

Penultimate draft. Forthcoming in *Canadian Journal of Philosophy*. Please do not cite without permission (skrobins@ku.edu)

Contiguity and the Causal Theory of Memory

Sarah K. Robins

skrobins@ku.edu

Department of Philosophy, University of Kansas, 3090 Wescoe Hall, 1445 Jayhawk Blvd., Lawrence, KS 66045, USA

Abstract

Bernecker (2010) argues for an account of memory causation in terms of contiguity. This *Contiguity View* is meant to solve *relearning* and *prompting*, wayward causation problems plaguing the causal theory of memory. I argue that Bernecker's Contiguity View fails in this task. Contiguity is too weak to prevent relearning and too strong to allow prompting. These failures illustrate a problem inherent in accounts of memory causation. Relearning and prompting are both causal relations, wayward only with respect to our interest in specifying remembering's requirements. Solving them requires saying more about remembering, not causation. I conclude by sketching such an account.

Keywords: Causal Theory of Memory; Causation; Spatio-Temporal Contiguity; Memory Traces; Remembering

Acknowledgements:

Many thanks to audiences at Wake Forest University and the University of Kansas for helpful feedback on previous drafts of this paper. Completion of this manuscript was supported by the University of Kansas General Research Fund allocation #2302343.

1. Introduction

The Causal Theory of Memory characterizes remembering as a causal process. The representation produced in remembering must have been caused by the previous experience being represented, and in the right way. Proponents of the causal theory thus devote much of their energy to articulating what this right way must be—i.e., providing an account of the nature of memory causation. Doing so requires ruling out a set of wayward causation cases, which allow an accurate representation of a past event to be produced without remembering. Most attention has been paid to two such cases: *relearning* and *strict prompting*. Relearning occurs when one reacquires forgotten information from another source. In cases of strict prompting, the ability to remember the past event is retained, but another source acts to produce the

Penultimate draft. Forthcoming in *Canadian Journal of Philosophy*. Please do not cite without permission (skrobins@ku.edu)

representation first, and this source contains all of the information that the trace would have provided in the prompt's absence.

Recently, Bernecker (2008; 2010) has proposed that the nature of memory causation should be characterized in terms of temporal contiguity, arguing that contiguity can solve these wayward causation problems.¹ Bernecker's *Contiguity View* appeals to a widely shared assumption about remembering, namely, that it requires memory traces because action at a distance is impossible. Learning and remembering are separated by a temporal gap, the argument goes, and so memory traces must form an uninterrupted causal chain between them, preventing appeal to action at a distance. This contiguity, Bernecker claims, explains the difference between remembering and relearning—relearning lacks the requisite contiguity. Further, so long as this contiguity is maintained, then remembering is compatible with strict prompting.

In this paper, I argue that Bernecker's Contiguity View fails to solve the problems of relearning and strict prompting. Contiguity does not uniquely characterize remembering. In the case of relearning, the contiguity requirement is too weak. Remembering and relearning are both connected (causally) to the previous event and so cannot be distinguished by appeal to contiguity alone. Contiguity is, on the other hand, too strong for cases of strict prompting: by requiring the trace to be the proximal cause of the representation in remembering, it rules out cases where the prompt acts preemptively. I conclude the discussion of each of these wayward causation cases with a gesture at what more must be said to sustain the causal theory of memory. Remembering is distinct from relearning because remembering is a causal relation to the past event that is

¹ At times, Bernecker refers to the required contiguity as *spatio-temporal* contiguity. Bernecker does not elaborate on the spatial component of this contiguity requirement, and so my argument here focuses on temporal contiguity. I understand temporal contiguity to be Bernecker's primary concern, as he states, "the stipulation of memory traces is motivated by the need to understand how a mental state can exert causal influence at a temporal distance" (2010, 130).

Penultimate draft. Forthcoming in *Canadian Journal of Philosophy*. Please do not cite without permission (skrobins@ku.edu)

retained within the same cognitive system in which it was first formed, as a memory trace that is available to the rememberer without re-perception. What matters for remembering is the activation of this capacity, regardless of when it becomes active relative to the representation. Remembering can occur even when the trace merely reinforces or sustains a representation brought about by other means. The prompt may act first so long as the trace still acts. In short, remembering is a causal process, but it is not merely a causal process. Remembering occurs when a person retains the capacity to represent information acquired from past events, in a way that allows the information to contribute—without re-perception—to future thought and action, either by producing or sustaining representations of those events.

2. Bernecker's Contiguity View

Bernecker's (2010) analysis of remembering has three key features. First, he defines remembering as a relation between two mental representations with propositional contents. He focuses on propositional memory exclusively, which he characterizes as including all mental representations where the "natural expression of its content involves a 'that' clause" (2010, 21). Second, Bernecker does not require the two mental representations in remembering to be had by the same subject.² If remembering requires the person learning and the person remembering to be one and the same, then memory would presuppose identity.³ Third, the contents of the two mental representations involved in remembering are not required to be identical; they are

² Exception is made for what Bernecker calls *introversive memories*, where the memory's content refers to one's own mental states (2010, 43).

³ Schechtman (2011) argues that Bernecker's argument is insufficient to block circularity objections to memory-based accounts of personal identity. Later in §3.1 argue that Bernecker should reject this condition (at least if he remains committed to solving the relearning problem), so nothing I say here turns on whether Bernecker's account can withstand Schechtman's objection.

Penultimate draft. Forthcoming in *Canadian Journal of Philosophy*. Please do not cite without permission (skrobins@ku.edu)

required only to be “sufficiently similar” (2010, 217). The required similarity is characterized in terms of entailment in relevance logic—the remembered representation must be something that could be inferred from the content of the earlier representation.⁴

With this analysis of remembering in hand, Bernecker goes on to claim that the relation between these two representations is best explained by assuming they are causally connected. He argues that non-causal accounts, such as evidentiary and simple retention views, fail to identify remembering uniquely. Bernecker claims that the causal view offers the best interpretation of the counterfactual that if one hadn’t represented a certain proposition in the past, then one would not represent a similar proposition at a later time. Together, the analysis of memory and the argument for remembering as a causal process yield the following condition on remembering:

Causal Condition: S’s representation at T2 that p is suitably causally connected to S’s representation at T1 that p* (2010, 128).

Bernecker recognizes that more must be said to explain what is required for this causal connection to be “suitable.” His primary concern is to avoid appeal to action at a distance. In order for the representation at T1 to cause the representation at T2, there must be some intermediary by which the causal process is sustained and the possibility of temporal gaps avoided. Bernecker appeals to memory traces as the requisite intermediaries:

Common sense speaks in favor of the stipulation of memory traces. For how else can past representations (or experiences) act at a temporal distance, if not through a continuous trace (or a series of traces)? (2008, 34)

Bernecker’s characterization of causation derives from his reading of Hume (1739), whom he interprets as imposing a contiguity requirement on all causal relations.⁵ He stops short

⁴ Others have argued that such “elliptical” (Malcolm 1963) or “impure” (Cusmariu 1980) memories should not be included in an analysis of remembering.

Penultimate draft. Forthcoming in *Canadian Journal of Philosophy*. Please do not cite without permission (skrobins@ku.edu)

of arguing that contiguity is a requirement on all causal connections. Instead, he offers the more modest proposal that “at least in the case of memory, contiguity is a necessary condition” (2010, 131). From the causal condition on remembering and the temporal contiguity requirement on memory causation, Bernecker derives the following trace condition on remembering:

Trace Condition: S’s representation at T1 that p^* and S’s representation at T2 that p are connected by a persisting memory trace or a contiguous set of memory traces (2010, 130).

Bernecker provides both a mental and physical description of the memory traces that sustain the contiguity between the representations at T1 and T2. Mentally, memory traces are either dispositional beliefs or subdoxastic states, which reflect the content retained from T1 until T2. Physically, traces supervene on intracerebral occurrences.⁶

Bernecker defends this Contiguity View by illustrating its ability to solve two wayward causation problems familiar to proponents of the Causal Theory of Memory. He identifies and addresses two, relearning and prompting:

Relearning: distinguishing the causal connection in remembering from the causal connection involved when one learns information, forgets it, and then reacquires it from another source.

Prompting: distinguishing cases where a prompt aids the trace in causing the representation from cases where a prompt causes the representation on its own, overriding or pre-empting the causal influence of the trace.

⁵ Bernecker acknowledges that his interpretation of Hume’s contiguity requirement is not shared universally. Others have defended the need for memory traces in similar fashion. Rosen (1975), for example, claims that the need for memory traces can be derived from a general, empirical principle of spatial and temporal contiguity. Bernecker also considers the principle to be an empirical one, citing cases of discovered contiguity as evidence in its favor.

⁶ The intracerebral requirement concerns only the initial formation of the memory trace. Later, in discussing certain kinds of memory transplant, Bernecker acknowledges traces can be transferred to other non-brain-based media. I discuss these concessions about traces and memory transplants in §3.1.

Penultimate draft. Forthcoming in *Canadian Journal of Philosophy*. Please do not cite without permission (skrobins@ku.edu)

In the next two sections, I describe Bernecker's proposed solutions to the relearning and prompting problems, respectively.

3. Contiguity and Relearning

Relearning occurs when someone acquires information, forgets it, and then learns it again. Say, for example, that Pete rides a rollercoaster and finds the experience terrifying, thinking to himself *rollercoasters are dangerous*. The experience was so traumatic, in fact, that Pete records the details of the event and the judgments he formed in his diary. A few years later Pete suffers an accident, the result of which is a traumatic brain injury and complete retrograde amnesia—he no longer remembers events from his past, including the rollercoaster ride and his previous journaling habit. At some point afterward, Pete finds his diary and re-reads the previous entry about the rollercoaster. *Rollercoasters are dangerous*, he thinks.

Such cases are meant to prod the intuition that relearning is a way of producing accurate representations of previously acquired information that is distinct from remembering. It is often assumed that the difference is best characterized by appeal to memory traces. Bernecker agrees, stating that in relearning, “the original memory traces are erased” (2010, 131). This claim will help to preclude certain cases of relearning—i.e., cases where the relearning derives from an independent source that also experienced the event in question. Suppose that Pete had ridden the rollercoaster alongside his friend Peggy, for example, and that his becoming reacquainted with the experience after his accident relied upon her account of their adventure. In such a case, the chain of events that leads to Pete's representation at T2 goes through Peggy's initial encounter—not Pete's. Remembering cannot occur because the route through Pete's memory traces is no longer available.

Penultimate draft. Forthcoming in *Canadian Journal of Philosophy*. Please do not cite without permission (skrobins@ku.edu)

There are, however, other cases of relearning that are not as easily handled by Bernecker's trace condition—cases, for example, like the one used at the opening of this section, where the opportunity for relearning derives from the would-be-rememberer's initial representation of the event. This form of relearning cannot be dismissed simply by saying that the original traces have been erased. Memory traces are fragile and finite, which is why Bernecker's trace condition allows for the contiguity of remembering to be sustained by a memory trace or *contiguous set of memory traces*. Loss of the original traces is not enough to rule out this form of relearning; the original traces are often erased in cases of remembering, too. Both remembering and this form of relearning are causal chains that begin with the would-be-rememberer's representation of the event in question—here, Pete's proclamation that *rollercoasters are dangerous* (p^* at T1). The causal chains end in the same place, too. In each case Pete believes, again, that *rollercoasters are dangerous* (p at T2).

Contiguity alone is not enough to tell the two cases apart. The relearning chain goes through a diary while the remembering chain does not, but both are causal chains. In other words, while it is true that Pete's relearning deviates from assumptions about what remembering requires, it is not *causally* deviant. It is wayward only with respect to our interest in saying what remembering requires. Both cases span a temporal distance and require something to sustain the causal relation between T1 and T2. So long as there are a series of states that do so—whether states of Pete or states of the diary—contiguity is maintained. The need for contiguity was used to justify the existence of memory traces. Contiguity cannot be appealed to, again, to do the additional work of sorting between ways of maintaining that contiguity over time.

Penultimate draft. Forthcoming in *Canadian Journal of Philosophy*. Please do not cite without permission (skrobins@ku.edu)

3.1 Enriching Contiguity

Can the contiguity requirement be enriched so as to preclude cases of relearning?

Bernecker elaborates on the nature of memory traces, which may provide his account with additional resources for sorting between these forms of contiguity. As noted in §2, Bernecker offers a description of traces. Traces supervene on intracerebral occurrences, supporting either dispositional beliefs or subdoxastic states.⁷ This characterization of traces may be right, but it alone will not help to distinguish memory traces from other mental states. Not all mental states are either beliefs or subdoxastic states, so this description distinguishes memory traces from some other mental states; still, there are many mental states that fit this description other than memory traces—the mental states that support relearning, for example.

Perhaps the trace condition can be saved by appealing to an apparent difference between remembering and relearning: in relearning, the information is located, at least temporarily, outside the rememberer. After the accident, the information about the rollercoaster is in the diary, not Pete. Thus one might suppose that cases of relearning can be excluded by stipulating that remembering involves *bodily* contiguity, so that remembering requires the trace to remain in physical contact with the rememberer. Proximity to the body is, however, incidental. Relearning can occur even when the information never leaves the subject's body. Pete could have tattooed the claim about rollercoasters on his arm instead of writing it in his diary, for instance.⁸ Further,

⁷ Bernecker elaborates on these intracerebral occurrences, too: they are “structural modifications at synapses (i.e. the area where the axon of one neuron connects with the dendrite of another neuron) that affect the ease with which neurons in a neural network can activate each other” (p. 132). My point above still holds: even if this is true of memory traces, it is not true of *only* memory traces. Many other mental states, including states of relearning, are likely to be supported by the same or similar neural processes.

⁸ Martin and Deutscher (1966), earlier proponents of the causal theory of memory, also deny the possibility of using a body requirement to characterize the remembering relation. They provide

Penultimate draft. Forthcoming in *Canadian Journal of Philosophy*. Please do not cite without permission (skrobins@ku.edu)

remembering need not involve physical contact between the memory and the body of the rememberer. It is conceivable that members of another species could store their memories externally, or even that advancements in neural prosthetics will someday allow us to rely on such external storage. Bernecker accepts both of these points, and so dismisses the possibility of using a body requirement to block cases of relearning. For identifying the causal process of remembering, he says, “the spatial location of the causal chain is inessential” (2010, 130). In other words, what matters is the *existence* of the contiguous causal chain, not its *location*. This allowance is fine for dealing with cases of relearning that derive from an outside source, as when Pete relearns of the rollercoaster ride from Peggy. In such a case, there are two distinct causal chains originating from the event at T1—one that comes to an end when Pete suffers his accident and another that culminates in relearning when Peggy testifies about the event to Pete. This allowance causes problems, however, in cases of relearning where the source of the relearning derives ultimately from the would-be rememberer’s original representation, as in the case of Pete and his diary.

Allowing external causal chains could reflect Bernecker’s endorsement of an Extended Mind Thesis, as proponents of this view argue that the mind often extends beyond the boundaries of the skin to include parts of the surrounding environment (e.g., Clark & Chalmers 1998; Clark 2008). If so, then we could look to the Extended Mind Theory for a principle of contiguity suitable for separating remembering and relearning. Bernecker does not address this possibility, but his subsequent discussion of memory transplants makes clear that his disinterest in the representation’s spatial location differs from that of the Extended Mind theorist. Bernecker does not want to stretch the mind’s circumference—he wants only to allow for the possibility that

the example of a student using a pen to inscribe answers to a chemistry exam into his palm (181-182).

Penultimate draft. Forthcoming in *Canadian Journal of Philosophy*. Please do not cite without permission (skrobins@ku.edu)

individual memory traces can persist when they are no longer connected to the body. So while the Extended Mind theorist and Bernecker would agree that some mental states could be located outside of the body, they diverge over their subsequent inferences regarding the relation between such states and the mind. For the Extended Mind theorist, if the representation is located outside of the person's body, then it follows that the person's mind is outside of her body. For Bernecker, if the representation is located outside of the body, then it is no longer part of the mind—it is now a free-floating memory trace eligible for transplant into any mind whatsoever.

Bernecker discusses the possibility of remembering in various memory transplant scenarios, and in so doing fleshes out further the kind of contiguity he envisions as characteristic of remembering. As I will show, even this enriched account of contiguity cannot identify remembering uniquely.

Since Bernecker's analysis of remembering does not require the rememberer to be identical to the person who had the original representation, he allows for traces to be removed from the brain, stored in a remote location, and then placed either into the brain of the original learner—an *intrapersonal* transplant—or into the brain of someone else— an *interpersonal* transplant. Even though transplants are possible, he requires memory traces to begin as states that supervene on neural states. He thus denies the possibility of memory transplants via *trace creation*, where traces are manufactured artificially and then implanted in a person's brain. If one could create a trace of riding a rollercoaster on the moon *ex nihilo*, for example, Pete could then receive the transplant but he would not be remembering. The worry is not simply that such "memories" could never be true, but rather that they would not bear the right relation to the world. Because created traces lack a causal connection to the events they purport to represent, Bernecker denies that created transplants support remembering. Once created via such an

Penultimate draft. Forthcoming in *Canadian Journal of Philosophy*. Please do not cite without permission (skrobins@ku.edu)

appropriate connection to the events they represent, however, traces are allowed to float free of the body in which they were formed and participate in remembering via transplant.

Bernecker even allows traces to float free of the initial medium of synaptic modification, as indicated by his claim that remembering is possible in cases of *trace replication*. Trace replication occurs when a trace is formed, and then at some later point the original trace medium is destroyed but the information carried by the trace is transferred to another host. Bernecker gives the example of downloading information from a trace to a computer and then erasing the trace. As he sees it, temporal contiguity is the critical feature of such cases: the replication must occur before or just as the old information is destroyed. In fact, the transfer would have to occur in this way in order to secure the causal relation. So long as temporal gaps are prevented, remembering via trace replication is possible. He offers the following as an example of the kind of transfer he intends to capture:

If trace replication involved in teletransportation consists in something like downloading information from a trace onto a computer, emailing a file to another computer, and copying the information into an “empty” trace, there is no good reason to deny that the causal chain is temporally and spatially contiguous (2010, 140).

It is curious that Bernecker would want to count such a case as remembering. At least, it is curious that he wants to include it *and* preclude relearning. The above example of trace replication appears at least directly analogous to the case of relearning involving Pete and his diary. In trace replication, a trace is formed, the information contained in the trace is copied onto another source, and the original trace is destroyed. This description fits relearning as well, at least the cases where the relearning derives from the person’s initial representation of the event and not from an independent source. In relearning, one acquires information, forgets it, and then learns it again. If the transfer of information in relearning does not violate temporal contiguity

Penultimate draft. Forthcoming in *Canadian Journal of Philosophy*. Please do not cite without permission (skrobins@ku.edu)

(i.e., if the information is transferred to the source from which the later relearning derives before the traces are damaged in forgetting), then—on Bernecker’s view—remembering should be possible. If Pete’s diary had been written on a personal computer, rather than with old-fashioned pen and paper, he could have relearned of his earlier thought about rollercoasters through the sort of replication process that Bernecker describes. Bernecker’s judgments about transplant cases further reveal the conflict between his stated aim of blocking relearning and the consequences of his reliance on temporal contiguity.

The conflict could be mitigated by disallowing cases of trace replication. But excluding this form of memory transplant will not, on its own, solve the relearning problem. The difficulties that arise in Bernecker’s treatment of transplants is symptomatic of a deeper problem in his characterization of remembering. According to Bernecker’s Contiguity View, the remembering relation is sustained when “there is no spatio-temporal gap in the causal chain and the content of the traces is the same as, or sufficiently similar to, the contents both of the past representation and the state of recall” (2010, 14). Once the initial trace has been formed, the remembering relation to that trace is sustained so long as two conditions are met: 1) causal contiguity and 2) content preservation.

Cases of relearning are a challenge for the causal theory of memory; they are designed to meet these conditions, matching remembering on both connection and content. Nothing about what is represented tells the two apart: remembering and relearning produce exactly similar representations. And, given that the result is a representation, the causal process leading up to it must be gapless, as causation cannot occur across a temporal distance. When Pete relearns that *rollercoasters are dangerous*, he does so through a chain of events that can be traced back to the initial event via the well-preserved contents of his diary. If Bernecker wants to rely on these two

Penultimate draft. Forthcoming in *Canadian Journal of Philosophy*. Please do not cite without permission (skrobins@ku.edu)

conditions alone, then he should accept relearning as a form of remembering. If Bernecker wants to block cases of relearning, then he must add another condition.

Bernecker does add an additional causal condition to his analysis of remembering. It is a *counterfactual condition*: “If S hadn’t represented at T1 that p* he wouldn’t represent at T2 that p” (2010, 148). This condition is of little use for sorting out cases of relearning. First, the counterfactual condition is invoked to deal with further wayward causation cases, not to address relearning. The trace condition was meant to have dispatched with relearning cases already. Second, it is unclear how the addition of a counterfactual strengthens the analysis. The counterfactual offers a way of describing the causal relation between the two events in question (here S’s representations at T1 and T2), but the existence of a causal relation between these two representations has already been established by the original causal condition. Even if the counterfactual could help to explain further why cases of relearning from an independent source do not qualify as instances of remembering (because in these cases the counterfactual is false), it will not help with the diary cases, which are captured by the counterfactual. If Pete hadn’t represented the rollercoaster ride at T1, he would not have written about it in his journal, making the relearning possible.

Bernecker resists adding further conditions because he wants to retain a feature of his initial analysis, namely, its compatibility with accounts of personal identity (as discussed in §2). If an account of remembering required the person who learns a representation to be the same person who later remembers that representation, then such an account would be circular. Since Bernecker’s view has no such requirement, it offers a way of grounding identity without circularity. In fact, Bernecker’s conditions on memory causation mirror those that identity theorists themselves propose. Shoemaker (1970), for example, proposes to build an account of

Penultimate draft. Forthcoming in *Canadian Journal of Philosophy*. Please do not cite without permission (skrobins@ku.edu)

personal identity up from a view of memory that requires (i) an initial experience, (ii) a content that matches the initial experience, and (iii) a causal connection between the experience and that content. Bernecker requires (i) traces to be created in the brain, (ii) the preservation of content sufficiently similar to the original, and (iii) a causal chain that exhibits temporal contiguity. The only difference is that Shoemaker, and other personal identity theorists, refer to these states as *quasi-memories*, in recognition of the fact that a state that meets these three requirements is similar to, but not the same as, remembering (Collins 1997). Memory builds in a commitment that the rememberer is the same person who had the learning event; quasi-memory does not.

It is, of course, open to Bernecker to allow cases of relearning into his account of remembering. It would go against the stated aims of his Contiguity View and his description of relearning as “the basic problem facing the causal theory of memory” (2010, 129). Nonetheless, it would allow him to retain the connection to theories of personal identity. But it’s worth pausing to consider the implications of such a move. What would a theory of remembering look like if temporal contiguity of content was the sole requirement on the remembering relation? This constraint can be met by nearly all instances of information transfer, across people, places, and times. If Pete has written in his diary that *rollercoasters are dangerous*, then not only can he remember this claim—by calling it to mind or via relearning, transplant, or replication—the same option is available to anyone who reads Pete’s diary, so long as they do so in a way that respects temporal contiguity. The diary can outlast Pete, making its contents accessible for others to remember long after Pete is gone and rollercoasters cease to exist. Even those who read the diary and think thoughts other than those that Pete had when he made the entry may also be remembering, so long as what they represent is inferentially related to Pete’s initial representation. I might read it and think that people who are afraid of rollercoasters are irrational,

Penultimate draft. Forthcoming in *Canadian Journal of Philosophy*. Please do not cite without permission (skrobins@ku.edu)

whereas you might do so and think that theme parks should do more to evaluate the safety of their rides. On Bernecker's Contiguity View, we are both remembering.⁹

Contiguity may pick out an interesting set of causal relations between mental representations, one that could be used to characterize the successful transmission of information across a wide range of contexts. It fails, however, to clarify the nature of the causal process that is involved in remembering uniquely. By allowing traces to persist across changes in location, medium, person, and content, Bernecker has created an account that is too permissive. If any credence is to be given to the idea of remembering as a distinct causal connection between two mental representations, it cannot be derived from considerations of contiguity alone. Contiguity is too weak to distinguish between remembering and relearning.

3.2 Cognitive Contiguity: Relearning as Re-Perceiving

Bernecker's Contiguity View is unable to distinguish remembering from relearning. This may lead one to worry that a similar fate awaits all causal theories of remembering, or to assume that avoiding this fate would require developing an altogether different theory of causation. These pessimistic conclusions can be avoided. Contiguity fails to solve the relearning problem because it is not a problem that can be solved by appeal to a principle of causation. Relearning is a deviant causal chain, but a causal chain nonetheless, and so tags along with whichever theory of causation one endorses.¹⁰ The problem is, in this way, internal to the causal theory, demanding for its solution a way of distinguishing between different sorts of causes. Solving it does not require saying more about *causation*; it requires saying more about *remembering*.

⁹ These worrisome implications could be prevented by disallowing cases of trace replication.

¹⁰ Perhaps a view of causal pluralism (e.g., Cartwright 2004) could avoid this worry. But even if there are multiple types of causation, it may be still be a stretch to say that there are so many types that remembering and relearning each have their own. I do not pursue this possibility here.

Penultimate draft. Forthcoming in *Canadian Journal of Philosophy*. Please do not cite without permission (skrobins@ku.edu)

The distinction between remembering and relearning is a *cognitive* distinction. To ask whether an accurate representation of a past event is an instance of remembering or an instance of relearning is to ask about the history of that representation's relation to the person who now represents it. Is the capacity to produce this representation something that has been retained, or something that has been regained? Remembering and relearning are both causal relations but they differ because, in the former case, the capacity to produce this representation stays within the same cognitive system from which the subsequent representation is produced. In relearning, the system loses this capacity and so the information must be reacquired.

My appeal to a cognitive system might invite worry, given that the boundaries of cognition are a matter of much debate in contemporary philosophy of mind. Defining and defending the border between perception and cognition is not without its difficulties. There are many accounts of the distinction between perception and cognition, as well as ongoing debates as to whether perception is cognitively penetrable (and if so, how thoroughly penetrable it may be).¹¹ A proponent of the causal theory of memory need not provide his or her own account of this distinction; doing so is beyond what is required for distinguishing between remembering and relearning. The difference between remembering and relearning is marked by the difference between retaining and reacquiring information, wherever that turns out to be. The Extended Mind Theory, discussed above, allows for broad consideration of what counts as a part of the cognitive system. For present purposes, it matters less where the cognitive boundary is located than what is used to establish it. And the boundary of the cognitive system that matters here is the perceptual border. Information that is relearned, rather than remembered, is available to the

¹¹ See, e.g., Zeimbekis & Raftopoulos (2015).

Penultimate draft. Forthcoming in *Canadian Journal of Philosophy*. Please do not cite without permission (skrobins@ku.edu)

person only via *re-perception*.¹² Remembering is often defined in contrast to perceiving, as information initially acquired from perception but no longer sustained by a perceptual connection to its source. This need not correspond to the body—a representation could count as inside the cognitive system even if it is outside the skin, or remains on the body but beyond the perceptual border, as would a tattoo. Remembering requires the retained capacity to produce representations of previously acquired information. Such a capacity is retained within a cognitive system if it is a capacity that can manifest without re-perception.

Some cases of relearning may resist capture by this distinction. A person may forget something and then relearn that piece of information by making inferences from other things he or she already knows.¹³ Pete may, for example, forget that he received a green bike for his 8th birthday, but know that he received a bike that year and that his favorite color was green and so infer that he received a green bike. The possibility of such cases makes clear that the cognitive contiguity requirement will require strengthening. This can be done by making further distinctions within the cognitive system—specifically, by distinguishing between memory and general reasoning/inference systems. The condition I have in mind would look something like Michaelian’s (2011) requirement that the causal connection required for remembering be sustained by a properly functioning memory system.¹⁴ The aim of Michaelian’s condition is to distinguish memory from both perception and inference, making it an attractive solution to the present problem. Appealing to a reliable memory system picks out the specific form of cognitive

¹² Levine (2009) explores a similar possibility in considering the case of Leonard, the amnesic protagonist of the film *Memento*.

¹³ I am grateful to an anonymous reviewer for introducing this challenge.

¹⁴ Michaelian’s condition: “the causal chain goes continuously via a reliable memory system (responsibly for the (re)construction of the trace and the current representation)” (2011, 335).

Penultimate draft. Forthcoming in *Canadian Journal of Philosophy*. Please do not cite without permission (skrobins@ku.edu)

contiguity required for distinguishing remembering from the myriad other ways one might reacquire past information that has otherwise been lost to forgetting.

Establishing a fully fleshed out alternative is beyond the scope of this paper. Here my aim has been to show that a simple contiguity account will not suffice for distinguishing remembering from relearning. There is a sense in which remembering requires contiguity—it relies upon the persistence of the medium in which the initial representation was formed. But this is not bare causal contiguity; it is cognitive, causal contiguity (and, most likely, a form of cognitive contiguity supplied uniquely by a memory system). Much remains to be said about cognitive contiguity, and its relation to claims about the mind's extension and the possibility of personal identity without circularity. For now, it is enough to show that relearning is a problem that can be solved, just not by appeal to causation alone.

4. Contiguity and Prompting

Even if Bernecker's Contiguity View is too weak to solve the relearning problem, it might still be able to solve another problem of wayward causation: strict prompting. Attempts at remembering often involve prompts, external aids that serve as reminders of previously acquired information. Although most prompts offer a harmless boost, there are troubling cases where the prompt's contribution to the representation threatens to override that of the memory trace because the prompt contains all of the information that the trace could possibly provide. These are *strict prompts*. To see the problem, imagine the earlier case of Pete, with a slight modification—he still takes the ride on the dangerous rollercoaster and writes about it in his diary, but he does not have the amnesia-producing accident afterward. Nonetheless, Pete's memory is far from perfect. Sometimes he needs gentle reminders to call past representations to

Penultimate draft. Forthcoming in *Canadian Journal of Philosophy*. Please do not cite without permission (skrobins@ku.edu)

mind. Now suppose Pete reads his diary, which prompts him to think that rollercoasters are scary. Is he remembering? Bernecker wants to allow that Pete could be remembering in this case, and so he proposes that remembering is compatible with strict prompting so long as the trace retains its contiguity to the representation produced.

Bernecker defines a prompt as any force other than the trace that plays a causal role in bringing about the representation, whether it is recruited to the process deliberately or automatically (2010, 141). He appeals to Martin and Deutscher's (1966) three-fold classification of prompts as partial, complete, or strict. Partial prompts, as the name suggests, provide only some of the information contained in the remembering representation. As Bernecker describes them, partial prompts are *insufficient* for producing the representation. Sufficiency here is to be understood as the degree of overlap between the content of the prompt and the content of the representation. If the prompt contained only the information *the name of the tallest rollercoaster in Texas begins with a T*, for example, then this would be insufficient for representing that *The Titan is the tallest rollercoaster in Texas*. Remembering is often initiated by partial prompts, ones that have varying degrees of overlap with the ensuing representation: brushing shoulders with a stranger reveals a scent that brings back images of childhood, hearing a song on the radio reminds you of your old college roommate, being asked for the capital of Ohio brings to mind Columbus.

Both complete and strict prompts are, in contrast to partial prompts, sufficient: they contain all of the content that could be in the remembering representation. Complete prompts are those that mirror the representation's content. These are prompts used to test recognition memory, where one is presented with an item and asked whether he or she has encountered it before. Peggy might, for example, ask Pete whether he knows that the *Titan* is the tallest

Penultimate draft. Forthcoming in *Canadian Journal of Philosophy*. Please do not cite without permission (skrobins@ku.edu)

rollercoaster in Texas. Even though cases of complete prompting are sufficient for the representation, the rememberer can still provide a unique contribution to the act of remembering by adding that the particular fact or event is something he or she has previously known or experienced. Cases of strict prompting include both the complete prompt and the additional recognition claim. If Peggy implores Pete to remember his past ride on the *Titan*, she might not only relay all of the details of this event to him, but also remind him that this is something that he previously experienced. Strict prompting is defined as a prompt that contains every possible detail that a rememberer could contribute to the representation.

Bernecker wants to allow all forms of prompting—partial, complete, and strict—to count as cases of remembering. He is explicit about his intent to include cases of strict prompting. As he states, “that a person requires strict prompting before he can recount certain events shouldn’t be ruled out as inadmissible for genuine memory” (2010: 142). The question of whether remembering occurs in a case where prompting is involved, Bernecker argues, should be determined by the strength of the trace’s contribution to the representation. He thus proposes the following condition to handle instances of prompting:

Causal Strength Condition: The memory trace is at least an INUS condition for S’s representation at T2 that p. If the memory trace is an independently sufficient condition, it is not pre-empted by another independently sufficient condition (2010, 144).

Bernecker’s causal strength condition harkens back to Mackie’s (1965) analysis of causes as INUS conditions for their effects. Mackie’s analysis claims that a cause is an *insufficient* but *necessary* part of a condition that is itself *unnecessary* but *sufficient* for the result—INUS, for short. Mackie’s analysis is helpful for determining the causal influence of the memory trace in cases of prompting because it is focused on singular causation and accounting for how we can

Penultimate draft. Forthcoming in *Canadian Journal of Philosophy*. Please do not cite without permission (skrobins@ku.edu)

identify the causal influence of a particular factor when the effect in question has multiple causes. Bernecker uses the INUS condition to explain what the trace's role must be when the trace is insufficient for bringing about the representation produced. In such cases, remembering is still possible when the trace acts alongside a prompt, so long as it plays a non-redundant role in the process—i.e., insofar as it is an insufficient but necessary component of producing the effect. Bernecker does not elaborate on the underlying notion of redundancy, but we can imagine it as something like the sufficiency claim above. A trace with the content *The name of the tallest rollercoaster in Texas begins with a T* would be non-redundant if it acted alongside a prompt such as *The answer ends in "itan"* to produce a representation that *The Titan is the tallest rollercoaster in Texas*. Either way, the trace must make a unique causal contribution to the representation produced.

The causal strength condition concludes with consideration of cases where the trace is sufficient for producing the representation. In such cases, the question is whether the trace serves as the proximal cause of the representation. So long as the trace is not pre-empted by some other sufficient cause, the requisite proximity is maintained and remembering occurs. The causal strength condition prohibits only remembering in cases where a (sufficient) trace is pre-empted by another independently sufficient condition. The restriction may appear minor, and sensible, but it appears to block the possibility of remembering in the case of strict prompting—a form of prompting Bernecker claims he wants to include. A strict prompt is, by definition, sufficient to bring about the representation, and further, it is active in bringing about the representation *prior* to any such activity from the trace. It is a pre-emptive and sufficient cause of the representation. For example, if Peggy asks Pete to remember the name of the tallest rollercoaster in Texas, but blurts out the answer before Pete can call the name *Titan* to mind, then Pete cannot be said to

Penultimate draft. Forthcoming in *Canadian Journal of Philosophy*. Please do not cite without permission (skrobins@ku.edu)

remember that *The Titan is the tallest rollercoaster in Texas*, even though he is now representing it. This conclusion could be avoided if the sufficiency were framed in terms of what the trace would have been sufficient to produce had the prompt not interfered.¹⁵ But this way of addressing cases of strict prompting is ruled out explicitly. For Bernecker, what matters is the way that the representation comes about in the circumstances:

Evaluations of whether seeming memories are genuine memories should be based on the causation that is actually operative rather than on whether there is independent sufficient causation held in reserve on a deviant route (2010, 143).

That Pete has a memory trace whose content would have been sufficient to produce this representation, had Peggy suppressed her outburst, is irrelevant.

Bernecker is inclined to reject cases of preemption by a sufficient prompt because they violate the contiguity requirement on remembering. If the prompt causes the representation, then the trace is not the representation's proximal cause, nor even a component of this proximal cause. There is no temporal gap, as the prompt has wedged its way in to the causal chain, but the prompt breaks the contiguity between the trace and the representation. The causal strength condition thus makes sense as a commitment entailed by Bernecker's commitment to contiguity, but it is in conflict with his stated aim of allowing the possibility of remembering in cases of strict prompting. Although Bernecker set out to find a condition on the trace's causal strength that would be compatible with strict prompting, his Contiguity View rules out that possibility. For solving the strict prompting problem, contiguity is too strong: it blocks accurate representations of past events that a theory of remembering should include.

Or should it? Bernecker does not provide independent motivation for the idea that remembering is compatible with strict prompting. He simply assumes this is so, following Martin

¹⁵ I am grateful to an anonymous reviewer for suggesting this possible response.

Penultimate draft. Forthcoming in *Canadian Journal of Philosophy*. Please do not cite without permission (skrobins@ku.edu)

and Deutscher's (1966) commitment to the inclusion of such cases. Given that Bernecker's account has been unsuccessful in this task, it makes sense to pause and ask whether one ought to try and make strict prompting and remembering compatible. In what follows, I argue that their compatibility ought to be maintained, at least in cases where the trace sustains the representation originally produced by a strict prompt.

There is an apparent difference between cases where a strict prompt acts alone to produce a representation and cases where the strictly prompted representation is then sustained by a memory trace. Consider again the case of Peggy's strict prompting of Pete regarding his ride on the *Titan*. In cases where strict prompting acts alone, Pete represents all of the information that Peggy provides, including Peggy's claim that he had this experience. If he trusts Peggy, Pete may even believe that he has ridden the *Titan* before. He simply takes Peggy at her word, but never goes beyond that representation to remembering. In contrast, when Pete's representation is strictly prompted by Peggy and then becomes reinforced by the trace, things change. Pete might suddenly exclaim, *now I remember!* This may only occur after exhaustive and repetitive prompting and it may do nothing to alter Pete's representation—he may simply parrot Peggy's prompt back to her once more, only this time with enthusiasm. Still, something in this latter case is different.

This may seem like an appeal to an experience of remembering, the sort of phenomenological feature that many consider unreliable. But even if such cases often (or always) involve an *aha!* feeling, this need not be their defining feature. The difference between a representation produced by strict prompting alone and a representation produced by strict prompting and then sustained by a trace is the role that this representation can play in the person's subsequent thought and behavior. When the representation of the *Titan* ride is sustained

Penultimate draft. Forthcoming in *Canadian Journal of Philosophy*. Please do not cite without permission (skrobins@ku.edu)

by the trace, Pete may go on to recall that this ride occurred during the same summer that he broke his arm. Or it may cause him to be in a bad mood because it reminds him of the strain on family vacations in the years leading up to his parents' divorce. Representations produced by strict prompting alone will likely produce associations as well, but the associations available will be different. Given that the trace has not been reactivated, the available associations will be restricted to the more general—other thoughts about rollercoasters, other times that Peggy's testimony has proven more reliable than his own, and so on.

The difference between these two cases reveals that remembering is important not only because of what it allows us to represent, but also because of what it allows us to *do* with the information once it is represented. This is why, even in cases of strict prompting, the influence of the memory trace makes a difference. The prompt may provide the information, but the reacquisition makes contact with a distinct set of further thoughts and actions than what is provided by the trace. The memory trace's associations are unique; the trace ties the acquisition of this information to a particular time in the person's past. The acquisition of this particular capacity is influenced by the capacities one has already and by the other capacities being acquired around the same time. The same piece of information acquired at two distinct times (once during remembering, once during relearning) will be situated differently within one's network of acquired ideas and capacities. This claim need not be read as a strong endorsement of a particular Associationist view of the mind. It requires only the more limited point that learning is a process of connecting new information with the information that one possesses already. This is why those with more experience in a given domain respond differently to new information in that domain than do those who are unfamiliar. A veteran barkeep can better retain a list of drink orders than a novice (Beach 1988) and a chess expert does not *try* to memorize the position of

Penultimate draft. Forthcoming in *Canadian Journal of Philosophy*. Please do not cite without permission (skrobins@ku.edu)

pieces on the board—she simply does so in considering her next move (Ericsson, Patel, & Kintsch 2000). It is also why particular pieces of information may be linked for an individual, even when there is no connection between them otherwise. It may have been while standing in line for the *Titan* that Pete first learned that sheep and other ruminant animals have four stomachs. And for this reason, visits to a farm may always remind him of rollercoasters. When and how information is acquired shapes the way an individual can put that information to later use.

It is now clear why the distinction between remembering and relearning matters. Remembering and relearning have distinct kinds of causal history. They differ in how they interact with other thoughts and in the routes through which they manifest in behavior, even though they can produce exactly similar representations. By remaining within the cognitive system, as I have argued that memory traces do, the capacity to remember retains its role within that system. Information that must be re-perceived has to earn a new role through the act of relearning. Even when the capacity to remember manifests only after strict prompting, the very fact that it plays a role matters for determinations of remembering. Remembering can occur even when the trace merely reinforces or sustains a representation brought about by other means.

By focusing on temporal contiguity, Bernecker privileges the order of events by which a representation is brought to mind. But timing is not the critical feature. What matters for remembering is that a capacity that one has retained since a previous event is activated in the process of thinking about that event, regardless of whether this activation occurs before, after, or alongside a prompt. At times, Bernecker recognizes this point, as is reflected in his stated desire to include cases of strict prompting. His commitment to contiguity, however, makes this

Penultimate draft. Forthcoming in *Canadian Journal of Philosophy*. Please do not cite without permission (skrobins@ku.edu)

impossible. Contiguity is too strong a requirement: it disallows any representation for which the trace is not the proximal cause.

5. The Causal Condition on Remembering

Bernecker has argued that commitment to remembering as a causal process requires endorsing a particular view of the nature of memory causation. He argues that memory causation is characterized by the temporal contiguity it provides between two mental representations with sufficiently similar contents. This *Contiguity View* is intended to provide solutions to the wayward causation problems of relearning and strict prompting. I have argued that Bernecker's Contiguity View fails to characterize remembering uniquely. Relearning shares temporal contiguity with remembering, and the possibility of strict prompting shows that remembering is possible even when temporal contiguity is violated. Bernecker's contiguity requirement is thus inconsistent with his own aim of providing an account of the nature of memory causation that can withstand wayward causation problems, and more importantly, it violates widely shared intuitions about what remembering requires.

The failure of Bernecker's Contiguity View need not lead to pessimism about the possibility of formulating a sustainable form of the Causal Theory of Memory. There are alternative ways of addressing the wayward causation problems that face a causal theory of remembering. The distinction between remembering and relearning is a cognitive distinction, one that can only be captured by appeal to the subject who represents the information about a past event. What matters is not the relation the states in this chain of events bear to one another—whether they sustain causal contact—but rather how they relate to the individual who represents the information. The distinction between remembering and relearning matters because

Penultimate draft. Forthcoming in *Canadian Journal of Philosophy*. Please do not cite without permission (skrobins@ku.edu)

the difference in their causal histories influences the connections they have to the subject's other capacities and interests, affording them distinct roles in her further thought and action. And this, in turn, is why the activation of the capacity to remember matters, regardless of whether it occurs before, during, or after a prompt. By being active, the trace connects the current representation with the subject's other capacities. Remembering occurs when a person retains a memory trace that allows her to represent information acquired from a past event, such that this information can contribute—without re-perception—to future thought and action, either by producing or sustaining representations of the event.

References

- Beach, K. D. (1988). The role of external mnemonic symbols in acquiring an occupation. In M. M. Gruenberg, P. E. Morris, & R. N. Sykes (Eds.) *Practical Aspects of Memory: Current Research and Issues* (pp. 342–346). Chichester, England: Wiley.
- Bernecker, S. (2008). *The Metaphysics of Memory*. New York, NY: Springer.
- Bernecker, S. (2010). *Memory: A Philosophical Study*. Oxford, UK: Oxford University Press.
- Cartwright, N. (2004). Causation: one word, many things. *Philosophy of Science*, 71, 805–819.
- Clark, A. (2008). *Supersizing the Mind: Embodiment, Action, and Cognitive Extension*. Oxford, UK: Oxford University Press.
- Clark, A. & Chalmers, D. J. (1998). The extended mind. *Analysis*, 58, 7–19.
- Collins, A. W. (1997). Personal identity and the coherence of q-memory. *Philosophical Quarterly*, 47, 73–8.
- Cusmariu, A. (1980). A definition of impure memory. *Philosophical Studies*, 38, 305–308.
- Ericsson, K. A., Patel, V. L., & Kintsch, W. (2000). How experts' adaptations to representative task demands account for the expertise effect in memory recall: Comment on Vicente and Wang (1998). *Psychological Review*, 107, 578–592.
- Hume, D. (1739/1978). *A Treatise of Human Nature*, L. A. Selby-Bigge (ed.). Oxford, UK: Clarendon.

Penultimate draft. Forthcoming in *Canadian Journal of Philosophy*. Please do not cite without permission (skrobins@ku.edu)

Levine, J. (2009). Leonard's system: Why doesn't it work? . In C. Grau (ed.) *Eternal Sunshine of the Spotless Mind (Philosophers on Film)*. New York, NY: Routledge (pp. 45–64).

Mackie, J. L. (1965). Causes and conditions, *American Philosophical Quarterly*, 2, 245–264.

Malcolm, N. (1963). *Knowledge and Certainty*. Ithaca, NY: Cornell University Press.

Martin, C. B., & Deutshcer, M. (1966). Remembering. *Philosophical Review*, 75, 161–196.

Michaelian, K. (2011). Generative Memory. *Philosophical Psychology*, 24, 323–342.

Schechtman, M. (2011). Memory and identity. *Philosophical Studies*, 153, 65–79.

Shoemaker, S. (1970). Persons and their pasts. *American Philosophical Quarterly*, 7, 269–285.

Sutton, J., & Windhorst, C. (2009). Extended and Constructive Remembering: Two Notes on Martin and Deutscher. *Crossroads: An Interdisciplinary Journal for the Study of History, Philosophy, Religion, and Classics*, 4, 79–91.

Zeimbekis, J., & Raftopoulos, A. (Eds., 2015). *The Cognitive Penetrability of Perception: New Philosophical Perspectives*. Oxford, UK: Oxford University Press.