Propping up the Causal Theory

Peter Langland-Hassan

Abstract: Martin and Deutscher's (1966) causal theory of remembering holds that a *memory trace* serves as a necessary causal link between any genuine episode of remembering and the event it enables one to recall. In recent years, the causal theory has come under fire from researchers across philosophy and cognitive science, who argue that results from the scientific study of memory are incompatible with the kinds of memory traces that Martin and Deutscher hold essential to remembering. Of special note, these critics observe, is that a single memory trace can be shaped by multiple past experiences. This appears to prevent traces from underwriting Martin and Deutscher's distinction between remembering an event and merely forming an accurate representation of it. This paper accepts such criticisms of the standard causal theory and, through considering the phenomenon *forgetting through repetition*, raises several others. A substantially revised causal theory is then developed, compatible with the thesis that individual memory traces are shaped by multiple past experiences. The key strategy is to conceive of episodic remembering not as the simple retrieval and projection of a static memory trace, but as a complex quasi-inferential process that makes use of multiple forms of information and cues—"prop-like" memory traces included—in generating the experience known as episodic remembering. When remembering is understood as a multi-componential process, there are a variety of ways in which a representation of the past may be appropriately causally dependent upon a prior experience of the event remembered.

I.

I cannot remember the Battle of Waterloo because I was not there. But, even if I had been there, accurately representing that event would not suffice for remembering it. To accurately visualize one's own recent heart surgery, based entirely on the surgeon's testimony, is not yet to remember it. Remembering an event requires something more. After years of relative calm, the philosophy of memory is now in turmoil over what more that might be. It is a crisis brought on by increased attention to the developing neuroscience of memory. And it is one mired in the question of where the philosophy of memory ends and its science begins.

For five decades, Martin and Deutscher's (1966) causal theory of memory held sway as an account of remembering's necessary and sufficient conditions. They began their analysis by proposing three conditions necessary for a person S to successfully[1] episodically remember an event e through use of a mental state r. Paraphrasing slightly[2]:

1) S must, with certain limits of accuracy, represent e through the use of r.
2) S must have observed e.
3) *S's* past experience of e must have been causally operative in producing a state or successive states in S finally operative in producing r (1966, p. 166).

In condition (3), now known as the *causal condition* in the theory, Martin and Deutscher specify that one's observing e must be causally "operative in" producing a state that is later operative in producing r.

The rub, and source of current controversy, concerns two amendments they add to condition 3 in order rule out apparently deviant causal chains. If successful, the amendments serve to transform conditions 1, 2 and 3 from merely necessary conditions to necessary *and* sufficient conditions for successful remembering. The first—which they dub clause 3b[3]— specifies that S's past experience of e:

---

[1] If one hears the verb 'remember' as factive, then speaking of "successful" remembering is arguably redundant. However, some (e.g. Fernández (2019) and Michaelian (2016c)) favor a non-factive understanding of 'remember.' For that group, there is an important distinction between remembering and successful remembering. In specifying that our topic is successful episodic remembering, I aim to bring both groups into conversation around a single question: what are the necessary and sufficient conditions for successful remembering?

[2] The differences between the paraphrase and original are the following: Martin and Deutscher give criteria for remembering "something" or "a thing" (p. 166) as opposed to an event, though their broader discussion clarifies that memory of events is their primary explanandum (see, e.g., their focus on "remembering that$_2$" over "remembering that$_1$" (pp. 162-3)), and this is indeed how their theory is typically interpreted. Also, Martin & Deutsher speak simply of the need to "represent e" and of S's "representation of e" without assigning a variable *(r)* to refer to that representation. Finally, they provide a slightly different account of how internal (mental) events are remembered than external events, while the present version focuses only on external events so as not to engage debates about the nature of introspection.

[3] Clause 3, above, becomes clause 3a when Martin and Deutscher add sub-clauses 3b and 3c.

Is operative in producing the state (or the successive set of states) in him which is finally operative in producing the representation [r], *in* the circumstances in which he is prompted (p. 185, emphasis in original).

The role of this condition is to ensure that the immediate cause of *r* (one's representation of *e*) is a state that was itself produced during one's observation of *e*, thereby excluding deviant causal chains of the sort involved in *relearning*. In Martin & Deutscher's example of relearning, Kent is in an accident and tells his friend Gray what the experience was like. Unlucky Kent later has a second accident that produces in him total amnesia with respect to the first. Gray then informs Kent about the first accident in detail, enabling Kent again to accurately represent that event. Still later, Kent forgets that Gray has helped him "relearn" his own history and takes himself to remember the first accident. Intuitively, Kent does not remember the accident. He has instead (re)*learned* about it, in much the way anyone might learn about—and thereby come to accurately represent—the accident simply by listening to Gray's detailed account. However, it remains true that Kent's experience of the first accident was causally operative in producing his later representation of it, in line with condition 3. His being in the accident caused him to tell Gray about it, after all, such that Gray could later inform him of his own forgotten past. The problem in this and other cases of relearning seems to be that there is no mediating state (or successive set of states[4]) in Kent that was *both* produced by the observation of *e* (his first accident) and also operative in the very circumstance where Kent later generates *r*, his representation of *e*. Condition 3b stipulates that there must be such a state, or succession of states.

Martin and Deutscher go on further to require, in condition 3c, that this mediating state is a "structural analogue" of the thing remembered. This condition is now seen by most as too restrictive in its specification of the representational format of the state, treading too far into the realm of what must be empirically discovered (Bernecker, 2010; Michaelian, 2011; Michaelian & Robins, 2018). A more neutral and still widely endorsed version of the condition holds simply that the mediating state must be a *memory trace*. While there is no single conception of what it is to be a memory trace that is both precise and commonly shared, most assume that traces are

---

[4] We can add that all the states in this succession must be states *of the same kind*, as seems to be Martin & Deutcher's intent. If the states in the succession can be of different kinds, this would allow the succession to travel from Kent's brain to Gray's and back, via testimony, which is what condition 3b seeks to forbid. In any case, the revised causal theory I develop later does not appeal to the notion of 'successive states,' so I will not linger on it here.

mental representations of some kind that causally mediate between an experience of an event and the mental representation by which the event is recalled (Barry & Maguire, 2019; De Brigard, 2014b; Michaelian & Robins, 2018; Moscovitch & Nadel, 2019; Robins, 2017).[5]

Going forward, our focus will be on the still popular condition 3b, interpreted as to invoke a memory trace as the relevant state that is both produced by e and operative in the very circumstance in which r is generated. In order for r to be a successful remembering of $e,$ 3b requires that there is a memory trace, $t,$ that connects the causal dots between e and $r,$ mapping a discrete causal pathway from one to the other. It is in virtue of memory trace $t's$ being produced by an experience of event e—and *only* event e—that t's subsequent triggering of r qualifies r as a remembering of $e,$ and not some other event. Sensibly, it seems, 3b does not allow there to be a memory trace that was produced prior to, or long after, the event it enables one to remember.

It is in specifying that each trace is *produced by* the event it enables one to remember that Martin and Deutscher ensure that each memory trace will enable the remembering of exactly one event—that is, so long as we assume that to be "produced" is to be brought into existence, and that a thing can only be brought into existence once.[6] In that sense—and to coin a term—3b requires that traces are *monogamous*, each enabling the remembering of exactly one event (specifically, the event that produced it). Trace monogamy is assumed, if at times only implicitly, by the many who have continued to defend versions of the causal theory (Bernecker, 2010; Cheng & Werning, 2016; Michaelian, 2011).

The issue now disrupting the philosophy of memory, and which serves as the fulcrum for this paper's main argument, is that many now working in the cognitive science of memory hold that traces are in fact *promiscuous*, being both caused by and supporting the remembering of multiple distinct events.[7] On this view, to be spelled out in more detail below, individual traces are less like recordings of particular past events and more like reusable stage props that can be

---

[5] Recently, however, some have proposed views on which traces are "contentless"—and thus not *representations*—yet still causally mediate between an experience of event e and its subsequent remembering in r (Hutto & Peeters, 2018; Perrin, 2018; Werning, 2020).

[6] I will consider, in Section Two, the alternative possibility that a single trace is "produced by" perceptions of multiple distinct (if similar) events, all of which it may allow one to remember. We will see that allowing for this possibility creates gaps in the causal theory of memory as it is standardly developed.

[7] To be sure, scientific work on memory is not univocal on this point, with some theorists still pursuing—and finding some empirical support for—a more traditional, monogamous view of traces (Najenson, 2021; Robins, 2020).

called upon in the reconstructive representation of many distinct past events. This view of traces is especially common among those who see in remembering a form of "mental time travel," not different in kind from episodic future thought and episodic counterfactual thought (Addis, 2020; De Brigard, 2014a; Michaelian, 2016c; Schacter, Addis, & Buckner, 2007; Tulving, 1983). Some who advocate for a distributed view of memory traces, in line with connectionist principles in cognitive science, also appear committed (if at times unwittingly) to the promiscuity of traces, as argued by Robins (2016) and Michaelian and Sant'Anna (2019). If we can indeed make use of a single stored trace in the successful remembering of temporally distinct events, then traces are sometimes used in the remembering of events that did not produce them.[8] To accept this is to reject principle 3b.

However, it is far from clear that the causal theory itself can survive without 3b. The central thought behind any causal theory is that episodes of successful remembering must be *appropriately caused* by the events they represent (Debus, 2017). The possibility of deviant causal chains of the sort involved in relearning shows that the *mere* causal dependence of *r* on one's experience of *e* is not sufficient for remembering. Yet it has proven difficult to see what it could mean for an episode to be "appropriately" caused *other than* for there to be a discrete causal chain, leading from experience to recollection, where a monogamous memory trace provides the links in-between.

In the face of this obstacle, some advocates for the mental time travel framework have rejected the causal theory outright, omitting a causal condition from their accounts of successful remembering (Addis, 2020; Michaelian, 2016c; Michaelian & Sant'Anna, 2019). Michaelian (2016c) has developed the most sophisticated version this non-causal "simulationist" account of episodic remembering, on which successful remembering occurs whenever a reliably functioning "episodic construction system" issues in an accurate representation of an event from one's past. Notably, Michaelian (2016b) openly struggles to distinguish remembering from relearning and has more recently come to think that the two should not be distinguished (Michaelian, 2021). Most, however, still want to find room for the idea that successful remembering requires an

---

[8] This entailment holds so long as we assume that a particular trace can only be produced by the perception of one event (see fn. 6), but not if we allow that a particular trace can be produced by multiple past perceptions (because it is a kind of summation of those perceptions). However, going the latter route (while defending a causal theory) creates significant problems that I will discuss in Section Two.

appropriate causal relation between the act of remembering and the event remembered—at least as a necessary condition (Debus, 2017). To that end, some have proposed that monogamous memory traces are in fact consistent with contemporary memory science and, in particular, with the kinds of reconstruction and reconsolidation processes that occur during memory storage and retrieval (Robins, 2020). Others have defended revisionary, non-representational accounts of traces that take on board some of the simulationists' critiques, while still allowing traces to play the role of providing a distinct ("monogamous") causal pathway between an experience and its remembering (Werning, 2020).

Here I want to chart an alternative path forward for the causal theory—one consistent both with the promiscuity of memory traces and the highly constructive nature of episodic remembering. While I will offer reasons of my own for thinking that traces are indeed promiscuous, my primary aim is not to establish that point but to show how the promiscuity of traces is compatible with a substantially amended causal theory of remembering. That theory is developed and defended in the final three sections of the paper (Sections 7, 8, and 9). Prior to that, there are two important elements of scene-setting required. First, I will describe two different metaphors for thinking about memory traces—the prop theory and the replay theory—that help to clarify current debates and provide important tools for the discussion to come. Then, in Section 3, I will describe the familiar if underappreciated phenomenon of *forgetting through repetition*. Meditation on this phenomenon (in Section 4, 5, and 6) reveals challenges for both the standard causal theory and the new simulationism that rejects causal connections as essential to successful remembering. Importantly, this meditation also points the way to a substantially revised causal theory (Sections 7, 8 and 9) that better harmonizes with 21st Century cognitive science.

II.

To begin, it will be useful to have in hand two competing metaphors for understanding the relationship between episodic remembering and memory traces.[9] The first metaphor—which I will call the *replay theory*—is modelled on the view of memory traces found in standard

---

[9] These metaphors are also sketched and put to somewhat different use in Langland-Hassan (forthcoming).

versions of the causal theory, as defended, for example, by Martin & Deutscher (1966) and Bernecker (2010). I will understand a "standard" causal theory as one that explains appropriate causation by appeal to a monogamous memory trace. The second metaphor—which I will call the *prop theory*—is modelled on what I take to be the view of *simulationists*, such as Michaelian (2016b) and Addis (2020), who see remembering and imagining the future as the exercise of a single faculty for "mental time travel." The prop theory also aims to cohere with the view that traces are distributed patterns of connection weightings among neurons in neural networks, where the weightings definitive of a single trace may be shaped by multiple past experiences (De Brigard, 2014a; Robins, 2016; Sutton, 1998). Both metaphors are idealized and admittedly picturesque renderings of the theories they take as inspiration and are designed to highlight their differences. In fleshing them out, I will discuss some of the evidence brought to bear in support of each, focusing especially on motivations for the prop theory.

On both metaphors we can conceive of remembering as an event that occurs on a stage, where on-stage events are analogized to the contents of the conscious mind. According to the replay theory—our first metaphor—the stage features a large screen on which images are projected. Lucy, our would-be rememberer, counts as successfully remembering an event whenever a memory trace, recorded during some event in her past, is projected on the screen. The replay theory has it that Lucy's memory traces are stored offstage in a large library of such traces. For each trace in the library, there is exactly one event it enables Lucy to remember. Thus, on the replay theory, traces are monogamous. Each trace stores information about the event it allows Lucy to remember by being a sketchy, typically incomplete, recording of her past perceptual (or introspective) experience of the event—a recording made at the time of the event and dutifully stored ever since, more or less unaltered, in an offstage library of such traces. Thanks to the monogamy of traces on the replay theory, the question of *which* event Lucy is remembering is settled by the identity of the memory trace that is retrieved and projected. This allows for Lucy to make substantial mistakes about what she is remembering. If, for instance, Lucy is projecting a trace recorded at her eighth birthday party, then it is her eighth birthday she is remembering, even if she takes herself to be remembering her sixth birthday party.

This feature of the replay view harmonizes with Martin and Deutscher's example of the painter who, taking himself to be imagining and painting a purely invented scene, is (Martin &

Deutscher claim) in fact remembering a scene from his childhood (1966, p. 167-168). A not always recognized price of this approach is a commitment to the monogamy of memory traces. To see this, consider how things change if the causal theorist allows that a memory trace t can be used to remember two distinct events, e and e', in cases where perceptions of both e and e' were partially responsible for the production of t. (I am now suspending the earlier assumption that a single trace can only be produced by the experience of one event.) In that case, the fact that one's representation r was caused by t will not settle the question of which event (e or e') was remembered with r—and thus will not settle whether e was remembered successfully. Thus, to allow that traces are promiscuous, even in the limited way just proposed, would require Martin & Deutscher to give up their claim of having offered sufficiency criteria for the remembering of e. This is the "gap" in their account mentioned in footnotes six and eight.[10]

The replay view, with its commitment to monogamous traces, has a tidy explanation of the difference between remembering e and (merely) accurately representing e: the former occurs whenever a trace recorded at e is projected, whereas the latter occurs when e is accurately represented through any other means. The replay theory nevertheless remains compatible with "constructive" versions of the causal theory—such as Michaelian's (2011)—that allow successful remembering to feature *some* content not originally represented during the event remembered and that does not derive from the projection of a memory trace. Precisely how much content can derive from other sources while the event remains a successful remembering can be left open.

On the second metaphor—the prop theory—memory traces are more like stage sets used in a theater production, where each set is composed of various (reusable) props. The prop theory also holds that memory traces (and the props that compose them) spend most of their lives offstage, in storage. However, just as a theater company will reuse the same kitchen set for each scene that occurs in the kitchen, so too, on this metaphor for remembering, are individual props reused in the representation of distinct events. On the prop theory, traces are promiscuous. Lucy episodically remembers an event when her cognitive system assembles one or more props (that

---

[10] The door remains open to modified causal theories that do not take causation by a memory trace to be sufficient for remembering and which allow for imaginings to be caused by traces. However, in order to provide sufficiency conditions for successful remembering, any such causal theory will then need to add a condition of a sort not found in Martin and Deutscher's classical account—such as an appeal to the agent's intentions. I develop two such approaches below (Sections Eight and Nine).

together compose a particular set) in a representation of an event from her past. Each set enables Lucy to represent one or more past events by representing features of those events, during onstage performances of the past. Due to the promiscuity of traces, on the prop theory, the question of *which* event Lucy is remembering is not answered by the memory trace (or traces) currently on stage. Nor does the prop theory, as so far sketched, offer an account of the difference between Lucy's remembering a past event and (merely) accurately representing it.

An important feature of the prop theory is that props can be used for purposes other than remembering. Lucy sometimes uses props not in the representation of her own past but of likely future events, or of ways things could have gone differently. When she does so, the results are not cases of remembering but, rather, cases of imagining the future, or imagining counterfactual pasts. This is despite the fact that the very same props (and the sets they compose) might have been used in an act of successful remembering. This aspect of the prop theory meshes with the simulationist's claim that imagining the future and episodic remembering are fundamentally the same kind of constructive (and reconstructive) process, oriented in opposing temporal directions (Addis, Wong, & Schacter, 2007; De Brigard, 2014a; Michaelian, 2016a).

Unlike the replay view—where traces are not substantially altered after they are first created—the prop theory holds that work continues on props day to day, with details added or features removed. Entirely new props and sets are constructed only as needed to accurately depict one's life events. When Lucy visits Alaska for the first time, the stage crew gets to work in creating sets corresponding to her new experiences. Otherwise, when she stays home, the stage crew busies itself with maintaining and elaborating existing sets. If Lucy starts spending more time in her living room, her existing living room props are strengthened and elaborated in response to those experiences. As a result of this continuing work, it is impossible to pinpoint any particular past experience that is the causal source for the living room props. This feature of the prop theory coheres with the simulationist's vision of memory traces and remembering— expressed by Michaelian and Sant'Anna (2019) —on which "contents originating in experiences of multiple earlier events may contribute to the content of a single retrieved representation…because they make their way into the trace resulting from the experience of the remembered event" (p. 20). The view that a single trace may indeed be shaped by multiple distinct experiences marks a trend in neuroscience that sees memory traces as subject to

processes of *reconsolidation*, during which stored traces are "reopened" so as to be updated and refined in the light of new experiences (Dudai, 2012; Nadel, 2007; Sara, 2000).[11] It also coheres with the view of some neuroscientists that rapid structural and functional changes in the hippocampus are both normal and at odds with the neural storage of a memory trace that remains static from encoding to retrieval (Barry & Maguire, 2019).[12]

Notably, the prop theory takes on board Sarah Robins's (2016) critique of recent causal theories that view memory traces as distributed patterns of activation in connectionist neural networks. (Robins's targets include Bernecker (2010) and Michaelian (2011), who are themselves inspired by Sutton's (1998) account of traces.) When we look closely at the details of such accounts, she argues, we see that they leave the causal theory, as articulated by these theorists, unable to accomplish its central aim of distinguishing remembering from deviant causal chains as are involved in relearning. The problem is that:

> Distributed traces do not have individually distinguishable causal histories…similar memories are blended together in overlapping patterns. Individual traces do not leave a lasting, distinctive mark on the network by which their unique causal influence on a subsequent representation could be detected (p. 3008-9).

According to Robins, if we accept the distributed conception of traces proposed by Bernecker, Michaelian, Sutton, and others, there is no longer a single causal pathway from experience to remembering that can underwrite the distinction between genuine remembering and various deviant causal chains. While Robins does not endorse the distributed conception of traces—she aims, instead, to reveal its unappreciated consequences—the prop theory accepts the kind of trace-promiscuity the distributed conception seems to require and grants Robins's point that there is no discrete causal chain leading from each trace back to exactly one event during which it was produced.

---

[11] It is possible, however, for there to be traces that are subject to process of reconsolidation processes yet that are still only used to remember one event. In that sort of situation, reconsolidation need not result in the kind of promiscuity that is problematic for standard causal theories. Whether, and how often, reconsolidation in fact results in trace promiscuity remains unsettled by current empirical work.

[12] "Just as the hippocampus is necessary for *de novo* construction of imagined scenarios," Barry and Maguire propose, "it could perform the same function for remote memories long after the original hippocampal trace has decayed…its fleeting role in consolidation may be subservient to its primary function as a constructor of scenes, whether past, future, or fictive" (p. 134, 139). See Moscovitch and Nadel (2019) for a competing account of the role of the hippocampus more amenable to its storage of persisting memory traces.

What, then, makes for *successful* remembering on the prop theory? How can the prop theory distinguish remembering an event from (merely) accurately representing it? To date, Michaelian (2016b, 2016c, 2021) has offered the most detailed response to such queries. Recognizing that a view on which traces are promiscuous cannot explain successful remembering in the way familiar to causal theories (and the replay metaphor), Michaelian holds that to remember an event is simply "to imagine an episode belonging to [one's] personal past," where the representation is "produced by a properly functioning episodic construction system which aims to simulate an episode from the personal past" (2016c, p. 97). Such imaginings, he explains, "need not draw on stored information ultimately originating in experience of the relevant episode" (2016c, p. 111). Michaelian's simulationism can still distinguish successful remembering from at least some cases of (mere) accurate representation, as there are conceivable situations where one's episodic construction system is not functioning properly but still happens to produce an accurate representation of an event from one's past.

Nevertheless, Michaelian leaves the simulationist—and prop theorist—without any obvious means for excluding cases of relearning from the class of rememberings. For relearning about an event, via testimony, seems to involve following a procedure that tends to produce accurate representations of one's past. Admitting this lacuna as an initially counterintuitive result, Michaelian maintains that, on reflection, there is no compelling reason no to count relearning as successful remembering (Michaelian, 2021). In a second radical departure from the causal theory, he further proposes that an event can be remembered that one did not, in fact, experience—so long as one was indeed present for the event (Michaelian, 2016c, p. 118). It is a short step to that conclusion once one is convinced that remembering an event does not require storing information gathered from the event itself. It is in part the apparent promiscuity of traces—which he finds to be suggested by multiple lines of empirical research—that pushes Michaelian toward these revisionary conclusions.

It is no refutation of Michaelian's simulationist approach that it issues in counterintuitive results. His aim is to articulate a suitable theory of remembering consistent with often surprising results in cognitive science, not to vindicate our pre-theoretic intuitions about remembering. Nevertheless, it is worth exploring whether these same results can be made consistent with some less counterintuitive theory of remembering—one that respects the powerful thought that we are

only remembering an event if the event appropriately caused our representation of it. This will be my project going forward, in pursuit of a causal theory consistent with the kind of trace promiscuity envisioned by the prop theory. I will approach the task indirectly, through meditation on the phenomenon of forgetting through repetition.

<div align="center">III.</div>

I have only made saag paneer once in my life. It was about six months ago, and I remember it well: marinating and browning the cubed paneer, then setting it aside; measuring out the spices; mixing the spinach together with garlic and onions in a large pan, and so on. By contrast, I have made steel-cut oatmeal many times—about once a week for the last three months, always following the same recipe. It is easy for me to visualize the process of making steel-cut oatmeal because I have carried it out so many times, in the same kitchen, in much the same way. While I know that I first made steel-cut oatmeal about three months ago—shortly after it was recommended to me by a friend—I can no longer remember making it that day. That is, while I know *that* I first made steel-cut oatmeal about three months ago, I cannot remember the actual *event* of making oatmeal that day—even if I am confident that I can accurately visualize what it was like when I did so, and even if I can visualize it as accurately as I can my episode of making saag paneer, six months ago. Further, I take it that, had I not made steel-cut oatmeal many times since that fateful day three months ago, but had *only* made it that day, I would now not only be able to accurately visualize making it but would remember making it as well. After all, it was much longer ago that I made saag paneer, and I remember that event very well. So it seems that the reason I cannot, now, remember the first time I made steel-cut oatmeal is *not* that I failed to encode and store information about my oatmeal-making that day. The problem is that I have made oatmeal too many other times since then. That is why I can no longer remember the first time in particular. The more frequently we have observed a certain kind of event, the more difficult it is to remember *particular* events of that kind, despite its being easier to accurately represent events of that general kind. This is the phenomenon of forgetting through repetition (hereafter, "FTR").

Another example: if you have seen the film *The Godfather* only once in your life, you likely remember the event of doing so. But if you have watched *The Godfather* fifteen times,

because it is your favorite film, it will be difficult to remember any particular event of watching it, even if you now have a much richer capacity to represent it than someone who has seen it only once. Because we are susceptible to forgetting through repetition, we often forget particular events that we would have remembered had we not subsequently observed more events of the same kind.

I will put FTR to three uses. First, it will be used to cast doubt on the standard causal theory, with its commitment to the replay theory of monogamous traces. Then it will feature in an argument against Michaelian's simulationist, non-causal account of successful remembering, discussed earlier. (There we will consider the reply that forgetting through repetition is not really forgetting.) And, third, I will use it to motivate a new causal theory of remembering that is compatible with memory traces of the promiscuous prop theory variety. This new version of the causal theory gains support through its ability to explain why we are susceptible to forgetting through repetition.

<div align="center">IV.</div>

Let us consider what questions forgetting through repetition raises for the replay theory of traces and the standard causal theories that invoke them. On the replay theory, remembering an event is a matter of retrieving and projecting a memory trace that was produced in an observation of that event (and no other). Forgetting, then, is a matter either of losing any such trace, or of losing the ability to retrieve and activate it. Consider the example of my having made steel-cut oatmeal many times. I do not think that I can remember the first time I ever made that dish, even if I am confident that I can accurately visualize the event. However, a couple of days after having first made steel cut oatmeal, I *could* still remember doing so, just as I can now remember the one time that I made saag paneer, six months ago. On the replay theory, one of two things must have happened: either the memory trace encoding my first making of oatmeal was erased; or, it has remained intact but is now inaccessible. With respect to the first option, not enough time has passed to make it likely that the trace would have been erased, just in the normal course of things, had I *not* made oatmeal many other times since. That is, it seems clear that the trace would still be available—that I would still be able to *remember* making oatmeal for the first time—had I not made oatmeal many times since. After all, I can still clearly remember

making saag paneer, which was much longer ago. But it is also not obvious why simply having more experiences of the same kind would serve to *erase* a memory trace (of the "replay" sort).

The most plausible account of such forgetting, for the replay theory, may be that I retain multiple distinct memory traces of making steel-cut oatmeal—one need not commit to how many—and that their similarity (in terms of content) makes it difficult to tell one from the other upon retrieval. It is this difficulty in retrieval, brought on by a similarity of contents among multiple traces, that clouds my sense that I can remember any particular episode of making oatmeal. For this to be the correct explanation, we must assume (as is in any case plausible) that the content of a memory trace does not typically include the time and date of the event it records. Otherwise, such time-and-date "stamps" would provide an easy means for distinguishing one content-similar trace from another. Yet, going this route, it is hard to see what rationale there would be for the brain to store multiple distinct representations with essentially the same content, without *also* having the means to distinguish among them on retrieval. Surely, having just one stored representation with the content *C* is as good as having fifteen that one cannot distinguish among. Thus, the proposed explanation of forgetting through repetition attributes a peculiar, self-defeating inefficiency to our memory systems. We routinely encode a new memory trace even though the net effect will be that we can remember *fewer* events than we could before we encoded it.[13]

We are left with a puzzle. While it is implausible that each memory trace possesses a distinct time-and-date stamp indicating its moment of production, it seems equally unlikely that the brain stores multiple content-indistinguishable traces relating to distinct events.[14] But if an inability to distinguish among multiple content-similar traces is not what causes forgetting

---

[13] If the replay theorist holds that traces can remain accessible despite being indistinguishable—like a marble in a bag of 100 exactly similar ones remains (in some sense) accessible if indistinguishable—then the problematic result is that we store new traces even though the net effect will be that we *think* we can remember fewer events than we could before (because we now cannot distinguish certain traces). From an epistemic standpoint, this is as bad as not being able to remember the events (while retaining the ability to accurately represent them).

[14] It may seem that this criticism clashes with a popular neuroscientific theory of memory known as the Multiple Trace Theory (MTT) (Moscovitch et al., 2005). However, the MTT holds that multiple content-similar traces are stored with respect to a *single* event and not with respect to many distinct events (Moscovitch et al., 2005, p. 42). In such a situation, the multiplicity of traces can increase the likelihood that the event represented is remembered, without the subject's inability to distinguish among them impairing her ability to remember the event (as would occur if the many content-similar traces represented distinct events). Note also that having many traces all of which represent the same event is not to have promiscuous traces, so long as each trace only represents one event.

through repetition, the replay theory is left without any obvious means for explaining the phenomenon.

A remaining possibility worth considering is that only one memory trace (of the replay theory variety) with a certain type of content is ever preserved, with its encoding serving to erase any previous trace with that type of content. One problem with this proposal is that it seems to get the phenomenology of forgetting through repetition wrong. In cases of FTR, we do not seem to remember just one instance of a certain kind of event to the exclusion of others. I do not, for instance, have a keen sense of remembering my most recent episode of oatmeal making, or yesterday's shower, to the exclusion of other such episodes. Second—and more decisively—this way of accounting for FTR leaves us without an explanation of why we get *better* at remembering certain types of events after experiencing them multiple times, despite getting worse at remembering individual episodes of that kind. As earlier noted, after watching *The Godfather* fifteen times, we will likely have forgotten the first event of seeing it, even if we can far more accurately remember scenes from the movie in general. If a causal theorist explains FTR by proposing that we delete prior traces relating to the film each time we watch it, then we should know the film no better after fifteen viewings than after one.[15] Obviously, that is not the case. But neither can the causal theorist (who invokes replay-style traces) hold that we have a single set of traces relating to the film that are elaborated and strengthened during each viewing, without giving up on the monogamy of traces. And we have seen that trace monogamy is essential to the causal theory's explanation of successful remembering.

Forgetting through repetition thus poses a serious challenge to the replay theory of traces and to any causal theory that relies on monogamous traces. Combined with the empirical motivations outlined above for accepting the prop theory's view of traces, we have good reason to pursue alternatives.

V.

What can the prop theory say about forgetting through repetition? Why, on the prop theory, am I unlikely to remember making oatmeal three months ago if I have made it many

---

[15] Thanks to Ellie Magill for helping me to fine-tune this challenge.

times since (but not if I haven't)? We should first consider the response of a simulationist who may hold that (so-called) forgetting through repetition is not *forgetting* at all. We saw, above, that in allowing for the promiscuity of traces, Michaelian also holds that someone can successfully remember an event *e* without there being any causal relation in place between a prior experience of *e* and the state *r,* by which *e* is remembered. All that matters is that the subject makes use of a reliably functioning episodic construction system that aims to represent the personal past, and that it does so accurately. If we are convinced that appropriate causal relations are not necessary to successful remembering—and thus not required for the proper functioning of the episodic construction system—it seems reasonable to conclude that, when I make use of my existing "making oatmeal at home" props to represent the first time I made steel-cut oatmeal, I am successfully remembering my past. Even if I now no longer have the sense that I can remember the first time I made oatmeal, that impression may be incorrect. For the simulationist, it may be that we often remember events without believing that we do. As earlier remarked, we should not be surprised that an empirically-informed account of memory—as the simulationist aims to provide—requires some revision in our ordinary intuitions about what counts as remembering (or forgetting).

However, this revision of intuitions would result in a significant expansion of the number of events that any person can now remember. For example: I know that I have exercised every day this week and that I have always put on my running shoes before doing so. Ordinarily, I would not say that I can remember each time I tied my running shoes this week. And yet, having a wealth of available props, honed through past experience, for accurately representing the tying of my running shoes, it would seem that I can indeed remember each distinct episode of tying them this week, if the simulationist is right. To do so, I simply aim to represent my shoe-tying of a certain day and use my episodic construction system to summon the suitable props. Likewise, suppose that I keep a log of my exercise activities and can see from it that I exercised 459 days ago. Knowing that I have had the same running shoes for two years, I can now also remember tying my shoes that day by using the same familiar set of props, summoned with the same aim of representing that particular past event. Indeed, it would seem that no event need be forgotten, so long as one has at hand props for accurately representing the event and knows one experienced it—be it a daily shower, a locking of the back door, or a brushing of teeth.

This is not a knock-down objection against simulationism, as the simulationist may be prepared to group acts as remembering in very counterintuitive ways. But it is suggestive of a deeper quandary: like everyone, the simulationist ultimately owes an account of how the successful episodic remembering of events in familiar settings occurs in the many cases where, as a matter of act, no outside information (from diaries, or testimony) is available to guide the remembering—as may occur, for instance, when you go to the local grocery store and, an hour later, are able to report back on what you did there. In such situations, how do you manage to assemble the right props, most of which you likely had *before* the visit (having been there many times)? The working hypothesis of the causal theory is that the fact that you observed the grocery store during the most recent visit will be an essential part of the explanation of how you are able to remember the visit an hour later—that if you hadn't observed the grocery store today, you would not now seem to remember having been there. The simulationist may wish to say that the same kind of causal dependency holds, at least in this (very common) sort of case. But supposing that traces are promiscuous and that you therefore had props sufficient for accurately representing the visit well before today, it is not at all obvious why the observations you happened to make during today's visit should themselves be explanatorily relevant to your ability to remember the visit. For the simulationist (and for anyone who accepts the promiscuity of traces), *what is it* about the way our memory systems work that so often leaves the remembering of events in familiar settings causally dependent on our having observed the event remembered? The worry is that, in answering, the simulationist will need to bring back in monogamous memory traces; and, once they are back on the table, simulationism itself may founder.

But never fear. In Sections 8 and 9, I offer alternative answers that, while advanced in defense of a causal theory, may be of use to the simulationist as well.

VI.

Assuming once again that forgetting through repetition is indeed forgetting, we can return to the question of what the prop theory should say about why it occurs. Reflection on FTR suggests that we are more likely to remember an event if we have witnessed that type of event exactly once. Of course, we sometimes remember distinct events that are similar enough to be

considered of the same type. And we also forget many events of a kind we have only experienced once. Nevertheless, it remains a familiar phenomenon that experiencing a certain type of event more than decreases the probability that we will remember specific events of that type.

On the prop theory, no changes of a kind that would explain FTR are made to the props themselves when I go from having made steel cut oatmeal once to having made it two, three, or four times. The props may become more detailed, or easier to access. Such gains do not explain why we forget specific events. We should instead search for something that is psychologically *lost* in the process of experiencing a certain kind of event more than once. One such item is the belief that we have observed the relevant type of event only once. We should thus consider the possibility that believing that we have only once witnessed a certain type of event is often an important psychological factor in the process of remembering it. Having such a belief will not suffice for being able to remember the event, of course. I might come to believe that I once ate a sea urchin, based on a friend's testimony, while having no memory of doing so. One needs, in addition, some stored props that are content-relevant. Perhaps, then, finding myself with several props detailing my kitchen as it looks when saag paneer is cooked, and believing that I have made saag paneer only once, I may have the sense that, in summoning the props, I am remembering the time I made saag paneer. Important to my sense of remembering in such cases is that I believe that I did indeed make saag paneer only once, and that I don't think that belief was arrived at through an outside source (unless, that is, I am the simulationist of Section 5). Extrapolating from such cases, a common process (but not the only process) by which we remember a specific event, consistent with the prop theory, may be by following what I will call the *One-off heuristic*:

> *One-off heuristic:* when attempting to remember event *e,* search for two things together: a set of stored props *P* that seem appropriately content-related to *e* and a belief that you have observed an event of *e's* kind exactly once. When these two states are found together, use *P* to generate a representation *r* of *e.* (But only believe you are

remembering if you don't believe that the aforementioned belief was acquired from an outside source.)[16]

Let us set aside, for the moment, the question of whether (and why) accurate representations *r* of *e* that are arrived at by following this heuristic should be considered *successful* rememberings. Our focus now is on the question of why forgetting through repetition occurs at all (on the prop theory). An answer the prop theorist can offer is that, when we have experienced a certain type of event multiple times, we tend to believe that we have. And when we believe that we have experienced a certain type of event multiple times, we can no longer employ the one-off heuristic for remembering events, as that heuristic requires a belief that we have experienced the relevant type of event exactly once. Let us call this sort of belief—that one has experienced a certain sort of event only once—a *one-off belief.* The cause of FTR, then, is not that we gain too many traces that are too similar to be distinguished, but that we lose a one-off belief of the right kind. It is the new unavailability of the one-off heuristic, resulting from loss of a one-off belief, that leads us to *forget* what we could previously remember. Forgetting in this way is consistent with retaining the ability to generate props that accurately represent forgotten instances of some type of event, as occurs in cases of FTR.

One might worry, however, that employing the one-off heuristic would require us to generate implausibly many beliefs. It might seem to require that every time we have an experience we form a belief that we have done so. However, the one-off heuristic's explanatory value requires only that, for many of the events that we could at one time remember but now cannot remember, there was a related one-off belief. Because most experiences are never stored in long term memory, there is no need to explain how they are forgotten by appeal to a lost one-off belief. Thus, the one-off heuristic is compatible with our not generating related one-off beliefs for most of our experiences.

However, even if the one-off heuristic were a generally reliable procedure for representing events in one's past, we have yet to see when (and why) following it should be viewed as a way of (genuinely, successfully) *remembering.* It has simply been proposed that our

---

[16] This parenthetical condition may be rejected by the simulationist who denies that successful remembering requires appropriate causation. The point of including it here is to capture a commitment that seems to guide attempts at remembering in ordinary individuals not steeped in the philosophy or science of memory.

sense that we are remembering an event, in some contexts, could depend upon our ability to follow the one-off heuristic. On the hypothesis that we make use of the one-off heuristic to remember specific events, we can explain why I seem to remember making saag paneer, why I don't think I can remember the first time I made steel-cut oatmeal, and why I probably would not think I could remember the first time I made saag paneer had I made it many times since. Yet the metaphysical question of what it is to successfully remember an event has not been meaningfully connected to this explanation. With that in mind, we can now return to the causal theory to see whether its requirements for successful remembering can be aligned with the one-off heuristic's role as an epistemic practice. This will be an important step toward the final goal of defending a causal theory that is consistent with the promiscuity of memory traces.

VII.

Recall Martin and Deutscher's principle 3—the causal condition—on successful remembering:

3) *S's* past experience of *e* was causally operative in producing a state or successive states in *S* finally operative in producing *r*.

Earlier we saw that, even in combination with Martin and Deutscher's principles (1) and (2), (3) fails to articulate sufficiency criteria for remembering. This was shown by cases of relearning and other (apparently) deviant causal chains (such as non-memorial retention through "suggestive states," discussed below), where S's past experience of *e* is causally operative in bringing about the state that is finally operative in producing *r*, but not in the right way. Our immediate reaction to cases of relearning and other would-be deviant causal chains suggests that we (think we) know *something* about how memory works that excludes them as cases of successful remembering. It is in the attempt to render that putative knowledge explicit, in 3b and 3c, that Martin and Deutscher arguably go too far, coming into conflict with the contemporary cognitive science of memory—both in its rejection of memory traces as "structural analogues" of what they represent and in its more recent questioning of the idea that traces are monogamous.

A common reaction by those still intent to defend some version of a causal theory is to replace 3b and 3c with a more general, if less informative, version of the causal condition—one that simply invokes "appropriate" causal relationships, without trying to specify what qualifies

causation as appropriate (Debus, 2016). The assumption in so doing is that, once we come to understand what appropriate causation consists in, we will understand why it is that intuitively deviant causal chains, such as are involved in relearning, are not instances of appropriate causation. Arguably, a causal theory that invokes "appropriate" causation, without specifying the nature of appropriate causation, is more of a research program than it is a theory. Let us call this theory-cum-research-program the *Generic Causal Theory* and understand it as follows:

> *Generic Causal Theory (GCT)*: *S* successfully episodically remembers event *e,* through the use of mental state *r,* iff:

> (1') *S*, within certain limits of accuracy, represents *e* through the use of *r*.

> (2') *S* observed *e.*

> (3') *S's* past observation of *e* is appropriately causally related to *r.*

In what follows I will take GCT as a working hypothesis about what is required for successful remembering—one that leaves open the crucial question of what in fact constitutes "appropriate" causation. It is, to be sure, a question-begging hypothesis, by the lights of contemporary simulationists. The point in advancing it is to see whether there are *new* ways to defend it— including new ways to articulate what constitutes *appropriate* causation—in light of the preceding discussion of memory traces and forgetting through repetition.

Importantly, the Generic Causal Theory does not require that there is just *one* sort of appropriate causal relationship between an event and its remembering. For all we know *a priori,* there may be several kinds of appropriate causal relationship. I will go on to suggest two. These will only be *suggestions*, however: "how-possibly" stories that articulate new forms a causal theory can take while remaining consistent with the promiscuity of memory traces. Yet my aim will be to describe processes that are not just possible but also *plausible,* due to their having the following three features: 1) consistency with trace promiscuity; 2) an ability to explain forgetting through repetition; and 3) an ability to rule out of relearning (and other intuitively deviant chains) as instances of successful remembering.

VIII.

The one-off heuristic is the following rule of thumb: when trying to remember some event *e*, search for two things together—a set of stored props *P* that seem appropriately content-related to *e*, and a belief that you have observed an event of *e's* kind exactly once. When these are found together, use *P* to generate a representation *r* of *e*. Past causal theories have placed the burden of appropriately causally linking *r* to *e* on a persisting, monogamous memory trace. The one-off heuristic suggests a different possibility: it may be the one-off *belief* that appropriately causally links *r* to *e*. Specifically, the belief that I have only once experienced an event of *e's* kind may be caused by the observation of *e* and maintained until it is operative in generating *r*, via the one-off heuristic. Understanding this as one way for *r*'s causation to be "appropriate," we can express the proposal more formally as:

> *Appropriate Relation One* (AR1): S's observation of *e* produces a belief in *S* that *S* has experienced an *e*-kind of event[17] only once. S also has stored props *P* that can accurately enough represent *e*. (These props may or may not have been generated prior to *e*.)
>
> The aforementioned belief and props are retained by S from the time of *e* until such time as they are operative in producing *r*, a representation of *e*.

While the principle is long, the idea behind it is simple. When we normally follow the one-off heuristic—in the way associated with genuine remembering—all of the following are true: 1) experiencing an event *e* has caused us to believe that we have witnessed such an event only once, 2) we have suitable traces for representing such an event, and 3) we have retained the mentioned belief and traces from the time of *e* until such time as they are operative in generating *r*, a remembering of *e*. Importantly, the notion of memory trace employed here is that of the prop theory. There is no assumption that the traces employed in the remembering of *e* were produced in the observation of *e*. All that is required is that S has memory traces that accurately enough represent *e* and that such traces have been retained at least since the occurrence of *e*.[18] Within

---

[17] It an open question how finely-grained events are to be typed here, but the working assumption is that they will be very finely grained. For instance, an "e kind of event" will not simply be *cooking dinner,* but, rather, *cooking thus and such type of meal in thus and such setting*, where such contents are represented by mental imagery. In general, the spatial outlay of properties and objects in a represented environment will be the features most highly relevant to individuating the type of event in question.

[18] Why not allow that such traces/props can be first acquired *after e*? The only reason I see for not doing so is that this would create the possibility of there being times after *e* during which one could not remember *e,* despite being able to remember *e* after those times. This would be counterintuitive. Nevertheless, on the view I am suggesting,

Appropriate Relation 1, it is the one-off belief—produced in the observation of *e*—that bears the weight of providing a discrete causal chain leading from *r* back to event *e*.

So, for example, when I remember making saag paneer, the conditions required for engaging in AR1 are met: making saag paneer has caused me to believe that I have done so once; I have props suitable for representing my making of saag paneer; I have retained this belief and these props at least since the time I made saag paneer; and, they are now operative in my producing *r,* a representation of that event. It is in virtue of the one-off belief's causal role in the process of remembering that the counterfactual important to all causal theories holds: had I not observed myself making saag paneer, I would not now seem to remember having made it. For without the observation of *e*, there would be no one-off belief; and, without the one-off belief, there would be no remembering. Thus, when the conditions in AR1 are met, we have both followed the one-off heuristic *and* satisfied the requirements of Generic Causal Theory.[19]

In response, it may be objected that we do not yet know whether the kind of causation invoked in AR1 is indeed *appropriate* causation. To my mind, the final verdict on whether such causation is appropriate will hang on the empirical question of whether humans very often make use of the one-off heuristic in what are otherwise considered uncontroversial cases of remembering an event.[20] Yet there is preliminary support for the idea that it constitutes appropriate causation in the fact that AR1 avoids counting as appropriate the intuitively deviant causal chains that have traditionally plagued causal theories. Consider first relearning. Following the one-off heuristic when relearning has occurred is not an instance of AR1 because the relevant one-off belief and props are not *retained by S* from the time of *e* until such time as they are operative in producing *r*. Instead, they are generated anew based on testimony. It is true that AR1 will require an account of what it is to retain a belief or a prop across multiple

---

such remembering would not be different in any deep sense than what I am prepared to count as successful remembering.

[19] It is true, as a referee notes, that AR1 does not explain how one knows that one's one-off belief was caused by the remembered event, and so does not explain how one knows that one is remembering (and not simply seeming to remember). There are things to be said here, appealing to context and other background beliefs. For now, I simply note that the traditional causal theory—invoking monogamous memory traces—likewise does not explain how one knows whether one is in fact remembering, or only seeming to remember. Nor does Michaelian's simulation theory explain this metacognitive feat in its basic account of what it is to remember (2016c, Ch. 6).

[20] It is not question-begging to assume the existence of some set of episodes that, pre-theoretically, all sides agree are cases of successful remembering (while perhaps disagreeing about other episodes). For episodic remembering to be a possible subject of scientific study, there must be some such set that constitutes paradigmatic cases of the phenomenon that all sides wish to explain.

days; but these are not special difficulties. It is commonly agreed that humans retain beliefs from day to day, independent of any puzzles about remembering. Likewise, it is an assumption of all contemporary accounts of memory that a memory trace can persist from day to day. Such a view only becomes controversial when combined with the view that traces are monogamous, which is not required by AR1.

Now consider another sort of deviant chain: non-memorial retention. These are situations where the remembering subject retains a state caused by observing *e* and where that state is later operative in producing an apparent remembering of *e*—yet where it is not the right sort of retained state to count as the subject's remembering *e*. Martin and Deutscher give the example of someone who has an experience of *e* that creates in him a certain "suggestible" (quasi-hypnotic) state *t* which, in turn, causes him to believe whatever he is told about his past. Someone later tells this person, in detail, about the very event *e* that caused in him the suggestible state *t*. Because *e* has caused him to (continue to be) in a state *t,* which causes him to believe whatever he is told about his past, he accepts the testimony and forms an accurate representation *r* of event *e*. Here state *t* is both produced by observing *e* and is retained by the subject until it later causes him to form the accurate representation *r* of *e* (albeit with considerable "prompting")—presumably without his remembering *e*. It was in response to this sort of challenge that Martin & Deutcher stipulated that the retained state must be a "structural analog" of the event remembered. AR1 solves the problem instead by specifying that the retained states are beliefs and prop-like memory traces, while omitting any controversial claims about the representational format or neural realization of those states. AR1 is compatible with a wide variety of accounts of the nature of beliefs and memory traces. All it assumes is that there will be some fact of the matter concerning whether someone has retained a certain belief and certain props for some period, and that a belief may be caused by an observation of a certain event.[21]

---

[21] As a referee observes, it is possible that a particular one-off belief would be promiscuous in being caused by experiences of multiple distinct events—if, for instance, one unwittingly formed the belief "I have experienced this sort of event only once" in response to two distinct experiences. Such a belief could not feature in AR1. Yet we have reason to think that promiscuity among one-off beliefs will be uncommon, as their contents are always single-event-specific. Thus, when such a belief is promiscuous, something has gone wrong, and we should expect related errors in remembering. By contrast, it is in the nature of traces (on the prop theory) to represent general, repeatable outlays of properties, and to be strengthened by repeated perceptions of such outlays.

A second sort of worry one may have about AR1 is that—my denials notwithstanding—it still understands appropriate causation in terms of a persisting (monogamous) memory trace, and that it therefore does not allow us to see how a causal theory could be true even if traces are promiscuous. It is worth clarifying why AR1, as I have formulated it, does not rely upon there being a monogamous relationship between the prop-like traces used in an act of remembering and the event remembered. It simply states that I must have had props of the right kind since the time of the event itself, which includes the possibility that the props were created *before* the event (see also fn. 18). Consider again the example of my remembering making saag paneer. On reflection, it seems possible that the props I use in remembering making that event were, in fact, encoded and stored prior to my making saag paneer. After all, many of them are simply props of my kitchen, formed through countless observations prior to my ever making saag paneer. Likewise, the incorporated props of pots and pans could easily have been first encoded long before that one cooking episode. Even the props of spinach as it looks with the paneer and spices mixed could (for all I know) have been first encoded and stored prior to my making saag paneer, when I observed my wife making it several times before I ever made the attempt. So, when I reflect on the matter, I really cannot say for certain whether any of the props I used to reconstruct my cooking of saag paneer had their genesis—or were instead merely strengthened—during that event. Indeed, it now feels to me somewhat improbable that any of the props I use to construct a representation of the event my cooking of saag paneer truly do derive *solely* from that one event.[22] Perhaps my intuitive sense of that is wrong and the props do indeed solely derive from my own cooking episode. The important point is that this needn't be the case. Their promiscuity—which seems plausible enough—would be consistent with my successfully remembering the event, in accordance with GCT and AR1.

Stepping back, the picture so far sketched is this: we have, in GCT, a schematic account of the necessary and sufficient conditions for remembering that is silent on the question of how to understand appropriate causation. Appropriate Relation One (AR1), inspired by the one-off heuristic, is a proposal for one form of appropriate causation, supported in its ability to explain forgetting through repetition, its consistency with trace promiscuity, and its not including some

---

[22] As I am conceiving of things, the set of props that are used, combined with the way they are arranged, determines the "type" of event that is being remembered. This makes it possible for me to only remember one episode of the type *I am making saag paneer,* even if the many of the same of props, arranged differently, could be used in the remembering of a different type of event (e.g. *my wife is making saag paneer*).

famously problematic (intuitively) deviant causal chains among its instances. However, AR1 has a significant limitation: it requires that the remembering subject believes she has experienced the sort of event remembered only once. If satisfying AR1 were the only way for a remembering to be appropriately causally related to the event remembered, then (according to GCT) we could not successfully remember any event that is of a kind that we think we have experienced multiple times. I think there are cases where we can indeed remember such events. The last order of business, then, is to describe a (plausible) second form of appropriate causation that is consistent with this ability. Because it is also consistent with the promiscuity of traces, it will help to bolster the general case that the Generic Causal Theory can hold even if traces are promiscuous.

IX.

We have seen that one way a causal theory of memory can be true while memory traces are as the prop theory sees them is for some state *other than* a memory trace (namely, a persisting one-off belief) to causally link a remembering to the event remembered. With that in mind, we should consider the thesis—popular in the contemporary scientific study of memory— that episodic remembering involves distinct, interacting components of *recollection* and *feelings of familiarity* (Bastin et al., 2019; Jacoby, Kelley, & Dywan, 1989; Whittlesea, 1993; Whittlesea & Williams, 2000; Yonelinas, 2002). Recollection is viewed as a process by which one generates detailed qualitative and spatial representations of past events (Bastin et al., 2019; Yonelinas, Aly, Wang, & Koen, 2010). In present terms, we can think of recollection loosely as the use of a memory trace (of the prop theory variety) in the representation of a past event. A feeling of familiarity, by contrast, is a "feeling of oldness indicating that something has been previously experienced," and is not itself conceived of as bearing information about the object or event remembered—other than, perhaps, that it has been encountered previously (Bastin et al., 2019, p. 2). Despite its informational impoverishment, the feeling of familiarity has been called the "*sin qua non* of remembering"—that without which a mental process of representing the past seems only to be a bit of guessing or problem solving (Whittlesea, 1993, p. 1235). The idea that these two processes are distinct yet equally involved in ordinary acts of episodic remembering is now supported by multiple lines of research, including dissociations revealed in behavioral tasks (Whittlesea & Williams, 2000), dissociations in neuropsychological cases (Aggleton et al., 2000; Bastin et al., 2019), and neuroimaging indicating distinct neural substrates for each (Yonelinas,

2002). The process of recollection—which draws on memory traces—is held to rely on the hippocampus and prefrontal cortex, while familiarity is thought to be modulated by areas surrounding the hippocampus (Aggleton & Brown, 1999; Maguire & Mullally, 2013; Moscovitch & Nadel, 2019).

While there are different reasons that a stimulus may generate a feeling of familiarity, the most commonly cited is the *fluency* with which a stimulus is processed, where fluency is the speed and ease with which the stimulus is processed. Fluency can be increased by a variety of factors, including the clarity of a stimulus, its similarity to other recently perceived stimuli, or whether the stimulus itself has been encountered recently. In what follows, I want to show how fluency can provide a necessary second ingredient to memory traces in a way that constitutes a second kind of appropriate causal relation (despite the promiscuity of traces).[23] First, a little more background on fluency and its relation to feelings of familiarity.

Building on the work of Jacoby et al. (1989), Whittlesea and colleagues (Whittlesea, 1993; Whittlesea, Jacoby, & Girard, 1990; Whittlesea & Williams, 2000) have shown that feelings of familiarity can be boosted by increasing the fluency of a stimulus, even in cases where the stimulus is not in fact familiar. In one of their experiments (numerous variations of which have been conducted over the years) participants were briefly shown a list of words followed by a target word. They were required to utter the target word aloud and indicate whether it was on the previously viewed list. The target word was always displayed with varying degrees of dynamic noise, making it easier or more difficult to decipher as a function of the noise level. (However, participants were unaware that the degree of noise was adjusted on each trial.) On trials with a lower degree of noise, participants identified and uttered the word faster, indicating greater fluency in processing the stimulus when noise was lower. However, on these higher fluency trials, participants were also more likely to falsely claim that the stimulus was repeated from the earlier viewed list. The explanation Whittlesea proposed, now widely endorsed (Bastin et al., 2019), is that stimulus fluency feeds into a kind of subconscious

---

[23] For another recent philosophical account of remembering that makes appeal to fluency and familiarity, see Perrin (2018). While Perrin labels his a "procedural causal" account of remembering, his approach differs from that developed here in not preserving a discrete causal pathway from each act of remembering to the event remembered, and in not articulating an essential role for memory traces.

inference concerning the source of the fluency. In the absence of evidence indicating other likely sources, such fluency is presumed to result from recent past perception of the stimulus.

"Interpreting fluent performance as an effect of prior experience is a practical heuristic," Whittlesea explains, "because one of the primary effects of experiencing an object is to facilitate later interactions with the object….When the past appears to be a likely source of current ease of processing, a feeling of familiarity will emerge" (1993, p. 1235, p. 1248). Whittlesea (1993) showed that when participants are made aware that perceptual clarity is being manipulated, and, therefore, that fluency may derive from something other than recent experience with the stimulus, they are no longer likely to falsely claim that relatively clear stimuli were previously viewed. Nevertheless, Whittlesea notes, "use of this heuristic leaves one open to the possibility of illusions of familiarity" of the sort elicited in his experiments (Whittlesea & Williams, 2001a, 2001b).

In short, the feeling of familiarity generated by fluency, and the episodic remembering that depends upon it, is the result of a sophisticated yet swift and unconscious inferential process, involving appraisals of how likely it is that the fluency could result from some source other than recent perception of the stimulus. Within this paradigm, processing fluency for a stimulus results from physiological changes in the brain that themselves occur due the recent perception of the stimulus; this fluency is *perishable*, lasting only for a limited time period (more on this in a moment).

Let *F* be a fluency-for-stimulus-*s* state of the brain, caused by a recent perception of stimulus *s*. *F* is to be distinguished from any memory trace *T* that might also be caused by the perception of *s*. Let *C* be contextual beliefs or appraisals on the basis of which *F* is (subconsciously) deemed to result from a recent perception of *s*. Remembering the event of observing *s* may involve following the *fluency heuristic*:

> *Fluency heuristic*: *S's* recent experience of *e* causes *F,* a perishable state of fluency, in *S* with respect to representations of *e-kinds of events*; S has traces *T* suitable for representing *e*-kinds-of events. *S* has contextual beliefs or appraisals *C* on the basis of which *F* is (subconsciously) deemed to result from a recent perception of *e*. On the basis of *T, F* and *C, S* generates *r,* an episodic remembering of *e*.

An important assumption behind this heuristic is that the memory traces (T) retrieved in the attempt to remember *e* are relevantly similar to a recent past perceptual experience one has had of *e*. This similarity—in terms of content, format, and underlying neural regions—is why fluency is experienced with respect to the act of remembering. There is ample experimental evidence for the reuse of representational mechanisms and neural areas in the perception and imagistic representation of like stimuli (Anderson, 2010; Nanay, 2018; Pearson & Kosslyn, 2015; Stanley, Gessell, & De Brigard, 2019).

To see how the fluency heuristic may be at work in a successful case of episodic memory—one consistent with the promiscuity of traces—consider the following example.

> *Locking the door*: Jonah locks the door to his apartment and walks down the hallway to the elevator. Suddenly unsure of whether he locked the door, he tries to remember doing so. Because he has locked the door countless times before, and witnessed himself doing it, he has stored props available that accurately represent him locking the door. He had these props/traces well in advance of locking the door just now. However, because he has recently seen himself locking the door, he temporarily experiences fluency in retrieving the relevant props, as the props themselves are similar to the kind of experience he had in seeing himself lock the door. This fluency generates a feeling of familiarity, which, combined with the availability of suitable props, leads Jonah to episodically remember—with the help of the relevant props—locking the door.

We have here an example of successful remembering that is consistent both with the Generic Causal Theory and the promiscuity of memory traces. In *Locking the Door*, Jonah explicitly makes use of a memory trace that was not itself produced in a perception of the event it enables him to remember. However, the causal condition (3)—that *S's* past observation of *e* is appropriately causally related to *r*—is satisfied, because F, the (perishable) state of having fluency with respect to perceptual representations of his backdoor being locked, was caused by his recent experience of *e* and is itself operative in producing *r*. Because Jonah's fluency with respect to perceptual and imagistic representations of his door being locked causally depends upon his recent observation of himself locking the door, and his act of remembering locking the door itself depends upon the familiarity generated by such fluency, we can say that, had Jonah

not observed himself locking the door, he would not now seem to remember doing so. This is the causal dependency of *r* on an observation of *e* that 3 requires.

To head off a reasonable objection[24], it is important to emphasize that the notion of fluency in play here—and the related feeling of familiarity it engenders—is a highly time-sensitive, "perishable" one, typically engendered by a recent perception of a stimulus that is similar to the one that is currently being perceived or remembered. There are other, less time-sensitive, kinds of fluency discussed in the literature, as manifest when familiar faces "pop out" of a group of unfamiliar faces, or when common English words are easily identified on a list including non-words. In those cases, the familiar stimuli (known faces, or genuine words) are processed more quickly and easily than the non-familiar stimuli, even if one hasn't very recently seen the presented faces or words. However, there is another sense of fluency and familiarity invoked in the studies cited above that is highly time sensitive and perishable. In the kinds of studies carried out by Jacoby et al. (1989) and others (see, e.g., Besson et al. (2015) and Westerman et al. (2002)), participants first study a list of commonly used words, then later are asked which in a list of subsequently presented words were on the first list. By manipulating fluency for the stimuli (e.g. by showing masked primes of those words), experiments show that participants rely on fluency in making such determinations. But note that, in one sense, *all* of the words presented during the recall phase will be familiar, and will be processed with fluency, just because they are common English words that the participant has seen countless times before. This long-lasting form of fluency is not the kind being manipulated in the experiment. Instead, there appears to be a distinct, time-sensitive, form of fluency that one has for a stimulus only if it has been relatively recently perceived. (This is the only kind of fluency it would make sense to rely upon during such tasks, after all.) It must give rise, in turn, to a correspondingly distinct feeling of familiarity—one grounded not simply in the word's being a common word, but also in the fact that it was recently perceived. So, in proposing that Jonah relies upon fluency, and a related feeling of familiarity, I have in mind these time-sensitive, perishable notions of fluency and familiarity. These are the only kinds that can aid him in determining whether he has relatively recently locked the door.

---

[24] I'm indebted to two anonymous reviewers for this journal for raising this issue.

The final step is to assure ourselves that this sort of causation-by-time-sensitive-fluency is indeed appropriate (and non-deviant). The case should raise no eyebrows, as it is already well-established that fluency and feelings of familiarity play a pivotal role in remembering. It is reasonable to accept, as a working hypothesis, that the sort of fluency-based heuristic exploited by Jonah above is used commonly enough in paradigmatic instances of remembering—such as are studied in experimental contexts—to be considered one of the appropriate ways for an apparent remembering to be related to the past event it represents. We can formalize this idea as a second kind of appropriate causation relation:

> *Appropriate Relation Two (AR2)*: S's recent experience of *e* causes perishable fluency *F* in *S* with respect to representations of *e*-kinds of events; S has traces *T* that accurately enough represent *e*-kinds of events. *S* also has contextual beliefs or appraisals *C* on the basis of which *F* is (subconsciously) deemed to result from a recent prior experience of *e*. F, C and T are retained by *S* until such time as they are operative in producing *r,* an episodic remembering of *e*.

Here we apply the same strategy of relying on a mental state—or combination of states—*other than* a memory trace to provide the discrete causal chain leading from the perception of an event to its remembering. AR2 contrasts with AR1 in being less reliant on personal-level (i.e., easily reportable) states like beliefs and in incorporating more perishable, subconscious states. Yet AR2 leaves it equally clear how memory traces may be promiscuous. Jonah might make use of the very same props a week later when, again worried that he absentmindedly left the door unlocked, he triggers them in a successful remembering of locking his door—one causally dependent on his having locked it, due to a newly generated (and again perishable) state of fluency *F* that is operative in his remembering.

The present proposal concerning the role of fluency and feelings of familiarity in remembering remains speculative. But it is plausible enough—and sufficiently inspired by themes in contemporary memory science—to serve our purposes. Our aim has been to provide a clear and empirically supported account of how it could be that the causal theory of memory is consistent with the promiscuity of traces—a *how possibly* story, bordering on how-plausibly. If the two appropriate relations proposed here —AR1 and AR2—are wrong in their details, they

nevertheless paint a picture of how other, similar accounts might succeed in their place. On the view I am proposing, satisfaction of *either* AR1 or AR2 is sufficient for an apparent remembering to be appropriately causally related to the event it represents. Neither is a necessary condition for appropriate causation. In addition—and importantly—we should expect to discover other conditions (AR3, AR4…) that also constitute appropriate causation, because they describe the kinds of causally dependent heuristics that are replied upon by humans in many ordinary cases of remembering.

For now, we have in AR1 and AR2 the general shape of an account of how we remember that is consistent both with a causal theory's strictures on successful remembering and the promiscuity of traces. The key strategy is to conceive of episodic remembering not as the simple retrieval and projection of a static memory trace, but as a complex quasi-inferential process that makes use of multiple forms of information and cues—prop theory traces included—in generating the experience we intuitively identify as episodic remembering. When remembering is understood as a multi-componential process, there is a disjunction of "appropriate" ways in which a putative remembering may be causally dependent upon a prior experience of the event remembered. This serves to satisfy our philosophical sense that, when we successfully remember an event, we would not have done so had we not observed the event.

Works Cited

Addis, Donna R. (2020). Mental Time Travel? A Neurocognitive Model of Event Simulation. *Review of Philosophy and Psychology, 11*(2), 233-259. doi:10.1007/s13164-020-00470-0

Addis, Donna R., Wong, Alana T., & Schacter, Daniel L. (2007). Remembering the past and imagining the future: common and distinct neural substrates during event construction and elaboration. *Neuropsychologia, 45*(7), 1363-1377.

Aggleton, J. P., & Brown, Malcom W. (1999). Episodic memory, amnesia, and the hippocampal-anterior thalamic axis. *Behav Brain Sci, 22*(3), 425-444; discussion 444-489.

Aggleton, J. P., McMackin, D., Carpenter, K., Hornak, J., Kapur, N., Halpin, S., . . . Gaffan, D. (2000). Differential cognitive effects of colloid cysts in the third ventricle that spare or compromise the fornix. *Brain, 123*(4), 800-815. doi:10.1093/brain/123.4.800

Anderson, Michael L. (2010). Neural reuse: A fundamental organizational principle of the brain. *Behavioral and Brain Sciences, 33*(4), 245-266. doi:10.1017/S0140525X10000853

Barry, Daniel. N., & Maguire, Eleanor. A. (2019). Remote Memory and the Hippocampus: A Constructive Critique. *Trends Cogn Sci, 23*(2), 128-142. doi:10.1016/j.tics.2018.11.005

Bastin, Christine, Besson, Gabriel, Simon, Jessica, Delhaye, Emma, Geurten, Marie, Willems, Sylvie, & Salmon, Eric. (2019). An integrative memory model of recollection and familiarity to understand memory deficits. *Behav Brain Sci, 42*, e281. doi:10.1017/s0140525x19000621

Bernecker, Sven. (2010). *Memory: A Philosophical Study*. Oxford: Oxford University Press.

Besson, Gabriel, Ceccaldi, Mathieu, Tramoni, Eve, Felician, Olivier, Didic, Mira, & Barbeau, Emmanuel J. (2015). Fast, but not slow, familiarity is preserved in patients with amnestic mild cognitive impairment. *Cortex, 65*, 36-49. doi:https://doi.org/10.1016/j.cortex.2014.10.020

Cheng, Sen, & Werning, Markus. (2016). What is episodic memory if it is a natural kind? *Synthese, 193*(5), 1345-1385. doi:10.1007/s11229-014-0628-6

De Brigard, Felipe. (2014a). Is memory for remembering? Recollection as a form of episodic hypothetical thinking. *Synthese, 191*(2), 155-185. doi:10.1007/s11229-013-0247-7

De Brigard, Felipe. (2014b). The Nature of Memory Traces. *Philosophy Compass, 9*(6), 402-414. doi:10.1111/phc3.12133

Debus, Dorothea. (2017). Memory causation. In S. Bernecker & K. Michaelian (Eds.), *The Routledge handbook of philosophy of memory* (pp. 63-75). London: Routledge.

Dudai, Yadin. (2012). The Restless Engram: Consolidations Never End. *Annual Review of Neuroscience, 35*(1), 227-247. doi:10.1146/annurev-neuro-062111-150500

Fernández, Jordi. (2019). *Memory: A self-referential account*. New York: Oxford University Press.

Hutto, Daniel D., & Peeters, Anco. (2018). The roots of remembering: Radically enactive recollecting. In Kourken Michaelian, Dorothea Debus, & Denis Perrin (Eds.), *New Directions in the Philosophy of Memory* (pp. 97-118): Routledge.

Jacoby, Lawrence L, Kelley, Colleen M, & Dywan, Jane. (1989). Memory attributions. In H.L. Roediger & F.I.M. Craik (Eds.), *Varieties of memory and consciousness: Essays in honour of Endel Tulving* (pp. 391-422): Erlbaum.

Langland-Hassan, P. (forthcoming). Remembering, Imagining, and Memory Traces: Toward a Continuist Causal Theory. In C.J. McCarroll, K. Michaelian, & André Sant'Anna (Eds.), *Current Controversies in the Philosophy of Memory*. London: Routledge.

Maguire, Eleanor A, & Mullally, Sinéad L. (2013). The hippocampus: a manifesto for change. *Journal of Experimental Psychology: General, 142*(4), 1180.

Martin, C.B., & Deutscher, Max. (1966). Remembering. *The Philosophical Review, 75*(2), 161-196. doi:10.2307/2183082

Michaelian, Kourken. (2011). Generative memory. *Philosophical Psychology, 24*(3), 323-342. doi:10.1080/09515089.2011.559623

Michaelian, Kourken. (2016a). Against Discontinuism. In K. Michaelian, S.B. Klein, & K.K. Szpunar (Eds.), *Seeing the Future: Theoretical Perspectives on Future-Oriented Mental Time Travel*. Oxford: Oxford University Press.

Michaelian, Kourken. (2016b). Confabulating, misremembering, relearning: The simulation theory of memory and unsuccessful remembering. *Frontiers in Psychology, 7*, 1857.

Michaelian, Kourken. (2016c). *Mental Time Travel: Episodic Memory and Our Knowledge of the Personal Past*. Cambridge, MA: MIT Press.

Michaelian, Kourken. (2021). Imagining the past reliably and unreliably: towards a virtue theory of memory. *Synthese*. doi:10.1007/s11229-021-03125-4

Michaelian, Kourken, & Robins, Sarah. (2018). Beyond the causal theory? Fifty years after Martin and Deutscher. In K. Michaelian, Dorothea Debus, & D. Perrin (Eds.), *New Directions in the Philosophy of Memory* (pp. 13-32). London: Routledge.

Michaelian, Kourken, & Sant'Anna, André. (2019). Memory without content? Radical enactivism and (post)causal theories of memory. *Synthese*. doi:10.1007/s11229-019-02119-7

Moscovitch, M., & Nadel, Lynn. (2019). Sculpting Remote Memory: Enduring Hippocampal Traces and vmPFC Reconstructive Processes. *Trends in cognitive sciences, 23*(8), 634-635. doi:10.1016/j.tics.2019.05.001

Moscovitch, M., Rosenbaum, R. S., Gilboa, A., Addis, D. R., Westmacott, R., Grady, C., . . . Nadel, L. (2005). Functional neuroanatomy of remote episodic, semantic and spatial memory: a unified account based on multiple trace theory. *J Anat, 207*(1), 35-66. doi:JOA421 [pii]

10.1111/j.1469-7580.2005.00421.x

Nadel, Lynn. (2007). Consolidation: the demise of the fixed trace. In H.L. Roediger, Y. Dudai, & S.M. Fitzpatrick (Eds.), *Science of Memory Concepts* (pp. 177-182). Oxford: Oxford University Press.

Najenson, Jonathan. (2021). What have we learned about the engram? *Synthese*. doi:10.1007/s11229-021-03216-2

Nanay, Bence. (2018). Multimodal mental imagery. *Cortex, 105*, 125-134.

Pearson, Joel, & Kosslyn, Stephen M. (2015). The heterogeneity of mental representation: Ending the imagery debate. *Proceedings of the National Academy of Sciences, 112*(33), 10089-10092. doi:10.1073/pnas.1504933112

Perrin, Denis. (2018). A case for procedural causality in episodic recollection. In K. Michaelian, Dorothea Debus, & D. Perrin (Eds.), *New Directions in the Philosophy of Memory* (pp. 33-51). London: Routledge.

Robins, Sarah. (2016). Representing the past: memory traces and the causal theory of memory. *Philosophical Studies, 173*(11), 2993-3013. doi:10.1007/s11098-016-0647-x

Robins, Sarah. (2017). Memory traces. In K. Michaelian & S. Bernecker (Eds.), *The Routledge handbook of philosophy of memory* (pp. 76-87). London: Routledge.

Robins, Sarah. (2020). Fixed Engrams and Neural Dynamics. *Philosophy of science*.

Sara, Susan J. (2000). Retrieval and reconsolidation: toward a neurobiology of remembering. *Learning & Memory, 7*(2), 73-84.

Schacter, D. L., Addis, Donna R., & Buckner, Randy L. (2007). Remembering the past to imagine the future: the prospective brain. *Nat Rev Neurosci, 8*(9), 657-661. doi:nrn2213 [pii]

10.1038/nrn2213

Stanley, Matthew L., Gessell, Bryce, & De Brigard, Felipe. (2019). Network Modularity as a Foundation for Neural Reuse. *Philosophy of science, 86*(1), 23-46. doi:10.1086/701037

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

Tulving, Endel. (1983). *Elements of Episodic Memory*. Oxford: Oxford University Press.

Werning, Markus. (2020). Predicting the Past from Minimal Traces: Episodic Memory and its Distinction from Imagination and Preservation. *Review of Philosophy and Psychology*, 1-33.

Westerman, Deanne L., Lloyd, Marianne E., & Miller, Jeremy K. (2002). The attribution of perceptual fluency in recognition memory: the role of expectation. *Journal of Memory and Language, 47*(4), 607-617. doi:https://doi.org/10.1016/S0749-596X(02)00022-0

Whittlesea, Bruce W. A. (1993). Illusions of familiarity. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 19*(6), 1235-1253. doi:10.1037/0278-7393.19.6.1235

Whittlesea, Bruce W. A., Jacoby, Larry L., & Girard, Krista. (1990). Illusions of immediate memory: Evidence of an attributional basis for feelings of familiarity and perceptual quality. *Journal of Memory and Language, 29*(6), 716-732. doi:https://doi.org/10.1016/0749-596X(90)90045-2

Whittlesea, Bruce W. A., & Williams, Lisa D. (2000). The source of feelings of familiarity: The discrepancy-attribution hypothesis. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 26*(3), 547-565. doi:10.1037/0278-7393.26.3.547

Whittlesea, Bruce W. A., & Williams, Lisa D. (2001a). The discrepancy-attribution hypothesis: I. The heuristic basis of feelings and familiarity. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 27*(1), 3-13. doi:10.1037/0278-7393.27.1.3

Whittlesea, Bruce W. A., & Williams, Lisa D. (2001b). The discrepancy-attribution hypothesis: II. Expectation, uncertainty, surprise, and feelings of familiarity. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 27*(1), 14-33. doi:10.1037/0278-7393.27.1.14

Yonelinas, Andrew P. (2002). The Nature of Recollection and Familiarity: A Review of 30 Years of Research. *Journal of Memory and Language, 46*(3), 441-517. doi:https://doi.org/10.1006/jmla.2002.2864

Yonelinas, Andrew P., Aly, Mariam, Wang, Wei-Chun, & Koen, Joshua D. (2010). Recollection and familiarity: Examining controversial assumptions and new directions. *Hippocampus, 20*(11), 1178-1194. doi:10.1002/hipo.20864