Transactive Memory Reconstructed:
Rethinking Wegner's Research Program

BRYCE HUEBNER

Bryce Huebner is an Associate Professor of Philosophy at Georgetown University. His research primarily focuses on issues in in the philosophy of the cognitive, biological, and social sciences. He is the author of *Macrocognition: A theory of distributed minds and collective intentionality* (Oxford University Press, 2014), as well as several articles addressing issues the structure of human agency, the kinds of minds we find in our world, and the implications of cognitive scientific research for ethical theorizing.

Abstract: In this paper, I argue that recent research on episodic memory supports a limited defense of the phenomena that Daniel Wegner has termed transactive memory. Building on psychological and neurological research, targeting both individual and shared memory, I argue that individuals can collaboratively work to construct shared episodic memories. In some cases, this yields memories that are distributed across multiple individuals, instead of being housed in individual brains.

When we first met Otto, he was experiencing early symptoms of Alzheimer's disease (Clark & Chalmers 1998). Although he had lived in the same apartment since the 1970s, he was having a hard time getting around his beloved city. But he could still remember *where to look* for information, even when he couldn't remember *where things were*. This allowed him to use a notebook to get around. When he decided to see an exhibit at the MOMA, he could check his notebook to see where the museum was, and set out with his notebook in hand. David Chalmers (2008) has recently claimed that many people rely on iPhones in a similar way. We adeptly navigate unfamiliar cities, and successfully follow through on our plans with the help of these prosthetic devices. Daniel Wegner (Sparrow et al 2011; Wegner 2012; Wegner & Ward 2013) has recently advanced a similar claim, suggesting that the use of the Internet has greatly expanded our cognitive capacities. I remain skeptical. Yet I believe that a traditional approach to cognition funds a limited, but no less important strategy for defending the existence of socially extended memory. In this paper, I focus on the *collaborative* reconstruction of shared episodic memories; my aim is to show that in some cases networks of interfaced cognitive mechanisms, housed in different brains, function as distributed cognitive systems that can remember past experience.

Let me put all of my cards on the table. My preferred approach to cognition starts from the perspective of individual entities, and I am a committed representationalist. In Section 1 I will offer a rough sketch of what I think an individualist and representationalist should be committed to. In Section 2, I will then turn to an interesting set of data suggesting that memories are sometimes distributed among the member of a group; and in Section 3, I will discuss a more recent, and I believe failed attempt to extend this approach to show that our use of Internet resources also extends cognition. To address the differences between these cases, I then offer a brief sketch of the architecture of individual memory in Section 4; and in the final two sections I argue that there is reason to believe that some groups of people function as distributed memory systems, while there is little reason to believe that our Internet usage extends the human mind.

1. Situated memories

Herbert Simon (1996) famously argued that the patterns that emerge as an ant walks across the beach do so because the ant constantly updates its behavior in light of the patterns in the shifting sands. His deeper insight was that information doesn't always need to be represented in the nervous system to be used; it can often be left in the environment and exploited when doing so becomes necessary, at least so long as we have ways to access and use it. We often structure our informational environment in ways that allow us to minimize the difficulties inherent in navigating our world; and we often minimize the amount of cognition that is necessary in a particular case by learning to exploit the information that is readily available in our material and social world. In many cases, we can rely on perception rather than memory to deal with environmental contingencies. "To take a homely example, it would be silly, for most purposes, to try to keep track of what shelf everything in the refrigerator is currently on; if and when you want something, *just look*" (Haugeland 1998a, 219). There is a broad and expanding consensus that we often exploit the physical and social structure of our world when we expect that the information we need will be available when we look for it again (cf., Bechtel 2009; Clark 2003; Rupert 2010). And there is reason to believe that the people around us frequently function as prosthetic devices, which allow us to fill in gaps in our own abilities by relying on skills possessed by others (Kosslyn 2006).

There are also cases where internal and external resources become so dynamically intermingled that it becomes hard to rule out the possibility that a distributed cognitive system collectively produces some type of goal-directed behavior. Consider Evelyn Tribble's (2011; cf., Tribble & Keene 2011) discussion of the actors in early modern England, who relied on the physical and informational structure of the theaters they worked in to remember their parts. These actors played roles in as many as five different shows per week, and they often learned *at least* one new part each week. To carry out these cognitively demanding tasks, they often relied on skeletally structured scripts that could be fleshed out using cues embedded in the structure of the theater, and the structure of their interactions with other actors. These actors clearly had an amazing ability to exploit the information available in the theater; but more importantly, the ability to remember these parts is *only made intelligible* by treating this cognitive task as the capacity of a socially situated agent, who remembered his parts in the context of theatrical practices that were available in early modern England. These memories depended on dynamic and socially situated patterns of engagement with the structure of the theater, and with other actors. So any plausible explanation of how they were able to carry out this amazing feat must be situated within a broader understanding of the mind-body-world relations that were indicative of this practice.

Nonetheless, there is room for dispute about the import of such cases. They clearly reveal that remembering can go beyond what is encoded by internal mechanisms. But a firmly committed internalist may argue that all of the cognitive work is carried out by internal processes, which are updated dynamically against the information that is present in the theater, and in the interactions with other actors. This is the upshot of Simon's attempt to extend his claim about the ant to provide an understanding of human cognition. We rely on the structure of the supermarkets we frequent in making our shopping decisions; and when the organization of that store changes it becomes harder for us to remember what we intended to buy. We routinely rely on our friends and family members to help us remember what happened to us in the past. But in many cases, this is because we only encoded a skeletal representation of our world, which can be elaborated in light of the physical and social aspects of the material space in which we live.

I call attention to this individualist hypothesis because it helps to make it clear why it is hard to demonstrate that memory extends beyond the boundaries of skin and skull. Such claims require disentangling the *environmental influences* on a system's performance, from the *inputs into* a cognitive system and the informational *interfaces within* that system (De Jaegher, Di Paolo, & Gallagher 2010; Weiskopf 2010). As Dan Weiskopf (2010) argues, the traditional cognitive scientific approach to cognitive systems attempts to uncover computational architectures that can be characterized in terms of: 1) a representational vocabulary; 2) a set of operations carried out over these representations; and, 3) the networks of control structures and resources that determine patterns of activation and inhibition within a system (cf., Kaplan 2012; Huebner 2014; Rupert 2010). From this perspective, cognitive systems are best understood as networks of interconnected mechanisms, which take input from transducers (which witlessly and reflexively map physical inputs onto representational vocabularies), carry out cascades of computational processing, and produce system-level behavior by passing the resulting representations back to *effectors* (which witlessly and reflexively map representations onto physical outputs). Everything in between the transducers and effectors is part of the cognitive system.

On this view of cognitive systems, which I take as my starting point, a specification of the boundaries around a cognitive system requires answering several interrelated questions: Which mechanisms are employed in solving a particular cognitive task; How to these mechanisms behave; How are these mechanisms spatially and temporally organized, in a way that constitutes an integrated network; and How does this network interact with things outside of itself? This approach will reveal the existence of cognitive systems wherever there is an integrated network of mechanisms that facilitates solving a particular cognitive task. And it makes it an empirical question whether a network of spatially distributed mechanisms, that cross the boundaries between skin and skull, and responsible for the production of goal-directed, unified, system-level behavior. Establishing the existence of an such a system is no mean task, and it cannot be accomplished by appeal coarsely articulated claims about what a system does. Indeed, establishing the existence of such systems requires showing that agents are “so intertwined with entities outside [themselves] that the *responsible system* includes one or more cognitive agents and their environment” (Bechtel 2009, 156; *emphasis mine*). The question, then, is whether there are cases where agents are intertwined with others in this way.

2. Transactive memories

After spending a substantial amount of time together, some people begin to treat one another as more than just romantic partners. They start to finish one another's sentences; they start to treat one another (if only implicitly) as cognitive resources; they learn to solve problems together; they collaborate to negotiate difficult and unfamiliar situations; and they work together to reconstruct the shared memories that they care about. These practices often emerge naturally and organically, and in a fairly familiar way. We learn about what our close friends and partners remember, and we build memory networks that allow us to access a wider range of information than we would have had access to on our own. Sometimes this happens slowly; sometimes it happens more rapidly. In many cases, we learn to exploit each other as cognitive tools. But there are many ways to engage with others, ranging along a spectrum anchored by pure exploitation and legitimate collaboration. And in an interesting set of cases, we find transactive memory systems: couples whose thinking is so integrated that they remember past events *as a couple*.

In the 1980s and 1990s, Dan Wegner and his colleagues began to examine cases of transactive memory empirically, and Wegner started to build a theoretical account of the conditions under which transactive memory systems would emerge (Wegner et al 1985; Wegner et al 1991; Wegner & Wegner 1995). In one study, they asked romantic couples and ad-hoc pairs who hadn't met previously to remember a list of items that was divided into several different categories (Wegner et al 1991). Half of the couples and half of the ad-hoc pairs were told which categories each member should focus on; the remainder were given no guidance whatsoever. The romantic couples that were given no guidance remembered the most items (M=31.40), and there was little overlap in the items remembered by each member. Ad-hoc pairs who were told which categories to focus on also did pretty well (M=30.14); ad hoc pairs that were not given guidance remembered even fewer items (M=27.64), and there was more overlap in the items remembered by each person. And most strikingly, when the members of romantic couples were assigned categories, they remembered *significantly fewer items* than every other type of group (M=23.75).

Wegner argues that as people spend time with one another, they come to rely on (often tacit) assumptions about functional specialization and about the distribution of memories. This allows them to distribute cognitive labor by allocating different types of information processing to each partner. When an experimenter assigns new categories, this produces interference in *the transactive memory system,* blocking the encoding of new memories. Roughly, this is like asking someone to memorize a list of words while remembering a set of unfamiliar commands—but the effect occurs at the level of the group, not at the individual level.

Anecdotal support for claims about this sort of functional specialization and this sort of distribution of memory can be derived from the cognitive impact of losing a partner or close friend. When a friend, loved one, or partner stops being part of a functionally specialized transactive memory system, they often seem to leave behind something that looks a lot like an index that references this missing information. Subsequent attempts at recall then tend to produce a feeling-of knowing—yielding something like the 404-error that occurs when a web-page is no longer found on the Internet. These feelings-of-knowing provide phenomenological support for the existence of transactive memory systems; but they also suggest a useful strategy for thinking about the functional architecture of transactive memory systems.

Furthermore, basic principles of hierarchical organization suggest that a computational system that *encodes* information topically, and *stores* it in associative networks, can rely on meta-memories to indicate the location of different kinds of information, and use this structure to facilitate rapid and reliable *retrieval*. But nothing precludes the possibility that meta-memories will designate other systems as the storage location for a particular kind of information. Indeed, networked computational systems often work by placing duplicate directories on multiple machines, each specifying where the information is stored (or how it is distributed in the case of peer-to-peer networks using the BitTorrent protocol). Each computer can then rely on a virtual memory that spans the entire network, yielding an increase in processing speed, and a simultaneous decrease on memory load for each machine. And it yields these advantages without a corresponding decrease in the number of tasks that are executable by the network as a whole. With this picture in mind, Wegner (1995) contends that it is a trivial consequence of a computational approach to memory that social groups can function as computational networks.

In stable long-term partnerships, strategies for allocating information can develop that are sensitive to the idiosyncratic capacities of individuals; while in less stable partnerships, specialization can be explicitly negotiated, or can emerge implicitly as a result of biased assumptions or perceptions of genuine expertise. This produces a functionally differentiated structure, wherein information is distributed throughout the system (often with some types of information encoded redundantly). In a functionally integrated system of this sort, an individual can rely on a virtual memory that spans the entire network. Wegner argues that when we retrieve a memory, we reflexively check to see whether it is stored endogenously or exogenously; if it is stored endogenously, we access it; but if it is stored exogenously, we execute a verbal query to retrieve that information from another person. Finally, he claims that these systems can be updated by adjusting the structure of meta-memories, or by adjusting the distribution of information. Since these processes do not always track one another, inconsistent updating strategies can yield feelings-of-knowing such as those I discussed above.

A functional architecture like this could allow exogenously referring meta-memories to play a role in an integrated network, implemented by generating informational interfaces between individuals. In such systems, a group may remember things that an individual cannot remember on their own; memories that are relevant to group projects might also be distributed in ways that allow the members of the group to draw on shared representational resources. While this view of memory critically depends on individual representations, it suggests that information may be encoded and stored in a virtual memory, which allows processing to be distributed across multiple individuals. This yields a computational sketch of an extended cognitive system. But it is important to remember that Wegner has only offered a rough characterization of how transactive memory systems work, omitting many of the structural aspects of a mechanistic explanation. If these aspects could be filled in, his analysis could be transformed to yield a mechanistic explanation of transactive memory systems (Piccinini & Craver 2011).[1] and my aim in Section 4 will be to return to this possibility. But first, I think that it is worth thinking carefully about what more this account entails, and about the kind of computational architecture that it depends upon. And in doing so, I believe that it will

1 Many philosophers and scientists working on extended or situated cognition have argued that we should adopt a mechanistic approach to cognition. Wilson (1994), Menary (2007, 2010), Hutchins (1995), and Kirchhoff (2013) have developed some of these suggestions, and Theiner (2013) has recently provided a "mechanism sketch" for different kinds of transactive memory systems. There are many points upon which I would agree with these authors, and there is a great deal of convergence between Theiner's project and my own. But I approach these questions from a different direction. I hope to explain why representational systems that we must already posit to explain individual memory provide a plausible foundation for understanding transactive memory. For the purposes of this paper, I remain non-committal about the more exciting claims advanced by Theiner (though see Huebner 2014 for an initial statement of my position on these issues)

help to consider Wegner's more recent attempt to extend his account of transactive memory systems to our use of the Internet.[2]

3. Googling transactive memory

Many of us *exploit* the Internet as an epistemic resource: we use smartphones to get around unfamiliar cities; we quiet interpersonal disputes with Google searches; and, we spend hours scouring the Web to find hip and exciting underground restaurants. In each case, the Internet serves as a tool that we can use to navigate our high-tech world. It is also relatively clear that Internet use has modified our inferential landscape, allowing us to exploit a novel range of “context-specific correlations to simplify the problem solving process” (Rupert 2010, 180). But in a paper that builds upon the transactive memory framework, Betsy Sparrow, Jenny Liu, and Dan Wegner (2011) argue that the Internet has also come to function as an external memory for many people. They maintain that the Internet sometimes serves as *part* of a distributed cognitive system.

Sparrow and her colleagues used a priming study to show that people reflexively think about computers when presented with hard questions. On this basis, they argue that people store meta-memories designating computers as memory storage locations. In another study, they found that people are less likely to encode information if they think it will be available later using a search engine or database. They take this to be good evidence of functional specialization and resource allocation that integrates computers into transactive memory systems. Noting that Google is becoming ubiquitous in many of our lives, Sparrow and her colleagues (2011, 778) contend that we no longer need to remember much of anything, since we “are becoming symbiotic with our computer tools, growing into interconnected systems that remember less by knowing information than by knowing where the information can be found”. Consequently, we only need to remember things we that won't have access to, otherwise we can simply remember *where to look*.

In a popular treatment of these and other data, Dan Wegner and Adrian Ward (2013) argue that the “Internet has become the external hard drive for our memories”. This suggests that memories are stored on the Internet, and later retrieved for use. And in a recent opinion piece in the New York Times, Wegner (2012) suggests, in unambiguous terms, that this fact has an exciting implication:

> “Groups of people commonly depend on one another for memory in this way — not by all knowing the same thing, but by specializing. And now we’ve added our computing devices to the network, depending for memory not just on people but also on a cloud of linked people and specialized information-filled devices.
>
> We have all become a great cybermind. As long as we are connected to our machines through talk and keystrokes, we can all be part of the biggest, smartest mind ever.”

---

[2] I have relied on Wegner’s discussion of transactive memory systems for expository purposes in this paper, as they readily lend themselves to a discussion of the sort of computational architecture that could be used to implement a form of extended and distributed memory. There are other approaches to the phenomena of transactive memory systems, which rely more heavily on behavioral phenomena; and there are many cases in which working collaboratively generates a form of collaborative inhibition, rather than leading to more sophisticated cognitive capacities in groups. For an overview of the scientific data on transactive memory, see Ren & Argote (2011). For a philosophical treatment of the complexities involved in the production of collective memories, see Sutton et al (2010) and Theiner (2013).

To be fair, this is not a scientific paper, it is an editorial written for presentation to a popular audience. So the theoretical support for these interesting claims leaves much to be desired. Yet, these claims do suggest a philosophically interesting question: Do we use Internet resources as tools, or are we really becoming part of distributed cognitive systems, which include both human brains and Internet resources?

One easy way of defending the more radical claim would be to build on the suggestion that the "Internet has become the external hard drive for our memories", and suggest that memories are really just packets of information, which can be stored almost anywhere and retrieved using many different kinds of search strategies. If this were a plausible story, it would be clear how web-based information and web-searches could be treated as analogues of stored memories and retrieval strategies. And this would make it easy to treat the Internet as part of various transactive memory systems. Approaching things from this perspective would also help to explain why our web-use yields a system that looks like the transactive memory systems that Wegner discussed in the 80s and 90s. But this brings the problem with offering a coarse-grained sketch of transactive memory systems into stark relief.

Why should data like these establish the existence of a system that consists of a person and their web-searches, rather than supporting the more banal claim that we use the Internet as a tool for navigating our modern technologically rich environment? Put differently, why should the kinds of information retrieval strategies that occur in collaborative recall, web use, and individual memory all be seen as members of a single cognitive kind, remembering (Rupert 2004)? Suppose we decide to treat Chalmers-and-his-iPhone as a cognitive system. As Wegner and his colleagues suggest, he may grab his iPhone when he is faced with difficult empirical questions, and he may be less likely to encode information that he thinks will be available using a Google search or an iPhone application (Chalmers 2008). In carrying out his Google searches, he "ultimately may manipulate some data on a server stored in Helsinki", and it would take a lot more argument than we have seen thus far to establish that the processes carried out on these servers are part of a distributed *cognitive system* (Weiskopf 2010, 319). Yet, it would be a mistake to rule out the claim that such data manipulations are parts of extended cognitive systems on a priori grounds. This is why we must think more carefully about how individual memory works, and about what kinds of interfaces with external resources are possible.

4. Constructing memories

The rhetoric of Wegner's proposal, sits comfortably with the view that people 1) store representations of important events and things, 2) organize these representations associatively, and 3) remember them by retrieving these stored representations. This view also resonates with the story of Otto and his notebook, as it occurred in the earliest discussions of the extended mind. After all, Otto stores his memories in his notebook, and then retrieves them when doing so is necessary. As Sarah Robins (2012) notes, this 'archival' approach to memory also sits comfortably with an everyday understanding about how digital computers work.[3] They use distinct systems for central processing and data storage; and their central processors execute logical operations to manipulate stored representations, or to access representations that can be retrieved from external sources. As I noted about, this view would make it clear why manipulating linguistic and iconic representations on the Internet counts as remembering, because meta-memories facilitate the retrieval of representations stored on the web.[4] But as Wegner knew all too well, this is not a plausible account of individual memory.

As I noted in Section 1, there is theoretical reason to suppose that we often store skeletal representations that can be fleshed out when doing so is necessary. And nowhere is this insight more important than it is in the case of memory. Many of the things we encounter are irrelevant to our ongoing cognitive activity; others are too common to bother remembering; and many things change too frequently for it to be economical to maintain an up-to-date representation of the world (Haugeland 1998b). From a biologically point of view, there is little value to storing a robust memory, which can be literally recalled, where the most important features of the environment are likely to change before we need or want to examine them again (Bartlett 1932, 204). These considerations have yielded a rapidly expanding consensus that we often encode only enough information to specify the *gist* of a situation, rather than encoding every detail of an important situation, event, or practice in long-term memory (Bartlett 1932; Neisser 1967, 1981; Schacter & Addis 2007; cf., Sparrow et al 2011).[5] When the details of this argument are fleshed out, they cut against an archival approach to memory, suggesting that memory is “far more decisively an affair of construction rather than one of mere reproduction” (Bartlett 1932, 204). This fact is significant, because it provides a more promising strategy for grounding a mechanistic account of transactive memory systems (Or so I shall argue in Section 5).

Consider three interesting empirical phenomena that have been demonstrated in the case of individual memory:

- Fredric Bartlett (1932) asked students at Cambridge to read a Native American folktale. He found that many students recalled details of the story many years later, though they replaced the surreal aspects of the story with culturally grounded assumptions about what *should have happened*.

---

[3] For discussion of the distinction between archival and constructive views of memory, see Sutton (1998) and Robins (2012; submitted). The distinction between archival and constructive memories is also present, though less clearly articulated, in Bartlett (1932).

[4] To be clear, the more recent literature on transactive memory has expanded on other facets of Wegner’s research (cf., Ren & Argote 2011; Theiner 2013). Much of this work focuses on group-level phenomena as opposed to episodic and narrative memories, and this work has included a variety of strategies for distributing expertise throughout a team. My focus is on more familiar types of individual memory. It is a complicated issue to determine whether the group-level properties Theiner (2013) addresses belong to the same cognitive kinds as individual memories, and addressing this issue would take me too far astray from the main line of investigation in this paper. So, for now, I focus on ways to get individuals to constitute a transactive memory system, using only the kinds of cognitive systems that must be posited to explain ordinary individual memories.

5 There is little consensus about the nature or structure of the skeletal representations that memory requires—though Nadel & Moscovitch (1997) offer one promising account of a mechanism in the parahippocampal cortex that might carry out this task. For a review of the philosophical, psychological, and neuroscientific positions that been taken on the nature of memory traces, see de Brigard (2014).

- Ulrich Neisser (1981) analyzed John Dean's *Watergate* testimony, and found that Dean did not have the rich and vivid memories he claimed to have. Yet, he routinely provided a plausible report of what *should have happened*. Sometimes he captured the gist of a situation, but even where he didn't his memories were consistent with the facts at a deeper level: “he recalled the theme of a whole series of conversations, and expressed it in different events” (Neisser 1981, 22).
- Elizabeth Loftus' research on eyewitness testimony suggests that the color of a vehicle, the presence of a stop sign or a yield sign, and judgments about the speed at which a vehicle was traveling are all affected by gentle presentations of misinformation (e.g., Loftus 1975; Loftus & Palmer 1974). And Loftus and Pickrell (1995) found that people will falsely remember being lost in a supermarket as children after reading a set of series of stories drawn from their own lives along side a made-up story about being lost in a supermarket.

I contend that there are three accompanying insights that can be drawn from the results of this research. First, memory requires two different processes: one that encodes and stores details of our experience; and one that exploits culturally specific schemata and stereotypes to flesh out a more complete narrative (Bartlet 1932). Second, repeated and rehearsed events play an important role in structuring the things that we remember (Neisser 1981).[6] And, perhaps most importantly, people incorporate new information into their memories of past experiences, presumably because their memories are being constructed rather than simply recalled. These phenomena together suggest—though they do not establish—that memory is a re-constructive process, which relies on general knowledge of the world to flesh-out the details of skeletal representations.[7] As Neisser (1967, 285) famously put the point, memory is more like archeology than it is like archiving: “Out of a few stored bone chips, we remember a dinosaur.”

Research in neuroscience and neuropsychology provides support for the hypothesis that memory is largely a matter of reconstruction, offering insight into the mechanisms that facilitate episodic remembering (Buckner & Carroll 2007). Thinking about the past,

6 In contrast to the increase in semantic knowledge that comes through repeated exposure to similar situations, repeated exposure to similar events can sometimes yield more confusion, and can produce intractability for episodic memories. To take a familiar example, it's easy to find your car the first time you park it in an unfamiliar parking lot, but it is harder—though clearly not impossible—to do so after you have parked there many times (Gluck, Mercado, & Myers 2010, 86).

7 Another common type of data derives from the DRM paradigm for studying false memories (Deese, 1959; Roediger & McDermott, 1995). Participants are presented with a list of related items, and then asked if they remember items from this list, unrelated items, and related lures that were not on the list. The accuracy for items on the list obey a serial position curve with words at the beginning and the end remembered 70-80% of the time and items in the middle being recognized ~55% of the time. Using lists of 15 semantically related words, Roediger & McDermott (1995) found that participants rarely reported hearing unrelated words, but reported hearing the critical lures ~55% of the time! Their participants expressed a high degree of confidence in their memories, and claimed that the clearly remembered hearing the word. Subsequent studies have yielded participants who 'remember' where the word occurred on the list, and who are able to report the tone of voice in which they 'heard' the word. Moreover, similar results have been discovered for words that are related phonologically and orthographically, and subsequent studies have uncovered similar effects for non-linguistic stimuli like pictures and faces. These types of data are systematic and robust. But, what marks them as interesting is the fact that these memories aren't just mistakes. They are mistakes that can only be explained by assuming that participants “remember at least some of the semantic information conveyed by the previous list” (Robins submitted). For a fuller discussion of these data, see the insightful reading of this paradigm by Robins (2012; submitted)

imagining the future, and considering counterfactual possibilities make use of a common cortical network, which includes the hippocampus, the posterior cingulate, the inferior parietal lobe, the medial prefrontal cortex, and the lateral temporal cortex (Addis et al., 2007; Addis & Schacter, 2008; Buckner & Carroll 2007; Schacter et al., 2007). This *core network* plays a critical role in many tasks that require self-projection and ‘mental time travel’, and there reason to believe that neural systems can be redeployed in the service of multiple different tasks (Anderson, 2007; 2010; de Brigard 2013). But if a system is involved in remembering the past, imagining the future, and counterfactual thinking, it must be a system that is sensitive to the fact that the future is never exactly like the past or present. It must be a system that can “draw on the elements and gist of the past, and extract, recombine and reassemble them into imaginary events that never occurred in that exact form” (Schacter & Addis 2007, 27).

The emerging consensus is that episodic remembering depends on mechanisms for encoding and storing skeletal representations, as well as mechanisms for constructing counterfactual representations to flesh out these representations when doing so is necessary. By integrating skeletal representations of events with general representations of previously experienced and encountered situations, we can construct simplified mental models, which allow us to draw inferences about things that are not explicitly encoded in the structure of the stored representations (cf., Fauconnier & Turner 2003). There is little doubt that we store skeletal representations, which accurately depict some aspects of how our world is or how it was. But it is no less important that we rely on a system for constructing counterfactual representations to flesh out these skeletal representations.

Importantly, our ability to construct representations of counterfactual situations can draw on a wide range of biologically and socially significant information, including stored schemata, information that is available in our material and social environment, things we can deduce from cultural conventions, and tacit assumptions about the similarities and differences between our current and previous situation.[8] But importantly, we exploit a heterogeneous array of information in producing and structuring representations of counterfactual situations. Some of the information we use is stored in long-term memory, in the form of generic schemata that are entrenched as a result of frequently reuse and frequent rehearsal (Fauconnier & Turner 2003, 103). These schemata lie somewhere between episodic and semantic memories, representing events that have occurred repeatedly or common features of events (Neisser 1981). But regardless of where this information is located, the constructive work that we carry out in remembering the past runs by way of mechanisms that are dedicated to counter-factual modeling. And sometimes features of the material and social world in which we live can function as material anchors, which ground and constrain the conceptual models we build of possible situations (Hutchins 2005).

A great deal more empirical work would be necessary to provide a fully mechanistic account of remembering, and much of this work is ongoing in psychology, neuroscience, and neuropsychology. But assuming that this account of memory is roughly right, I contend that a constructive theory of memory provides a more plausible foundation for transactive memory research of the sort that Wegner was interested in.

[8] Developmental data suggests that this capacity emerges early. Five-year-old children have a clear sense of the features that must be shared and that can differ across possible worlds, as well as a sense of the kinds of things that are possible in different fictional worlds (Skolnick & Bloom 2006a, 2006b).

5. Transactive memory reconstructed

As Robins (2012, in prep) argues, a theoretically and empirically viable theory of memory must distinguish between the representations we store as memory traces and the representations we produce during the process of remembering. We store sparsely encoded representations of some significant experiences; but remembering is always an inferential process, which can rely on information from a variety of different sources. I contend that even if sparsely encoded memory traces are always stored in individual brains, it will not follow that the representations required for episodic memory must also be stored in individual brains. Put differently, I suggest that the representations we store as memory traces are housed in individual brains, while the representations that we produce during the process of collective remembering are often realized by shared inferential processes that cross the boundary between brain, body, and world. Indeed, I believe Wegner is right to claim that romantic partners and long-time friends sometimes become cognitively interdependent. Where this happens, the process of counter-factual elaboration required for the production of episodic memories sometimes becomes distributed between them. This yields an integrated memory system that can flexibly adapt to the demands placed on the system from outside, as well as the demands placed on the system by the members of such partnerships.

To get a sense of how this argument goes, let's return to the case of the couple who have spent many happy years together, who are attempting to remember something that happened *to them* long ago. Many of us have seen such couples cue, re-cue, and acknowledge one another's claims as they attempt to construct a plausible representation of what happened to them. Often without reflecting on how this process occurs, they collaboratively build upon the sparse representations that each has stored (or that each has constructed as a result of multiple repetitions of similar events); over the course of a conversation, they collaboratively build a plausible narrative that both agree to as a representation of a shared past event. Through this process of cuing, re-cuing, and narrative reconstruction, each partner adjusts and recalibrates, and they work together to produce a shared memory of their first date or their first conversation.

Fortunately, we need not be content with anecdotal support for this collaborative and constructive process of remembering. Celia Harris and her colleagues (2011, 291) have found that couples construct more detailed and elaborate representations of past events when they rely on cues and shifts between speakers in a dynamic and interactive manner. They have also found that collaborative reconstructions sometimes allow people to recall details that neither remembered on their own. In these cases, the process of cuing and re-cuing is reflexive and natural—and where it occurs, couples exhibit "more detailed, episodic, emotionally richer recall" (Harris et al 2011, 292).

I argue that a couple should be seen as a *transactive memory system* when the processes required for episodic remembering are distributed between the members of a couple in this way. Significant computations can occur in a linguistic medium, at the interface between individuals. This is made possible by the fact that our ability to model counterfactual situations relies on domain-general computations, which are a constitutive part of our ability to construct narrative representations of the past. Because these capacities exploit linguistically encoded representations, transactive memory systems can emerge when members use their ability to think out loud, as well as their abilities to use gestural transactions, to externalize the reconstructive processes required for producing episodic memories. Transactive memory systems arise because cognitive processing occurs at that interface between distinct computational systems. Of course, this doesn't prevent the individuals from constructing their own representation of the target situation as well; but appealing only to internal representations will not explain the behavior of the couple as a whole. The *transactions* play a critical role in the production of these memories. To see why, it will help to consider how this process breaks down.

Consider the functional distribution of information that Wegner treats as a necessary condition on transactive memory systems. As it turns out, functional specialization on its own is neither beneficial nor detrimental to recall (Harris et al 2011, 289). Where couples rely on compatible forms of expertise to facilitate cuing, integrating individual representations in ways that yield a more complete story, the distribution of information processing leads to more robust and more detailed memories. But where couples develop incompatible forms of expertise, which preclude collaborative recall, one person may speak as though they are the authority (Harris et al 2011, 287). These kinds of functional specialization lead to monologues and the constructions of a single person's memory, not to the collaborative construction of a shared representation of a past event. Many forms of correction (e.g., "No. You're remembering our third date, not our first") also lead to truncated or aborted attempts at narrative reconstruction. When one partner repeatedly corrects the other, couples "fail to provide a joint narrative for the event that they had been asked to describe" (Harris et al 2011, 289). Similar inhibitory effects arise when one member of the couple uses a weak acknowledgement of her partner's claims by continually saying 'yes' or 'uh-huh' as fillers. Such acknowledgments can easily be seen as tacit recognitions of one person's authority, or as dismissals of the cues that are being presented. In either case, this will block the kind of collaborative engagement that is necessary if a couple is to construct a shared narrative.

Finally, it turns out that some cues are likely to block shared recall for architectural reasons, rather than as a result of the dismissive or exploitative strategies that are adopted by the members of a couple. Recent research on socially shared retrieval-induced forgetting suggests that where a speaker recounts a memory, her partner may be unable to recover memories of related by unmentioned material (Cuc, Koppel, & Hirst 2007; Coman, Manier, & Hirst 2009). The problem is that the production of cues is blocked by the presentation of related memory, a familiar effect in individual recall (Anderson, Bjork, & Bjork 1994; Anderson & Levy 2002). In collaborative and conversational settings, the presentation of a related cue may block the retrieval of the information that is required for the continuation of a narrative. Unlike the other difficulties, this is an emergent effect on the transactive memory system rather than an organization of strategies that prevents the development of such a system.

6. Wegner reconstructed?

I believe these facts provide a foundation for thinking about how transactive memory systems arise in cases where people rely on collaborative transactions to distribute the process of memory construction.[9] Skeletal representations are retrieved, and through a

9 I would like to defend the following general claim about distributed cognition: collaboration is the mark of distributed cognition, while exploitation is the mark of a tool. The problem, however, is that there are many different ways in which human interact with the material and social aspects of their world, and the concepts at play in discussions of distributed cognition are all "accordion words which, by their expansion and contraction, generate so much philosophical music" (Sellars 1965, 158). There may be significant and

process of cuing and re-cuing, some couples and perhaps some small groups construct shared memories of significant shared episodes. But there are also many cases where individuals only rely on one others as informational resources, or as forms of social scaffolding. This occurs in a way that parallels our reliance on the material world as a source of information that can cue recall. Famously, Dean relied on newspaper clippings to trigger his memories about Nixon, but the newspapers were not part of a larger cognitive system. They served as input to the process of counterfactual construction, which was carried out by a network of interfaced mechanisms that are all internal to Dean's brain. Similarly, there is little reason to treat a person and their partner as a distributed cognitive system if they only use their partner to trigger the construction of memories. Where this occurs, the partner is not part of a larger cognitive system because they only serve as input to the process of counterfactual construction, which was carried out by a network of interfaced mechanisms that are all internal to one person's brain. With these facts in mind, we can now return to Wegner's claims about transactive memory.

Wegner and his colleagues (1991) found evidence of functional specialization in long-term couples, and they found a significant decrease in the number of items that were remembered when couples were told which items to focus on. These are interesting patterns in their own right. But unfortunately they can tell us little about the *transactions* that occurred between partners. Perhaps some of the groups in Wegner's studies relied on gestural cues or other types of transactions to produce shared genuinely memories. But there may also have been groups whose partners exploited one another as informational resources, relying on assumptions about what information they could count on their partners to store and nothing more. Importantly, there has been subsequent research on small work-groups that supports the claim that there is little difference on some tasks between the performance of a group who trains together and the performance of a group who learns explicitly which tasks the other members will perform (Moreland & Myaskovsky 2000; but cf., Ren & Argote 2011). I do not deny the possibility that genuinely distributed cognition could derive from functional specialization alone. But there is a potential lumping error that arises in adopting a coarse-grained functionalism as a foundation for thinking about transactive memory (cf., Craver 2007, 123). Group-level data on their own come up short, because they cannot distinguish between the process of construction and the process of retrieval, and they cannot distinguish cases where people *use* information from cases where people are parts of distributed and *transactive* memory systems. Such data reveal that specialization *suffices* for some kinds of increased memory in couples, and they demonstrate ways in

---

irresolvable differences in how people want to use the concept “memory”, for example. Some people may chose to adopt a narrower usage, which focuses on memories that are constructed using the sorts of mechanisms I have been discussing in this paper; while others may chose to adopt a more open-ended perspective, acknowledging that there are significant differences between internally constructed memories and memories that rely on 'exograms', while at the same time calling attention to the fact that “Brains like ours need media, objects, and other people to function fully as minds” (Sutton 2010, p. 205). I do not want to wade into disputes over which perspective is correct, here. And my pragmatist commitments suggest that there probably isn't an answer to be given as to which perspective is correct. My project in this paper, and indeed in my work more generally, is to see how far a traditional approach to cognitive science can be pushed in examining claims about distributed cognition. My claim is that even if you start with an individualist approach to cognitive systems, there are some cases where cognition extends beyond the boundaries of skin and skull. For a detailed, and far more expansive approach to the dimensions along which cognitive systems can be integrated with aspects of their material and social world, see Heersmink (2014).

which task demands impact specialized individual recall. But more must be done to demonstrate that these are cases of *transactive* memory.

We can make progress on this issue by thinking carefully about the implications of the fact that empirical research consistently reveals cases where groups are worse at remembering things than individuals would be if they were working on their own (see Harris et al 2011 for a review). In many cases, groups working in laboratory experiments don't cross-cue one another in ways that can produce transactive memories (Harris et al 2011). At best, they find ways of exploiting their assumptions about the things that other members are likely to remember—and it turns out that people are not always as good at figuring out what others will remember as they think they are. The strategies they adopt in trying to remember something often conflict, preventing them from developing an integrated processing solution that would allow them to perform *well* as a single cognitive system. This is part of the reason why friends and couples who collaborate effectively are able to produce more detailed episodic memories (Harris et al 2011). By collaborating on the construction of a shared narrative representation, friends and couples can implement an interfaced system that relies on the stored representations that happen to be spatially distributed.

This approach to transactive memory systems suggests that Wegner was wrong to claim that our frequent and pervasive use of Google searches and iPhones is sufficient to establish the existence of novel transactive memory systems. In these cases, the flow of information is unidirectional and exploitative. In these cases, we find a *person* who uses the informational resources, and who encounters information that is structured in a way that makes it a target for exploitation. But *exploitation* is the paradigmatic relation that obtains between a person and the tools that she uses. Sometimes when I use a web-search to trigger a memory, the information I acquire serves as an input into a process of counterfactual construction, which is carried out by a network of interfaced mechanisms that are internal to my brain; other times, I realize that it would be silly to try to keep track of everything that goes on in the web, so when I want something I *just look* (cf., Haugeland 1998a, 219)

That said, Chalmers-and-his-iPhone could come to constitute a cognitive system if there were *transactions* between Dave and his iPhone that yielded genuine informational integration. For this to occur, his iPhone (and its component Siri system, or a decendent) would have to posses the capacity to dynamically engage with Dave in a way that allowed for the cuing and re-cuing of skeletally-encoded representations in each system, and it would require the production of a shared episodic or counterfactual representation of the world *they* inhabited together by way of a shared process of counterfactual construction. Put differently, we would have to live in a world where people and iPhones workd together, *as a team*, to solve the problems they faced together.

While John Haugeland (1998a) suggested otherwise, I maintain that it is a mistake to think that we collaborate with *the road* to get to San Jose. We simply exploit the information we find as we drive down the road—and this information makes it a lot easier to San Jose than it would be without access to this information. The same thing holds for the way in which most of us have come to use the Internet. While Google doesn't know the way to San Jose, we can exploit Googlemaps to help us get there from almost anywhere—whether by car, by bike, or even by public transportation. But asking a friend, "Do you know the way to San Jose", may sometimes trigger the construction of a shared memory of a life lived together. And this kind of cognitive interdependence may play a critical role in a life worth remembering.[10]

7. Works Cited

Adams F. & K. Aizawa (2008). *The bounds of cognition*. Oxford: Blackwell.

Addis, D., A. Wong, & D. Schacter (2007). "Remembering the past and imagining the future," *Neuropsychologia*, 45: 1363-1377.

Addis, D., & D. Schacter (2008). "Effects of detail and temporal distance of past and future events on the engagements of a common neural network," *Hippocampus*, 18: 227-237.

Anderson, M., R. Bjork & E. Bjork (1994). "Remembering can cause forgetting,"Journal of Experimental Psychology, 20: 1063–1087.

Anderson, M., & B. Levy (2002). "Inhibitory processes and the control of memory retrieval," Trends in Cognitive Sciences, 6: 299–305.

Anderson, M. (2007). "The massive redeployment hypothesis and the functional topography of the brain," *Philosophical Psychology*, 21, 2: 143-174.

Anderson, M. (2010). "Neural reuse," *Behavioral and Brain Sciences*, 33, 4: 245-313.

Anderson, M., M. Richardson, & A. Chemero (2012). "Eroding the boundaries of cogntion," *Topics in Cognitive Science*, 4: 717-730.

Azulay, J., S. Mesure, B. Amblard, O. Blin, I. Sangla, & J. Pouget (1999). "Visual control of locomotion in Parkinson's disease," *Brain,* 122, 1: 111-120.

Bartlett, F. (1932). Remembering. (Cambridge: Cambridge University Press).

Bechtel, W. (2009). "Explanation: Mechanism, modularity, and situated cognition," in P. Robbins & M. Aydede (3ds.). *Cambridge handbook of situated cognition*. (Cambridge: Cambridge University Press): 155-170.

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

Chalmers, D. (2008). Foreword to Andy Clark's *Supersizing the mind*. (New York: Oxford University Press).

Clark, A. (2003). *Natural born cyborgs*. (Oxford: Oxford University press).

Clark, A. & D. Chalmers (1998). The extended mind, *Analysis* 58: 10-23.

Coman, A., D. Manier, & W. Hirst (2009). "Forgetting the unforgettable through conversation," *Psychological science*, 20, 5: 627-633.

Craver, C. (2007). *Explaining the brain*. (New York: Oxford University Press).

Cuc, A., J. Koppel & W. Hirst (2007). "Silence is not golden," *Psychological Science*, 18: 727–733.

De Brigard, F. (2013) "Is memory for remembering?" *Synthese,* 191 (2): 1-31.

De Brigard, F. (2014). "The nature of memory traces," Philosophy Compass, 9 (6): 402–414.

De Brigard, F., D. Addis, J. Ford, D. Schacter, & K. Giovanello (2013). "Remembering what could have happened," *Neuropsychologia*, 51 (12): 2401-14.

[10] I am grateful for the extensive comments that I received from the anonymous reviewers who have kindly read various drafts of this paper. I would also like to thank the audience at the Southern Society for Philosophy and Psychology where an early version of this paper was presented; and I would like to thank Andy Blitzer, Tony Chemero, Felipe De Brigard, Carrie Figdor, Ruth Kramer, Sarah Robins, John Sutton, and Georg Theiner for helpful discussions on this topic, and for opening my eyes to important issues that I would otherwise have missed.

De Jaegher, H., E. Di Paolo, & S. Gallagher (2010). “Can social interaction constitute social cognition?” *Trends in Cognitive Sciences*, 14, 10: 441-447.
Deese, J. (1959). “On the prediction of occurrence of particular verbal intrusions in immediate recall,” *Journal of Experimental Psychology*, 58: 17-22.
Fauconnier, G., & M. Turner (2003). *The Way We Think*. (New York: Basic Books).
Gluck, M., E. Mercado, & C. Myers (2010). *Learning and Memory*. (New York: Worth Publishers.
Harris, C. P. Keil, J. Sutton, A. Barnier, & J. D. McIlwain, “We Remember, We Forget,” *Discourse Processes 48*, 4: 267-303.
Haugeland, J. (1998a). “Mind embodied and embedded,” in *Having Thought*. (Cambridge: Harvard University Press): 207-237.
Haugeland, J. (1998b). “Representational genera,” in *Having Thought*. (Cambridge: Harvard University Press): 171-206.
Haugeland, J. (2004) “Closing the Last Loophole” *Inquiry,* 47: 254-266.
Heersmink, R. (2014). “Dimensions of integration in embedded and extended cognitive systems,” Phenomenology and the cognitive sciences, 13: 1-22.
Huebner, B. (2014). Macrocognition. New York: Oxford University Press.
Hutchins, E. (1995). *Cognition in the wild.* Cambridge: MIT Press.
Hutchins, E. (2005). "Material anchors for conceptual blends," *Journal of Pragmatics* 37, 10: 1555-1577.
Kaplan, DA. (2012). “How to demarcate the boundaries of cognition,” *Biology and philosophy,* 27 (4): *545-570.*
Kirchhoff, M. (2013). Extended Cognition & Constitution: Re-Evaluating the Constitutive Claim of Extended Cognition. *Philosophical Psychology*, 27 (2): 1-26.
Kosslyn, S. (2005). “Mental processes: An out-of-body existence?” *Edge, the world question center*, http://edge.org/q2005/q05_10.html (Accessed 1 Nov 2012).
Kosslyn, S. (2006). “On the evolution of human motivation,” in S. Platek, T. Shackelford, & J. Keenan (Eds.), *Evolutionary cognitive neuroscience*. (Cambridge: MIT Press).
Loftus, E. (1975). “Leading questions and the eyewitness report,” *Cognitive Psychology*, 7: 560-572.
Loftus, E & J. Palmer (1974). "Reconstruction of Automobile Destruction," *Journal of Verbal Learning and Verbal Behavior*, 13: 585–9.
Loftus, E. & J. Pickrell (1995). “The formation of false memories,” *Psychiatric Annals*, 25: 720-725.
Menary, R. (2007). *Cognitive Integration*. New York: Palgrave MacMillan.
Menary, R. (2010). “Cognitive integration and the extended mind,” *The extended mind,* Menary, R. (ed). Cambridge: MIT press, 227-243.
Moreland R., & L. Myaskovsky (2000). “Exploring the performance benefits of group training” *Organizational Behavior and Human Decision Processes*, 82, 1: 117-133.
Nadel, L. & M. Moscovitch (1997). “Memory consolidation, retrograde amnesia and the hippocampal complex,” *Current Opinion in Neurobiology, 7:* 217-227.
Neisser, U. (1967). *Cognitive Psychology*. (New York: Appleton).
Neisser, U. (1981). “John Dean's memory,” *Cognition*, 9: 1-22.
Piccinini, G. & C. Craver (2011) “Integrating Psychology and Neuroscience,” Synthese, 183: 283-311
Prokop T., M. Schubert, & M. Berger (1997). “Visual influence on human locomotion,” Experimental Brain Research, 114: 63–70.
Pylyshyn, Z. (1984). Computation and cognition. (Cambridge: MIT Press).

Ren Y. & L. Argote (2011): “Transactive memory systems 1985–2010: An integrative framework of key dimensions, antecedents, and consequences,” The Academy of Management Annals, 5: 189-229.

Robins, S. (2012). The constructive trace theory of memory (Unpublished doctoral thesis). Washington University in St. Louis. St Louis, Mo.

Robins, S. (submitted). “Misremembering and the constructive trace theory of Memory.”

Roediger, H & K. McDermott (1995). “Creating false memories,” *Journal of Experimental Psychology, 21*: 803-814.

Rowlands, M. (2010). New science of the mind. (Cambridge: MIT Press).

Rupert, R. (2004). “Challenges to the hypothesis of extended cognition,” *Journal of Philosophy* 101, 8: 389-428.

Rupert, R. (2010). *Cognitive systems and the extended mind*. (New York: Oxford University Press).

Schacter, D. & D. Addis (2007). “The cognitive neuroscience of constructive memory,” *Philosophical Transactions of the Royal Society B, 362, 773-786.*

Sellars, W. (1965). “Scientific realism or Irenic instrumentalism,” in R. Cohen & M. Wartofsky (eds). Boston studies in the philosophy of science, vol. 2, New York: Humanities Press, 171-204.

Silberstein, M. & T. Chemero (2011). “Complexity and extended phenomenological-cognitive systems,” *Topics in cognitive science*, 4 (1): 35-50.

Simon, H. (1996). *The Sciences of the Artificial*, (Cambridge: MIT Press).

Skolnick, D. & P. Bloom (2006). “What does Batman think about SpongeBob?” *Cognition*, 101, 1: B9-B18.

Skolnick, D. & P. Bloom (2006). “The intuitive cosmology of fictional worlds,” in S. Nichols (Ed.), *The Architecture of the Imagination* (Oxford: Oxford University Press).

Sparrow, B., J. Liu, & D. Wegner (2011) "Google Effects on Memory: Cognitive Consequences of Having Information at Our Fingertips," *Science* 333: 776-778.

Sutton, J. (1998). Philosophy and memory traces: Descartes to connectionism. Cambridge: Cambridge University Press.

Sutton, J. (2010). Exograms and interdisciplinarity: History, the extended mind and the civilizing process. In R. Menary (Ed.), The extended mind (pp. 189–225). Cambridge: MIT Press.

Sutton, J., et al. (2011). “The psychology of memory, extended cognition, and socially distributed remembering,” Phenomenology and the cognitive sciences 9, 4: 521-560.

Theiner, G. (2013). “Transactive memory systems: A mechanistic analysis of emergent group memory,” Review of Philosophy and Psychology 4, 1: 65-89.

Tribble, E. (2011). Cognition in the Globe. (New York: Palgrave Macmillan).

Tribble, E. & N. Keene (2011). Cognitive Ecologies and the History of Remembering. (New York: Palgrave Macmillan).

Wegner, D. (1995). “A computer network model of human transactive memory,” Social Cognition, 13, 1-21.

Wegner, D. (2012). “Don’t fear the cybermind,” The New York Times (5 August 2012): SR6.

Wegner, D., R. Erber, & P. Raymond (1991). “Transactive memory in close relationships,” *Journal of Personality and Social Psychology*, 61, 923-929.

Wegner, D., & D. Gilbert (2000). Social psychology: The science of human experience. *Subjective experience in social cognition and behavior.* H. Bless & J. P. Forgas (Eds.). Philadelphia: Psychology Press: 1-9.

Wegner, D., T. Giuliano, & P. Hertel (1985). “Cognitive interdependence in close relationships,” in W Ickes (ed), *Compatible and Incompatible relationships*. (New York: Springer-Verlag): 253-276

Wegner, D. & A. Ward (2013). “The Internet has become the external hard drive for our memories”, Scientific American, 309, 6.

Wegner, T. & Wegner, D. (1995). “Transactive memory,” in A Manstead and M Hewstone (eds), *The Blackwell Encyclopedia of Social Psychology*. (Oxford: Blackwell): 654-656.

Weiskopf, D. (2008) “Patrolling the mind’s boundaries,” *Erkenntnis* 68: 265–276.

Weiskopf, D. (2010). “The Goldilocks problem and extended cognition,” *Cognitive systems research* 11: 313-323.

Wilson, R. (1994). “Wide Computationalism,” *Mind* 103 (411): 351-72.

Wilson, R. (2010). “Review of *Cognitive systems and the extended mind,”* http://ndpr.nd.edu/news/24309-cognitive-systems-and-the-extended-mind/