already in a position to see that the picture that emerges when we take them into account (see Fig. 2) is considerably more complicated than the simple, purely preservative picture suggested by Clark and Chalmers' criteria. On this more realistic picture, biological memory does not have most of the features that Clark and Chalmers cite as rendering the Otto case plausible as a case in which an external resource serves as a memory for an agent; thus these features cannot be used as criteria for determining when an external resource counts as a memory for an agent. Criterion 2 (direct availability of the information in the resource) is unacceptable: on the realistic picture, though little information is rendered permanently unavailable once stored, much stored information is inaccessible to the agent at any given point in time. Criterion 3 (automatic endorsement upon retrieval) is likewise unacceptable: endorsement of retrieved records is not automatic but rather only one possible outcome of the metacognitive decision process involved in

---

[3] My review of these points here is necessarily brief; I have developed them in detail elsewhere (Craik, 2002; Michaelian, 2011, in press-a, in press-b).
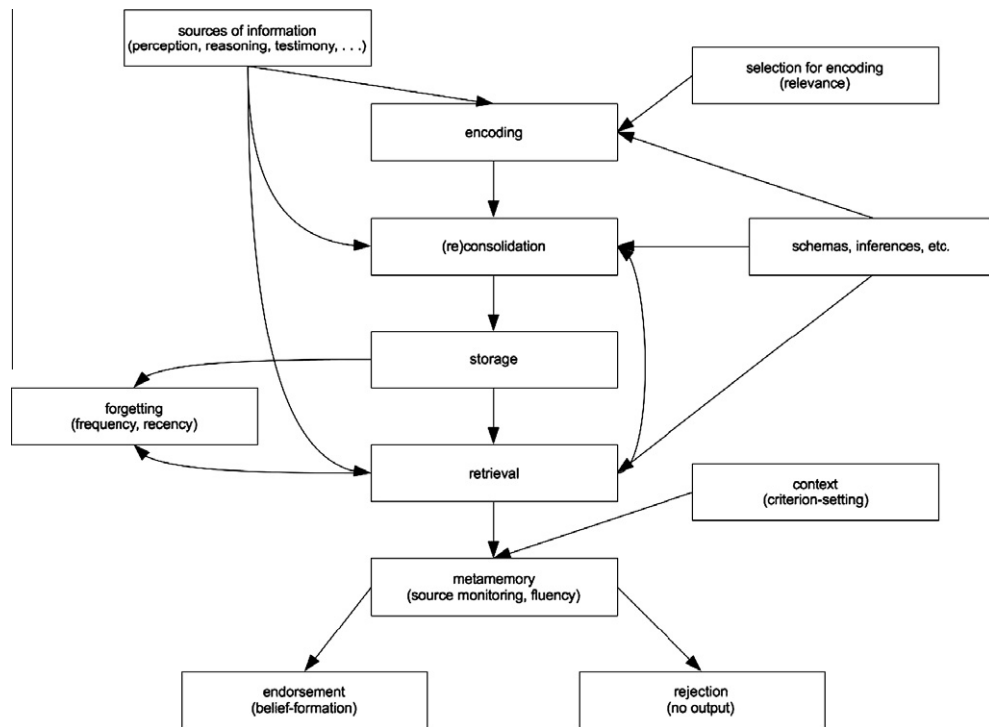
Fig. 2. A more realistic picture of memory: encoding is determined by the relevance of the information to the agent; non-endorsed records are stored. Information is transformed as it is encoded and consolidated, under the influence of schemas, heuristics, and other available information. Stored information can again be transformed under these influences if it is retrieved (reconsolidation). Though little or no information is permanently lost from memory, much information is rendered inaccessible through forgetting (determined at least in part by frequency and recency of retrieval). Records are not automatically endorsed upon retrieval—metamemory judgements (sensitive to both features of retrieved records and the retrieval process, on the one hand, and, on the other hand, to the agent's current context—e.g., the current importance of forming an accurate representation) determine endorsement/ rejection.

retrieval—many representations are rejected once retrieved. Finally, criterion 4 (storage as a consequence of past endorsement) is also unacceptable: storage is not determined by endorsement but rather by an assessment of the relevance or significance of incoming information. Only criterion 1 (constant access to the resource) is compatible with the more realistic picture.

One might defend Clark and Chalmers's argument by suggesting that their focus is ultimately not on memory but rather on belief, so that the failure of biological memory to satisfy their criteria is irrelevant to their argument. However, it is difficult to see what their criteria might be criteria for, if not memory. Moreover, while they do discuss dispositional beliefs, this is in the context of a focus on external resources as stores of dispositional beliefs. Thus my reading of their argument is a natural one.

Alternatively, one might defend Clark and Chalmers' argument by suggesting that their criteria are not meant to provide a general account of memory. Admittedly, they are clearly not attempting to provide anything like a definition or analysis of memory. But it is clear, given the role that the criteria play in their argument, that they take the criteria to refer to important, central features of memory (otherwise they would not be in a position to invoke them to argue that Otto's notebook counts as a memory). Thus they clearly take their criteria to be satisfied by biological memory.

My focus, however, is not on Clark and Chalmers' argument as such. I take the argument as my starting-point because it has provided an important reference point for subsequent debates around extended mind/extended cognition and, primarily, because the view of memory on which it relies is an extremely natural one. But the question of whether external memory is memory is of independent interest, in part due to its implications for the areas discussed in Section 3.

1.3. Modified criteria

Criterion 1 on its own is obviously inadequate, since an agent might easily have constant access to a resource that is clearly not a candidate for an external memory for her, e.g., someone else's notebook. Modifying the remaining criteria in a straightforward way to render them compatible with the more realistic picture of memory will produce something like the following criteria:

1. the agent has constant access to the resource (as before);
2. at least some of the information in the resource is available to him (in appropriate conditions);

3. he determines whether to endorse or reject retrieved information (but trusts at least some of the information stored in the resource);
4. information is stored as a function of its relevance to the agent.[4]

The modified criteria classify the Otto case in the same way as the original criteria: criterion 1 has not changed; the new criterion 2 is weaker than the original criterion 2; Otto determines whether to endorse information retrieved from the notebook (though in fact he choses to endorse all retrieved information); and presumably he does not store information in the notebook unless it is relevant to him.

But while the new criteria are an improvement on the original criteria, they remain unsatisfactory, for, while they take the points about forgetting and endorsement (at encoding and retrieval) into account, they still do not take into account the constructive character of memory. Rather than attempting to approach the question whether external memory is memory by continuing to refine the criteria, I suggest that we should consider directly whether Otto's notebook and similar resources perform the same *function* as biological memory, taking into account the constructive character of biological memory. Supporters of the extended mind hypothesis normally appeal to function when arguing for the hypothesis—according to the parity principle, what matters is only the function of a resource, not its location relative to the agent's body or brain—so they will grant that this move is legitimate. As Clark emphasizes, not any difference between internal and external memory matters, since internal and external memory might differ in terms of implementation—but differences at the level of function do matter (Clark, 2010; Sutton, 2010).

2. The function of memory

Given that the simple picture is adequate with respect to typical putative external memories, and given the deep differences between it and the more realistic picture, it should be suspected that external memory performs a function very different from that of biological memory. Any discussion of the function of memory will inevitably be somewhat speculative, since explicit discussions of the function of memory are rare in psychology (as Boyer points out (Boyer, 2009)). But the standard view in the philosophy of biology, according to which the function of a trait explains why it was selected for, implies that we need to take into account both the actual operation of memory (laws/regularities of its activity) and its evolution (how it could have been selected for, given what it actually does) (Bekoff & Allen, 1995). This is analogous to the situation with respect to the function of reasoning: as Mercier and Sperber argue, rather than assuming that we know what reasoning is for (increasing our stock of true beliefs), we should take the actual operation of reasoning, including apparent failures of rationality, to provide clues to the evolution of the capacity, while general features of the environment in which reasoning evolved can lead to a reassessment of apparent failures of rationality; together, these two constraints can begin to give us insight into the function of reasoning (producing convincing arguments, rather than true beliefs, according to Mercier & Sperber (2011)). In the case of a technology such as external memory, we should, in order to determine the function of the resources, consider both its actual operation and the way in which it was designed.

It is natural to assume that, at a sufficiently general level, we can identify a function (roughly, storage of information for future use) common to all memory systems, including external memory. I have argued elsewhere (Michaelian, 2011) that this is not the case for declarative vs. procedural systems. Here, I argue that it is not the case for internal (declarative) vs. external (declarative) systems. We should of course bear in mind that a given memory system might come to serve more than one function through exaptation (Gould & Vrba, 1982; Sherry & Schacter, 1987); but I will argue that, even taking this into account, simple preservation appears not to be among the functions of biological memory—while biological memory does preserve some information, it does so not as the performance of its function but rather in the service of its constructive function. This is in contrast to external memory, which is designed precisely to preserve information unchanged.

2.1. Against preservationism

While Clark and Chalmers say that Otto's notebook performs the function that is performed by memory in a cognitively healthy person, they do not explicitly specify the function that they take memory to perform. But the intuitive picture of memory suggested by their criteria is associated with a natural "preservationist" view of memory function. On this view (not to be confused with preservationism about memory justification (Lackey, 2005), a view in epistemology according to which memory can preserve but not generate the epistemic justification for a belief), both quantity and accuracy are central to the function of memory: the function of memory is to preserve as many (quantity) of the agents' beliefs (accuracy)[5] as possible.

While natural, preservationism is incorrect as a characterization of memory function: as I show in this section, the points about relevance and endorsement discussed above already pose problems for it, and the point about the constructive character of memory undermines it completely.[6]

---

[4] Strictly speaking, what matters here is predicted or potential relevance: since relevance cannot be determined with certainty ahead of time, selection of records for encoding and, subsequently, for preservation from forgetting, involves processes which in effect attempt to predict future relevance; see Section 2.1.1.

[5] I assume that the agent's other belief-producing processes are largely reliable.

[6] For additional critiques of preservationism and related views, see (Craik, 2002; Matthen, 2010; Michaelian, 2011).

2.1.1. Relevance

A crucial point to take into account when considering the function of memory is that all actual and possible cognizers are in the "finitary predicament" of having limited storage capacity and computing time (Cherniak, 1986). In the case of human memory, storage capacity does not appear to be a serious problem (relatively little information is permanently lost from memory once encoded, and we do not run out of storage space), but this precisely renders the second aspect of the finitary predicament more severe: absent forgetting (which is accomplished by rendering stored records inaccessible, and so can potentially be reversed), the agent would regularly be overwhelmed by records retrieved in response to a cue (Craik, 2002). Anderson's work on adaptive forgetting (as well as related work by Bjork and others (Benjamin, 2011; Bjork, 1989; Bjork & Bjork, 1988)) suggests that forgetting is governed at least in part by retrieval history, which the memory system uses to predict future need for records (Anderson, 1990, 1991; Anderson & Schooler, 2000). The function of memory thus cannot be described as that of preserving as many of the agent's beliefs as possible, since some beliefs must be eliminated—even given that we want to describe the function of the system in terms of preservation, the role of forgetting in the system undermines the quantity aspect of preservationism, since information will be preserved on the basis of its (predicted) relevance.

The role of relevance in determining selection of information for encoding does not undermine this point: selection is imperfect, and, more importantly, relevance changes over time, with changes in the agent's environment (and in the agent herself). The importance of relevance is reinforced by recent work on the influence of survival processing on encoding (Nairne, 2010; Nairne & Pandeirada, 2008; Nairne et al., 2007).[7] The survival processing approach is related to older adaptive approaches to memory, though it differs from them in details. On Anderson's influential approach, memory systems are designed, through forgetting, to reflect the occurrence of information in the environment (Anderson, 1990, 1991; Anderson & Schooler, 2000). Nairne critiques Anderson's approach on the ground that it fails to take into account that our memory systems should be designed to reflect the occurrence of fitness-relevant information, rather than information in general (Nairne, 2010, p. 19). The basic argument here is that "from a fitness perspective … it is more important to remember stimuli related to survival or to other fitness-relevant information such as finding a mate than to remember random occurrences in the environment" (Nairne et al., 2007, p. 269), so that we should expect that "memory systems are "tuned" to remember information that is processed for fitness, perhaps as a result of survival advantages accrued in the past" (Nairne et al., 2007, p. 263). This hypothesis is supported by the finding that survival processing enhances memory (even relative to other forms of meaningful processing) (Nairne, 2010, p. 1). The important point, for present purposes, is that, while the survival processing approach (like Anderson's approach) is compatible with a broadly preservative function for memory, this function is not the preservation of information in general, as preservationism would have it, but only a relevant subset of available information: even if we assume that the function of memory is to preserve the agent's beliefs, only relevant beliefs are selected for encoding, and only relevant beliefs are preserved from forgetting.

2.1.2. Endorsement

Since non-endorsed information can be relevant, we cannot, even given that we want to characterize the function of memory in broadly preservationist terms, say simply that its function is to preserve beliefs or even relevant beliefs. As noted above (Section 1.2), information need not be believed by the subject to be encoded, and stored information is not automatically believed by the subject upon retrieval. Thus records stored in memory do not count as dispositional beliefs (stored contents which tend to be endorsed by the subject when they are retrieved) in virtue of being stored in memory. Some records stored in memory will not count as dispositional beliefs, since the subject does not tend to endorse them upon retrieval. Other stored records will count as dispositional beliefs, but they will do so not in virtue of being stored in memory but rather in terms of the subject's separate disposition (grounded in a metacognitive capacity to evaluate the probable truth of a retrieved record (Michaelian, in press-b)) to endorse them. The memory system, in short, is simply not designed to store dispositional beliefs. This basic point stands regardless of how we eventually unpack relevance: whatever, precisely, determines which information is encoded and then preserved from forgetting, it is clear that non-endorsed information is stored.

Thus, while the role of relevance undermines the quantity aspect of preservationism, the role of endorsement begins to undermine the accuracy aspect of the view—given that agents have access to information from a variety of sources, including many unreliable sources (both internal and external), so that much relevant but inaccurate information will be stored in memory, it begins to look like memory does not function to maintain an accurate representation of past experience (episodic memory) and the world (semantic memory). Even given the assumption that the function of memory is to preserve *something*, it seems, we can say at most that its function is to preserve relevant information.

2.2. Episodic memory

While these two points (about relevance and endorsement) apply to declarative memory in general, caution is required concerning differences between the functions of episodic and semantic memory. As Sherry and Schacter point out (a similar point is made by Klein, Cosmides, Tooby, & Chance (2002)), if we have different memory systems, then we should expect that they serve different functions, for different and incompatible functions give rise to different memory systems (Sherry &

---

[7] Note that Nairne acknowledges that the mental time travel approach (discussed below) is promising and takes it to be compatible with the survival processing approach (Nairne, 2010).

Schacter, 1987, p. 443). Sherry and Schacter focus on the difference between declarative and procedural memory (though they do not use these terms), but the point holds also with respect to semantic and episodic memory.

2.2.1. Mental time travel

Most researchers distinguish episodic memory from the mere "episodic-like" or what-where-when (WWW) memory possessed by some animals (Clayton, Griffiths, Emery, & Dickinson, 2001; Griffiths, Dickinson, & Clayton, 1999; Roberts & Feeney, 2009; Suddendorf, 2003; Suddendorf & Corballis, 2007). In contrast to WWW memory, true episodic memory involves a distinctive phenomenology, autonoetic consciousness—the agent does not merely retrieve a representation of the event, but again has the subjective sense of experiencing it (Markowitsch & Staniloiu, 2010; Nyberg, Kim, Habib, Levine, & Tulving, 2010; Tulving, 1983; Vandekerckhove & Panksepp, 2009).[8] Semantic memory, in contrast, can be defined as memory for facts in general (as opposed to experienced episodes) and does not involve this sort of phenomenology.

An increasingly influential proposal on the general nature of episodic memory is that our capacity to remember the past is intimately bound up with our capacity to imagine the future—that remembering the past and imagining the future are in fact two aspects of a single general capacity for "mental time travel" (MTT), which "allows us not only to go back in time, but also to foresee, plan, and shape virtually any specific future event" (Suddendorf & Corballis, 2007, p. 299). Even a relatively conservative interpretation of the MTT hypothesis already poses problems for preservationism: given the hypothesis, it is a mistake to ask about the function of episodic memory as such; we must instead ask about the function of the system or capacity for MTT, and this will not be only or even primarily a matter of preservation of records of the agent's past experiences.[9]

Moreover, even given that we confine ourselves to MTT and do not abandon it for the more radical imagination hypothesis discussed below, there is reason to prefer a less conservative version of the MTT hypothesis, on which imagination of future events is the primary function of the system or capacity for MTT. As Suddendorf and Corbalis argue, "[t]he fact that episodic memory is fragmentary and fragile suggests that its adaptiveness may derive less from its role as an adequate record of personal history than from providing a "vocabulary" from which to construct planned future events (and perhaps to embellish events of the past)" (Suddendorf & Corballis, 2007, p. 303). This suggests that preserving accurate representations of past experiences is not part of the function of the relevant system, that "our ability to revisit the past may be only a design feature of our ability to conceive of the future" (Suddendorf & Corballis, 2007, p. 303). Given a more radical version of the MTT hypothesis, then, the function of the MTT system is (roughly) construction of representations of relevant possible future events, rather than preservation of representations of relevant past events; to the extent that information about past events is preserved, this is in the service of the production of representations of future events.

2.2.2. Imagination

Related to the MTT hypothesis are versions of the more radical "imagination" hypothesis, as we might refer to it, the claim that episodic memory is an aspect of a capacity which includes not only mental time travel but also other forms of imaginative construction of possible situations.[10] Buckner and Carroll, e.g., point out that there is increasing evidence (see (Nathan Spreng, Mar, & Kim, 2008) for a review) that episodic memory, prospection, theory of mind, and forms of navigation are products of the same brain network, which suggests that "these abilities . . . rely on a common set of processes by which past experiences are used adaptively to imagine perspectives and events beyond those that emerge from the immediate environment" (Buckner & Carroll, 2007, p. 49). They argue that "the adaptive value of episodic memory is not solely in its ability to afford mental reconstruction of the past but rather in its contribution to building mental models—simulations—of what might happen next or other perspectives on the immediate environment", that "[t]he flexibility of the core network might be its adaptive function, rather than the accuracy of the network to represent specific and exact configurations of past events" (Buckner & Carroll, 2007, p. 55).

Buckner and Carroll hypothesize that the common core capacity here is "self-projection", "the ability to shift perspectives from the immediate present to alternative perspectives" (Buckner & Carroll, 2007, p. 49). Memory, in turn, functions to support this capacity for self-projection: "self-projection relies largely on memory systems because past experiences serve as the foundation on which alternative perspectives and conceived futures are built" (Buckner & Carroll, 2007, p. 55). If something like the self-projection hypothesis is right, the function of episodic memory cannot be understood in terms of the preservation of accurate representations of past experience; episodic memory will, rather, function primarily to provide raw materials to be used in various forms of self-projection—in the performance of this function, something is preserved, obviously, but the function of episodic memory is not the preservation of accurate information.

An alternative to self-projection is the scene construction approach defended by Hassabis and Maguire (2007, 2009), who argue that the self-projection approach is unable to account for the commonalities of remembering with imagination of

---

[8] This phenomenology is shared by mental time travel into the future. Boyer (2008) proposes that, from a functional perspective, we can account for the characteristic phenomenology of mental time travel in terms of its ability to counteract our natural tendency to time-discounting, which thus aids long-term planning. If, however, the imagination hypothesis discussed below is right, this explanation does not appear to be available, since the phenomenology will be shared not only by mental time travel into the future but also by at least some other forms of imaginative construction of possible experiences.

[9] Note that Suddendorf and Corbalis also assign a role in prospection (mental time travel into the future) to semantic memory (Suddendorf & Corballis, 2007, p. 300); I return to this point in Section 2.3.

[10] The imagination hypothesis might seem to threaten to blur the distinction between memory and other constructive processes, but, while it does indeed highlight the possibility of intermediate or indeterminate cases, I assume that, in central cases, we can still distinguish between remembering and other constructive processes.

fictitious scenarios. They argue that episodic memory is one of a number of different forms of "scene construction"—"the process of mentally generating and maintaining a complex and coherent scene or event" (Hassabis & Maguire, 2007, p. 299)—and subserved by the same brain network; these include not only episodic future thinking but also navigation, theory of mind, mind wandering (default network), and imagining fictitious experiences.[11] The implications of this approach are similar to those of the self-projection approach: episodic memory, here, is viewed not as a system dedicated to preserving accurate records of the agent's past experience but rather as one activity of a more general "construction system" devoted to producing representations of various relevant possible scenes—the system can be used to produce representations of past experiences, but, even to the extent that it draws on stored information deriving from past experiences, its overall function is not the preservation of accurate such representations (which is not to say that it does not sometimes produce accurate representations).[12]

2.3. Semantic memory

If one of the approaches canvassed in Section 2.2 is right, then episodic memory does not have a preservative function. We have already seen (in Section 2.1) that simple preservationism is untenable with respect to declarative memory as a whole (including semantic memory), but it might be thought that a modified form of preservationism, on which the function of memory is to preserve relevant information, is correct with respect to semantic memory in particular; if so, it might be possible to argue there is enough similarity between the preservative function of external memory and the distinct preservative function of semantic memory to justify classifying the former as a type of memory. The problem with this suggestion is that is there is reason to take semantic memory, too, to have a constructive function: while information is obviously preserved in semantic memory (as in episodic memory), semantic memory does not appear to be designed to preserve accurate copies of (relevant) previously acquired representations.

The basic point about constructive processes made above (Fig. 2) applies to both episodic and semantic memory. As Sutton emphasizes (Sutton, 1998, 2010), the distributed, superpositional character of biological memory storage (brought out clearly by connectionist models of memory—see (McClelland, 1995, 2011)) means that the representations stored in biological memory are, in contrast to the representations stored in typical[13] external memories, highly malleable:

> The storage and organization of information in Otto's notebook is, in Donald's terms [see Donald, 1991], typically exogrammatic. Notably, information is stored there in discrete fashion, and representations in the notebook (linguistic or pictorial representations, for example) have no intrinsic dynamics or activity, are not intrinsically integrated with other stored information, and do no cognitive work in their standing or dispositional form. Representations in Inga's biological memory [Inga is an agent with a normal memory], in contrast, may well blend and interfere: according to connectionist accounts of memory, for example, non-occurrent standing representations, "stored" superpositionally in a single network's weight matrix, influence processing continually in a holistic fashion and are themselves subtly shaped by this ongoing history. Sutton (2010, p. 197)

A system of the sort described by such models is ill-suited for performing a properly preservative function (whether or not restricted by relevance), for the representations it stores are inherently unstable and evolve over time. If connectionist models are approximately right about this aspect of memory storage, it is doubtful that the function of semantic memory can be characterized in (restricted) preservationist terms.

The nature of exograms, in contrast, makes them suitable for precisely the sort of preservative tasks for which engrams are ill-suited:

- Perfectly faithful storage: external memory can be used where the storage of a faithful copy of an item is important (e.g., where it is crucial to avoid intrusions, and where highly similar items need to be kept distinct).
- Unstructured storage: storage where the information is unrelated to existing knowledge.
- "Just in case" storage: external memory can be used for relatively permanent storage of items which might or might not be needed again.
- Storage without forgetting: storage where the information might not be needed again for a long time but should not be forgotten in the meantime.

These features mean that external memory tends to be used for precisely the sort of purely preservative task which biological memory performs poorly. The corollary is that external memory performs the constructive tasks for which biological memory is designed only poorly: in practice, Otto will have difficulty locating desired records in his notebook, will have to deal with many irrelevant or outdated records, and will be confronted with a mass of disconnected, particular facts—while Clark and Chalmers are right to point out that patients with memory deficits often rely on external resources to compensate

---

[11] They acknowledge that episodic memory and episodic future thinking are especially tightly linked, since they have nearly all their component processes in common.

[12] I argue for the compatibility of construction and accuracy in Michaelian (2011, in press-a)

[13] Yeo (2007) discusses a proposed form of external memory designed to mimic features of biological memory, namely, its use of structured storage, and we can in principle imagine forms of external memory that employ some sort of distributed storage.

for those deficits, it is clear that external resources are normally poor substitutes for biological memory (otherwise, memory deficits would not be as devastating to patients as they are).

While I have so far been taking it for granted that the episodic and semantic systems are distinct, there is evidence that episodic and semantic memory in fact exist on a continuum, that they are related functions of a single declarative system. Theoretically, this view is supported by the components of processing approach to memory (Burianova, McIntosh, & Grady, 2010; Moscovitch, 1994; Rajah & McIntosh, 2005; Roediger, Buckner, & McDermott, 1999), e.g., provide brain-imagining evidence for overlap between semantic and episodic memory networks.[14] If an approach according to which there is a single declarative system is correct, semantic memory no more has a purely preservative function than does episodic memory: the function of declarative memory as a whole will be, roughly, the production of relevant representations, drawing on stored information, rather than the simple preservation of stored information.

The multiple memory systems hypothesis remains the majority view, however (Michaelian, 2011; Squire, 2009). Nevertheless, even given the multiple memory systems hypothesis, it is clear that episodic and semantic memory collaborate in various ways. Developing a suggestion in this spirit but focussing on MTT in particular (though he grants that episodic memory might be bound up with additional constructive activities—navigation, unrealistic forms of imagination, and so on), Szpunar argues that semantic memory plays an important role in prospection (MTT into the future). The key observation here is that episodic memory on its own might in a given case not provide all the information required to construct a representation of a future scenario, or that relying only on episodic memory might not be the most efficient way to construct it. The consequence is that:

> whether the information that is used to simulate personal future episodes is episodic or semantic (or even a combination of the two) would seem to depend on the relative accessibility of relevant memory representations. More specifically, abstracted (semantic) representations that are relevant to a given simulation should generally be more accessible than episodic representations of similar information. Episodic representations, on the other hand, might serve an important role in the simulation of future episodes when repeated experiences with specific aspects of a scenario are absent. Szpunar (2010, pp. 156–157)

Szpunar thus argues that it is possible that "specific future episodes may be constructed without the need to necessarily rely on the contents of episodic memory per se" (Szpunar, 2010, p. 157).[15]

The overall picture of the function of declarative memory that emerges here is compatible with the picture given by Fig. 2, as long as we do not conceive of "retrieval" as a matter of the production of a representation designed to duplicate a representation acquired earlier: it is not just that the system effectively produces a new representation which might fail to match an earlier representation but, crucially, that the system does not necessarily attempt to produce a representation which matches an earlier one; in some cases, the system does attempt to match either earlier experience (episodic memory) or knowledge (semantic memory), but in many cases the system is simply engaged in another, non-reproductive activity (prospection, imagination, and so on).

3. Conclusions

We can thus conclude that biological memory and external memory (not all possible external memories, but existing external memories) have different functions.[16] If so, then by the standard employed by extended mind theorists, "external memory" is not a type of memory, after all. This conclusion has implications not only for theorizing in the extended mind framework but also for questions around the ecological validity of experimental studies of memory and for the causal theory of memory in philosophy.

3.1. Extended mind theorizing

While the conclusion that external memory is not memory is compatible with the extended mind hypothesis, it does have implications for theorizing conducted in the extended mind framework. Given the conclusion, we should, rather than looking for cases in which an external resource serves as a replacement for or extension of biological memory, focus on *interactions* between biological memory and external memory, ways in which external memory supplements biological memory or compensates for its inability to perform tasks for which it was not designed. Cases like that of Otto, in particular, are not very realistic (since they assume that an external memory can simply be swapped in to replace a defective internal memory), and they thus risk misleading us; we would do better to focus on more realistic cases, cases representative of the ways in which external memory is actually used by agents (e.g., for tasks where source monitoring is not an issue, or where clutter avoidance (Harman, 1986) is unimportant).

---

[14] (Wirth et al., 2011) explores the involvement of semantic memory in the default network, providing brain-imaging evidence that semantic memory is actively involved in default network activity.

[15] Though he admits that the case of amnesic patients (who have trouble with MTT in general) poses a problem for this hypothesis.

[16] Note that the point about the constructive function of memory is compatible with the adaptivity of true belief (McKay & Dennett, 2009): not all products of construction are believed (simulation of possible outcomes, etc.), and construction can produce new accurate representations.

Something like this reorientation is advocated by Sutton and indeed is already reflected in much theorizing in the extended mind framework. Rather than emphasizing *parity* between inner and outer resources and processes, this extended mind theorizing emphasizes *complementarity* between diverse inner and outter resources: internal and external components of an extended system can "play different roles and have different properties while coupling in collective and complementary contributions to flexible thinking and acting" (Sutton, 2010, p. 194). The argument of this paper reinforces the correctness of this approach.

3.2. Ecological validity

If a focus on interactions between biological and external memory is appropriate in extended mind theorizing, an awareness of the role of such interactions should presumably also inform experimental work on memory. There is reason to be concerned about the ecological validity of experiments that require subjects to use biological memory to perform tasks for which external memory would normally be used. Standard list-learning paradigms, e.g., do not permit subjects to make use of external memories, while, in the wild, an agent faced with the task of remembering a list of items will often resort to recording the list in a notebook. While approaches that deprive subjects of the use of external memory obviously can tell us a great deal about memory, they risk misleading us about the operation of memory.

Experiments using the Deese–Roediger–Mcdermott (DRM) paradigm (DEESE, 1959; Gallo, 2010; Roediger & Mcdermott, 1995), e.g., regularly find that subjects presented with a list of thematically related words ("bed, dream, pyjamas, . . .") later falsely remember a non-presented lure word ("sleep"). While these results provide valuable insight into the principles governing the constructive, inferential operation of biological memory, there is a risk, if we take them in isolation, of being blind to the ways in which agents in the wild succeed in avoiding intrusions through the use of external memory.

3.3. The causal theory of memory

There is a tension between the account of memory's non-preservative function given here and the causal theory of memory in philosophy (first proposed by Martin & Deutscher (1966), this is still the dominant analysis of memory today (Bernecker, 2010)), which assumes that, in order for remembering to occur, the representation produced by retrieval must be highly similar to the relevant earlier representation: given that memory is not designed to preserve representations (more or less) unchanged, many representations produced by retrieval, though accurate, will differ significantly from the earlier representations from which they stem; given a traditional, restrictive similarity requirement, the causal theory will imply that, in such cases, the agent does not remember, regardless of the veridicality of the representation produced by retrieval—it will imply, in other words, that the memory is merely apparent. Though we might bite the bullet and accept this implication, doing so will require us to reclassify many intuitive cases of genuine memory as merely apparent. This tension persists even if the traditional causal theory is replaced with the causal theory of constructive memory developed in Michaelian (2011). The causal theory of constructive memory is designed to be compatible with the psychology of constructive memory. In particular, unlike the traditional causal theory, it allows that a process in which memory retrieval generates new content can count as remembering. But the causal theory of constructive memory still requires "sufficient" similarity between the representation produced by retrieval and the representation encoded earlier. Given the account of the non-preservative function of memory developed here, there will be many cases in which the representation produced by retrieval differs extensively from the relevant earlier representation, thus violating this requirement.

Rather than attempting to resolve this tension, I suggest, we should perhaps take it as a signal that the distinction between genuine and merely apparent memories should be dropped. Whatever criterion (beyond accuracy) we adopt for distinguishing between them, both "merely apparent" and "genuine" memories are produced by the same system, performing the same function, operating according to the same principles. Once we focus on the nature of this system and the principles governing its operation, it is no longer clear whether there is anything to be gained by deeming only some of the representations produced by it to be genuine memories, for the category plays no explanatory role: what matters is whether the system gets things right (producing an accurate representation either of an event experienced by the agent or of some fact), not whether it does so by reproducing a representation entertained earlier by the agent.

References

Adams, Fred, & Aizawa, Ken (2001). The bounds of cognition. *Philosophical Psychology, 14*(1), 43–64.

Adams, F., & Aizawa, K. (2010). *The bounds of cognition*. Wiley-Blackwell.

Anderson, J. R. (1990). *The adaptive character of thought*. Hillsdale: Erlbaum.

Anderson, J. R. (1991). Is human cognition adaptive? *Behavioral and Brain Sciences, 14*, 471–517.

Anderson, J. R., & Schooler, L. J. (2000). The adaptive nature of memory. In E. Tulving & F. I. M. Craik (Eds.), *Handbook of memory* (pp. 557–570). Oxford: Oxford University Press.

Arango-Muñoz, Santiago (2011). Two levels of metacognition. *Philosophia, 39*(1), 71–82.

Bekoff, M., & Allen, C. (1995). Teleology, function, design and the evolution of animal behaviour. *Trends in Ecology & Evolution, 10*(6), 253–255.

Benjamin, A. S. (Ed.). (2011). *Successful remembering and successful forgetting*. Psychology Press.

Bernecker, S. (2008). *The metaphysics of memory*. Springer.

Bernecker, S. (2010). *Memory: A philosophical study*. Oxford: Oxford University Press.

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

Bjork, E. L., & Bjork, R. A. (1988). On the adaptive aspects of retrieval failure in autobiographical memory. In M. M. Grueneberg, P. E. Morris, & R. N. Sykes (Eds.). *Practical aspects of memory: Current research and issues* (Vol. 1, pp. 283–288). New York: Wiley.

Boyer, P. (2008). Evolutionary economics of mental time travel? *Trends in Cognitive Sciences, 12*(6), 219–224.

Boyer, P. (2009). What are memories for? functions of recall in cognition and culture. In P. Boyer & J. V. Wertsch (Eds.), *Memory in mind and culture* (pp. 3–28). Cambridge: Cambridge University Press.

Buckner, R., & Carroll, D. (2007). Self-projection and the brain. *Trends in Cognitive Sciences, 11*(2), 49–57.

Burge, Tyler (1993). Content preservation. *The Philosophical Review, 102*(4), 457–488.

Burianova, Hana, McIntosh, Anthony R., & Grady, Cheryl L. (2010). A common functional brain network for autobiographical, episodic, and semantic memory retrieval. *NeuroImage, 49*(1), 865–874.

Cherniak, C. (1986). *Minimal rationality*. Cambridge: MIT Press.

Clark, A. (2008). *Supersizing the mind*. Oxford UP.

Clark, A. (2010). Memento's revenge: The extended mind, extended. In R. Menary (Ed.), *The extended mind*. Cambridge, MA: MIT.

Clayton, N. S., Griffiths, D. P., Emery, N. J., & Dickinson, A. (2001). Elements of episodiclike memory in animals. *Philosophical Transactions of the Royal Society of London. Series B: Biological Sciences, 356*(1413), 1483–1491.

Clark, Andy, & Chalmers, David (1998). The extended mind. *Analysis, 58*(1), 7–19.

Craik, Fergus I. (2002). Levels of processing: Past, present … and future? *Memory, 10*(5/6), 305–318.

DEESE, J. (1959). On the prediction of occurrence of particular verbal intrusions in immediate recall. *Journal of Experimental Psychology, 58*(1), 17–22.

Donald, M. (1991). *Origins of the modern mind*. Cambridge, MA: Harvard UP.

Dudai, Yadin (2004). The neurobiology of consolidations, or, how stable is the engram? *Annual Review of Psychology, 55*(1), 51–86.

Dudai, Yadin (2006). Reconsolidation: The advantage of being refocused. *Current Opinion in Neurobiology, 16*(2), 174–178.

Dunlosky, J., & Metcalfe, J. (2008). *Metacognition*. Sage.

Gallo, David A. (2010). False memories and fantastic beliefs: 15 years of the DRM illusion. *Memory & Cognition, 38*(7), 833–848.

Gould, Stephen J., & Vrba, Elisabeth S. (1982). Exaptation – A missing term in the science of form. *Paleobiology, 8*(1), 4–15.

Griffiths, Daniel, Dickinson, Anthony, & Clayton, Nicola (1999). Episodic memory: What can animals remember about their past? *Trends in Cognitive Sciences, 3*(2), 74–80.

Harman, G. (1986). *Change in view*. Cambridge: MIT Press.

Hassabis, D., & Maguire, E. (2007). Deconstructing episodic memory with construction. *Trends in Cognitive Sciences, 11*(7), 299–306.

Hassabis, Demis, & Maguire, Eleanor A. (2009). The construction system of the brain. *Philosophical Transactions of the Royal Society B: Biological Sciences, 364*(1521), 1263–1271.

Hertwig, Ralph, Herzog, Stefan M., Schooler, Lael J., & Reimer, Torsten (2008). Fluency heuristic: A model of how the mind exploits a by-product of information retrieval. *Journal of Experimental Psychology: Learning, memory, and cognition, 34*(5), 1191–1206.

Klein, Stanley B., Cosmides, Leda, Tooby, John, & Chance, Sarah (2002). Decisions and the evolution of memory: Multiple systems, multiple functions. *Psychological Review, 109*(2), 306–329.

Koriat, A. (2002). Metacognition research: An interim report. In T. J. Perfect & B. L. Schwartz (Eds.), *Applied metacognition* (pp. 261–286). Cambridge: Cambridge University Press.

Koriat, A., Goldsmith, M., & Pansky, A. (2000). Toward a psychology of memory accuracy. *Annual Review of Psychology, 51*, 481–537.

Kraemer, P. J., & Golding, J. M. (1997). Adaptive forgetting in animals. *Psychonomic Bulletin and Review, 4*, 480–491.

Lackey, Jennifer (2005). Memory as a generative epistemic source. *Philosophy and Phenomenological Research, 70*(3), 636–658.

Loftus, Elizabeth F. (2005). Planting misinformation in the human mind: A 30-year investigation of the malleability of memory. *Learning & Memory (Cold Spring Harbor, N.Y.), 12*(4), 361–366.

Markowitsch, Hans J., & Staniloiu, Angelica (2010). Memory, autonoetic consciousness, and the self. *Consciousness and cognition*.

Martin, C. B., & Deutscher, Max (1966). Remembering. *The Philosophical Review, 75*(2), 161–196.

Matthen, Mohan (2010). Is memory preservation? *Philosophical Studies, 148*(1), 3–14.

McClelland, J. L. (1995). Constructive memory and memory distortions: A parallel-distributed processing approach. In D. L. Schacter (Ed.), *Memory distortion* (pp. 69–90). Cambridge, MA: Harvard.

McClelland, J. L. (2011). Memory as a constructive process: The parallel-distributed processing approach. In S. Nalbantian, P. Matthews, & J. L. McClelland (Eds.), *The memory process: Neuroscientific and humanistic perspectives* (pp. 129–151). Cambridge, MA: MIT Press.

McKay, Ryan T., & Dennett, Daniel C. (2009). The evolution of misbelief. *Behavioral and Brain Sciences, 32*(06), 493–510.

Mercier, Hugo, & Sperber, Dan (2011). Why do humans reason? Arguments for an argumentative theory. *Behavioral and Brain Sciences, 34*(02), 57–74.

Metcalfe, J., & Dunlosky, J. (2008). Metamemory. In H. L. Roediger (Ed.), *Learning and memory: A comprehensive reference* (pp. 349–362). Elsevier.

Michaelian, Kourken (2011). Is memory a natural kind? *Memory Studies, 4*(2), 170–189.

Michaelian, Kourken (2011). Generative memory. *Philosophical Psychology, 24*(3), 323–342.

Michaelian, K. (in press-a). The information effect: Constructive memory, testimony, and epistemic luck. *Synthese*.

Michaelian, K. (in press-b). Metacognition and endorsement. *Mind & Language*.

Mitchell, K. J., & Johnson, M. K. (2000). Source monitoring: Attributing mental experiences. In E. Tulving & F. I. M. Craik (Eds.), *Oxford handbook of memory* (pp. 175–195). Oxford: Oxford University Press.

Mitchell, Karen J., & Johnson, Marcia K. (2009). Source monitoring 15 years later: What have we learned from fMRI about the neural mechanisms of source memory? *Psychological Bulletin, 135*(4), 638–677.

Moscovitch, M. (1994). Memory and working with memory: Evaluation of a component process model and comparisons with other models. In D. L. Schacter & E. Tulving (Eds.), *Memory systems* (pp. 269–310). Cambridge, MA: MIT.

Nairne, James S. (2010). *Adaptive memory evolutionary constraints on remembering* (Vol. 53). Elsevier, pp. 1–32.

Nairne, James S., & Pandeirada, Josefa N. S. (2008). Adaptive memory: Remembering with a stone-age brain. *Current Directions in Psychological Science, 17*(4), 239–243.

Nairne, James S., & Pandeirada, Josefa N. S. (2008). Adaptive memory: Is survival processing special? *Journal of Memory and Language, 59*(3), 377–385.

Nairne, James S., Thompson, Sarah R., & Pandeirada, Josefa N. S. (2007). Adaptive memory: Survival processing enhances retention. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 33*(2), 263–273.

Nathan Spreng, R., Mar, Raymond A., & Kim, Alice S. N. (2008). The common neural basis of autobiographical memory, prospection, navigation, theory of mind, and the default mode: A quantitative meta-analysis. *Journal of Cognitive Neuroscience, 21*(3), 489–510.

Nelson, T. O., & Narens, L. (1990). Metamemory: A theoretical framework and new findings. In G. Bower (Ed.). *The psychology of learning and motivation: Advances in research and theory* (vol. 26, pp. 125–173). New York: Academic Press.

Nelson, T. O., & Narens, L. (1994). Why investigate metacognition? In J. Metcalfe & A. P. Shimamura (Eds.), *Metacognition* (pp. 1–26). Cambridge: MIT Press.

Nyberg, Lars, Kim, Alice S. N., Habib, Reza, Levine, Brian, & Tulving, Endel (2010). Consciousness of subjective time in the brain. *Proceedings of the National Academy of Sciences, 107*(51), 22356–22359.

Proust, J. (2010). Metacognition. *Philosophy Compass, 5*, 989–998.

Rajah, M. N., & McIntosh, A. R. (2005). Overlap in the functional neural systems involved in semantic and episodic memory retrieval. *Journal of Cognitive Neuroscience, 17*(3), 470–482.

Roberts, William A., & Feeney, Miranda C. (2009). The comparative study of mental time travel. *Trends in Cognitive Sciences, 13*(6), 271–277.

Roediger, H. L., Buckner, R. L., & McDermott, K. B. (1999). Components of processing. In J. K. Foster & M. Jelicic (Eds.), *Memory: Systems, process, or function* (pp. 31–65). Oxford: Oxford University Press.

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

Rupert, Robert D. (2004). Challenges to the hypothesis of extended cognition. *The Journal of Philosophy, 101*(8).

Schacter, DanielL., & Addis, DonnaR. (2007). The cognitive neuroscience of constructive memory: Remembering the past and imagining the future. *Philosophical Transactions of the Royal Society B: Biological Sciences, 362*(1481), 773–786.

Schacter, D. L., Norman, K. A., & Koutstaal, W. (1998). The cognitive neuroscience of constructive memory. *Annual Review of Psychology, 49*, 289–318.

Shanton, K. (2011). Memory, knowledge, and epistemic competence. *Review of Philosophy and Psychology, 2*(1), 89–104.

Sherry, David F., & Schacter, Daniel L. (1987). The evolution of multiple memory systems. *Psychological Review, 94*(4), 439–454.

Squire, Larry R. (2009). Memory and brain systems: 19692009. *The Journal of Neuroscience, 29*(41), 12711–12716.

Suddendorf, T. (2003). Mental time travel in animals? *Trends in Cognitive Sciences, 7*(9), 391–396.

Suddendorf, Thomas, & Corballis, Michael C. (2007). The evolution of foresight: What is mental time travel, and is it unique to humans? *Behavioral and Brain Sciences, 30*(03), 299–313.

Sutton, J. (1998). *Philosophy and memory traces*. Cambridge: Cambridge.

Sutton, John (2003). Constructive memory and distributed cognition: Towards an interdisciplinary framework. In Boicho Kokinov & William Hirst (Eds.). *Constructive memory* (pp. 290–303). Sofia: New Bulgarian University.

Sutton, John (2010). Exograms and interdisciplinarity: History, the extended mind, and the civilizing process. In Richard Menary (Ed.), *The extended mind* (pp. 189–225). MIT.

Sutton, John, Harris, Celia, Keil, Paul, & Barnier, Amanda (2010). The psychology of memory, extended cognition, and socially distributed remembering. *Phenomenology and the Cognitive Sciences, 9*(4), 521–560.

Szpunar, Karl K. (2010). Episodic future thought. *Perspectives on Psychological Science, 5*(2), :142–162.

Theiner, G. (2009). Making sense of group memory: The curious case of transactive memory systems. In: W. Christensen, E. Schier, & J. Sutton (Eds.), *ASCS09: Proceedings of the 9th conference of the Australasian society for cognitive science, Sydney* (pp. 334–342). Macquarie Centre for Cognitive Science.

Theiner, Georg, Allen, Colin, & Goldstone, Robert L. (2010). Recognizing group cognition. *Cognitive Systems Research, 11*(4), :378–395.

Tulving, E. (1983). *Elements of episodic memory*. Oxford: Oxford University Press.

Vandekerckhove, Marie, & Panksepp, Jaak (2009). The flow of anoetic to noetic and autonoetic consciousness: A vision of unknowing (anoetic) and knowing (noetic) consciousness in the remembrance of things past and imagined futures. *Consciousness and Cognition, 18*(4), 1018–1028.

Vosgerau, Gottfried (2010). Memory and content? *Consciousness and Cognition, 19*(3), 838–846.

Wirth, Miranka, Jann, Kay, Dierks, Thomas, Federspiel, Andrea, Wiest, Roland, & Horn, Helge (2011). Semantic memory involvement in the default mode network: A functional neuroimaging study using independent component analysis. *NeuroImage, 54*(4), 3057–3066.

Yeo, Richard (2007). Before memex: Robert Hooke, John Locke, and Vannevar Bush on external memory. *Science in Context, 20*(01), 21–47.