Generative memory

Kourken Michaelian

This paper explores the implications of the psychology of constructive memory for philosophical theories of the metaphysics of memory and for a central question in the epistemology of memory. I first develop a general interpretation of the psychology of constructive memory. I then argue, on the basis of this interpretation, for an updated version of Martin and Deutscher's influential causal theory of memory. I conclude by sketching the implications of this updated theory for the question of memory's status as a generative epistemic source.

Keywords: Causal Theory of Memory; Constructive Memory; Epistemology of Memory

There has been relatively little philosophical work in recent years on the metaphysics of memory (the nature of memory in general). This is presumably because most philosophers have assumed that something close to Martin and Deutscher's (1966) causal theory of memory is right. Bernecker's (2008) recent work, e.g., the first book-length work on the metaphysics of memory to appear in some years, defends a causal theory of memory very much in the spirit of Martin and Deutscher's; and there have been no very recent attacks on the theory. While I, too, ultimately want to defend a theory in the spirit of the classical causal theory, I also maintain that our confidence in the causal theory has been to a certain extent unfounded, for the psychology of constructive memory poses a significant challenge to the causal theory of memory, and few philosophers concerned with the metaphysics of memory have so far taken empirical work on the constructive nature of memory into account.

This paper therefore explores the implications of the psychology of constructive memory for the causal theory of memory. In section 1, I develop a general interpretation of the psychology of constructive memory. In section 2, I argue, on the basis of this interpretation, for an updated version of the causal theory of memory and compare the updated theory to existing attempts to take the constructive

Kourken Michaelian is Postdoctoral fellow at the Institut Jean-Nicod.
Correspondence to: Kourken Michaelian, Institut Jean-Nicod, Pavillon Jardin, Ecole Normale Supérieure, 29 rue d'Ulm, Paris 75005, France. Email: kourken.michaelian@ens.fr

character of memory into account in the metaphysics of memory. I conclude in section 3 by sketching the implications of the updated theory for the question of memory's status as a generative epistemic source.

Before proceeding, a note on the concepts of generation and preservation of beliefs and belief-contents. The generation of new content occurs when memory produces content in addition to that which it took as input; this can occur either before retrieval, by means of transformation of content received from other sources, or at retrieval, by means of transformation of content stored by memory. The generation of a new belief occurs at retrieval when the agent accepts a retrieved record that she did not previously accept; this can occur either when the agent's memory stored a content that she did not previously accept or when retrieval generates a new content. Memory is preservative (of beliefs or contents) if it is not generative.

1. Constructive Memory and Doxastic Generation

Though it is a commonplace in psychology that memory is constructive, 'construction' and its cognates (as they are used in the context of memory research) are difficult to define precisely. A precise definition is perhaps not possible at this point, but we can describe the various ways in which memory is constructive.

Note that we should reject the obvious proposal that construction in memory is precisely a matter of content generation. The proposal is not implausible: memory is, after all, called 'constructive' in part because it is "supplemental," because "some of the changes that occur between study and test involve 'memory' for information that was not contained in the input" (Koriat, Goldsmith, & Pansky, 2000, p. 485). But to identify construction with content generation is (for reasons given below) to disregard too many of the phenomena usually regarded as exemplifiying construction in memory.

Though the proposal is to be rejected, it can be used to illustrate an important ambiguity in the concept of constructive memory. Content generation in memory can occur at either of two points: content other than that provided by the initial representation might be incorporated into the memory trace before it is retrieved; and content other than that provided by the memory trace might be incorporated into the representation resulting from retrieval. In order to eliminate this sort of ambiguity, Koriat, Goldsmith, and Pansky (2000; following Alba & Hasher, 1983) propose using 'construction' to refer to relevant processes occurring at encoding, and 'reconstruction' to refer to relevant processes occurring at retrieval. For reasons given below, I will modify this usage slightly, allowing 'construction' to refer also to relevant processes (if there are any) occurring during consolidation. (I will continue to refer to memory in general as 'constructive'.)

The ambiguity in talk of construction is familiar to psychologists; it is less familiar to philosophers. I emphasize it here primarily for dialectical reasons. As long as we are in the grip of a picture on which a memory trace or engram is a faithful record of experience, we will be inclined to suppose that remembering must be a matter of

retrieving the trace unaltered. But if we realize that the trace itself is constructed, it should begin to seem much less obvious that remembering should be a matter of retrieving it unaltered: if there is nothing lost by construction at one stage, there need be nothing lost by reconstruction at the other; if there is something gained by construction at one stage, there might also be something gained by reconstruction at the other.

I take my general account of construction and reconstruction in memory from an influential paper by Alba and Hasher (1983).[1] They describe four ways in which memory might be constructive:

> **Selection:** Only certain incoming stimuli are selected for encoding.
> **Abstraction:** The meaning of a message is abstracted from the syntactic and lexical features of the message.
> **Interpretation:** Relevant prior knowledge is invoked.
> **Integration:** A holistic representation is formed from the products of the selection, abstraction, and interpretation processes.

We can now see why it is a mistake to identify construction with content generation: both selection and abstraction reduce the quantity of information incorporated into the engram—they eliminate content. Interpretation, in contrast, makes additional content available. In integration, finally, the simplified content and the newly-available content are incorporated into a single memory trace. Alba and Hasher characterize reconstruction in retrieval as follows:

> **Reconstruction:** Whatever information was selected for representation and is still accessible is used, together with general knowledge, (roughly) to generate a hypothesis about what might have happened.

Among the most interesting examples of reconstructive processes are metamemory processes (monitoring and control processes involved in retrieval) (Smith, Shields, & Washburn, 2003). In source monitoring, e.g., the source of a memory trace is inferred (perhaps automatically and unconsciously) from certain features of the content (e.g., its level of detail) (Mitchell & Johnson, 2000). In fluency processing, the source of a memory trace is inferred (again, perhaps automatically and unconsciously) from features of the retrieval process (e.g., how easy it is) (Kelley & Jacoby, 1998).[2] In general, constructive and reconstructive memory processes are processes in which information is actively transformed by the memory system.

The constructive nature of encoding (especially selection and abstraction) is illustrated by the phenomenon of false recognition, in which subjects who study a number of words related to a non-presented theme word, e.g., are likely later to "recognize" the theme word (Roediger & McDermott, 1995). Fuzzy trace theory (Brainerd & Reyna, 2002) provides a plausible candidate explanation of this effect: according to the theory, both verbatim and gist traces are formed during encoding; gist traces are more readily accessible and so tend to be preferred; hence subjects are led falsely to recognize non-presented theme words (Koriat, Goldsmith, & Pansky, 2000, p. 492). The phenomenon of boundary extension, in which subjects remember having seen more of a scene than they saw in fact (Intraub, Bender, & Mangels, 1992),

provides another revealing illustration of the constructive nature of encoding (especially interpretation). According to one explanation of this effect, it is a consequence of the fact that information about the likely layout of the scene is automatically retrieved and then incorporated into the memory of the scene (Koriat, Goldsmith, & Pansky, 2000, p. 495; Schacter, Norman, & Koutstaal, 1998, p. 305). A final illustration (especially of integration) is provided by the superportrait phenomenon, in which caricatures are often recognized faster and more accurately than are faithful portraits. According to Rhodes (1996), this is likely the consequence of the fact that the relevant representations are in a certain sense highly schematic— they emphasize the distinctive properties of the things represented (Koriat, Goldsmith, & Pansky, 2000, p. 489).

The list of examples of construction can be extended almost indefinitely. There are deep differences among the various types of construction, and there are interesting theoretical debates about how best to explain them. But there is no significant debate over whether encoding is constructive. Some of the construction occurring at encoding seems, moreover, to involve content generation. Consider again the phenomenon of boundary extension. It is not that the subject first forms a belief to the effect that the scene before her eyes appears to have such-and-such a layout, and then infers that the scene must extend in certain ways. Nor is it that the initial sensory representation already contains the extended boundary (Park, Intraub, Yi, Widders, & Chun, 2007). The representation of the scene is modified automatically as a memory for the scene is formed. The modification proceeds in stages: the early stages perhaps reflect the operation of perceptual schemas; the later stages apparently involve normalization in memory (Intraub, Bender, & Mangels, 1992).

Some might argue that the fact that encoding in human memory is constructive is of little significance to the philosophical theory of memory: construction at encoding, they might suggest, merely reflects contingent features of the way in which memory is implemented in our species; the theory of memory thus need not take it into account, but should, instead, focus on what there is in common between constructive human memory and other, non-constructive possible types of memory. But while it might seem to us that we can easily imagine non-constructive forms of memory, it is unclear that these imaginings reflect anything more than a folk theory of memory. For it is probably not in the end a contingent feature of human memory that it is constructive. Consider, e.g., the case of memory for gist: in many cases, memory records not a literal representation of an experience, but, rather, its gist. Schacter and Addis (2007, p. 778) suggest that extraction of gist is an adaptive feature of memory, an economical way for a system with limited storage capacity to store the most important features of experience (the aspects of experience knowledge of which is most likely to be useful to the organism again in the future). This kind of reasoning is widespread in the constructive memory literature: constructive features of memory, features of memory that might appear to a philosopher's eye to be mere peculiarities of the human memory system are, it is argued, in fact adaptive. There is thus good reason to think that construction at encoding is part of the real essence of memory, that, in other words, it is not merely a contingent feature of certain memory systems.[3]

It is natural to suppose that even if construction occurs at encoding, a memory trace, once formed, will remain stable until retrieved (ignoring decay over time). But in fact a process of consolidation, during which the memory trace is not yet fully stabilized, intervenes between encoding and permanent storage (Ambrogi Lorenzini, Baldi, Bucherelli, Sacchetti, & Tassoni, 1999; Squire & Zola-Morgan, 1991). If this process (which unfolds over a period of many years) is interrupted, the memory can be partly or entirely lost.

There is reason to think that consolidation is no more a contingent feature of the human memory system than is constructive encoding. McClelland, McNaughton, and O'Reilly (1995) argue (on the basis of connectionist modelling results) that consolidation is required in order to allow new memories to be incorporated without distorting existing memories: "if the changes were made rapidly they would interfere with the system of structured knowledge built up from prior experience with other related material" (p. 435). On their influential proposal, the hippocampus repeatedly replays episodes to the neocortex, allowing the neocortex gradually to discover their common structure; this can account for the categorization of memories and the formation of memories for what specific experiences have in common. If something like this view is right, the philosophical theory of memory should take consolidation into account, and it should take it into account as a constructive process, one in which, moreover, content (e.g., generalizations) is potentially generated.

Note that even once consolidation is achieved, memory traces are not permanently stabilized. It is not just that reconstruction can occur during retrieval. It is, rather, that we must acknowledge also a process of reconsolidation: memories become malleable when retrieved, and a period of reconsolidation is required before they can be said simply to be stored again (Sara, 2000). As Dudai (2004) puts it, "it is not the time since encoding that determines the susceptibility of a trace to interventions, but rather the functional state of the trace: An active (retrieved) trace can be truncated, but also augmented; an inactive (stored) trace is immune to such manipulation" (pp. 68–69). The transition from dispositional to occurrent and back, in other words, need not leave a memory unchanged.

Memory, Schacter, and Addis (2007) write, "is not a literal reproduction of the past, but rather is a constructive process in which bits and pieces of information from various sources are pulled together" (p. 773). Just as perceptual illusions can be studied to reveal the normally invisible constructive workings of the perceptual system, memory distortions can be studied to reveal the normally invisible reconstructive workings of memory (Roediger, 1996). To this end, a wide range of memory distortions have been investigated. False recognition arising from misleading post-event suggestions (Zaragoza & Mitchell, 1996), e.g., might occur because thinking inaccurately about an event can create a representation of the event that cannot easily be distinguished from a representation of the actually witnessed event (Schacter, Norman, & Koutstaal, 1998, p. 294). Misleading retrieval cues can be incorporated into retrieved memories to produce false beliefs (as when students with high GPAs overestimate their marks for classes for which they received low marks) (Bahrick, Hall, & Berger, 1996; Schacter, Norman, & Koutstaal, 1998, p. 306).

Related phenomena include the retrospective bias and the ''knew it all along'' effect (Schacter, Norman, & Koutstaal, 1998, p. 306): in the former, recall is distorted to render memory consistent with present beliefs (Levine, 1997); in the latter, subjects adjust their memories of earlier probability estimates in light of their current knowledge of the occurrence or non-occurrence of the relevant events (Fischhoff & Beyth, 1975). Additional clues to the nature of reconstruction come from observations of various biases in spatial memory (Koriat, Goldsmith, & Pansky, 2000, p. 490): e.g., landmarks produce asymmetric distance estimates, suggesting that spatial information is subject to a sort of interpretation according to the demands of the context of retrieval (McNamara & Diwadkar, 1997). Perhaps the most dramatic example of reconstruction is the phenomenon of confabulation, in which subjects are led to recall entirely fictional events in great detail (Hyman, Husband, & Billings, 1995; Hyman & Pentland, 1996): merely imagining a fictional event increases the probability that the subject will later remember it as having occurred (Schacter, Norman, & Koutstaal, 1998, p. 307).

The source monitoring framework (developed primarily by Johnson and her colleagues) provides plausible explanations of confabulation and other memory distortions (Mitchell & Johnson, 2000). I discuss source monitoring in more detail below; for the moment, I point out only that a crucial element of the framework is the claim that retrieval involves reconstruction and, in particular, that it involves content-generating reconstruction. Nor is the source monitoring framework unique in this respect: any account of retrieval, if it takes seriously the claim that memory distortions reveal the normal workings of memory, will grant that retrieval involves reconstruction and, moreover, that it involves content-generating reconstruction. When all goes well, there is a close match between the retrieved memory (including the newly generated content) and the initial representation; the content generation involved in retrieval is thus normally invisible. But under certain circumstances (when memory is distorted) the retrieved memory will fail to match the initial representation in virtue of incorporating inappropriate newly generated content. The point is that the same reconstructive processes are at work in both cases: reconstruction in general (and content generation in particular) is not the exception but the rule.

Just as it would be a mistake to suppose that construction at encoding or consolidation is an idiosyncrasy of human memory, it would be a mistake (and for similar reasons) to suppose that reconstruction at retrieval is an idiosyncrasy of human memory. It is likely that the reconstructive nature of retrieval is an adaptive feature of memory. Recent work on episodic memory, e.g., tends to emphasize the involvement of the system in prospection (imagining the future; Buckner & Carroll, 2007) Schacter and Addis (2007) write: ''since the future is not an exact repetition of the past, simulation of future episodes may require a system that can draw on the past in a manner that flexibly extracts and recombines elements of previous experiences—a constructive rather than a reproductive system'' (p. 774). If this line of research is on the right track, then we should expect reconstructive retrieval to be an essential feature of memory.

The acknowledgment that memory is constructive and reconstructive raises a question about how its reliability is ensured: were memory simply to keep faithful records of experience, its reliability would be uninteresting; but given that memory is thoroughly (re)constructive, its reliability[4] requires some explanation. The complete story about the reliability of memory will cite many different features of the system; but in order to illustrate the general structure of the explanation, I want to focus on source monitoring in particular (Johnson, Hashtroudi, & Lindsay, 1993; Johnson & Raye, 1981, 2000; Mitchell & Johnson, 2000).[5]

It is crucial for the possibility of source monitoring that memory traces bear marks of their origins: an origin in experience, in particular, is indicated by "embeddedness in spatial and temporal context; embeddedness in supporting memories, knowledge, and beliefs; and the absence of consciousness of or memory for the cognitive operations producing the event or belief" (Johnson, 1988, p. 57).[6] It is crucial for the possibility of the failure of source monitoring that the marks in question are not infallible indicators of origins: "because of variability within . . . source types, the distributions of features of memories from different processes and events overlap" (Mitchell & Johnson, 2000, p. 180), so that some memories originating in fantasies, e.g., might be more detailed than some memories originating in experience.

The consequence of these two features of memory traces is that attention to properties of a stored belief or content can provide a reliable indication of its origin, though not a perfectly reliable indication. Such attention thus can provide the basis for reliable predictions about the origins of beliefs and contents:

> The characteristics of mental experience cannot serve as a precise signature, or "tag," that specifies its origin. Rather, remembering always involves judgments about how the quantity and quality of these characteristics compare to expectations about characteristics of memories from various sources. So, for example, if a mental experience had substantial perceptual detail (e.g., visual), one would tend to attribute it to a perceived event (e.g., something one saw), since, on average, memories from perceived events contain quite a bit of perceptual detail. (Mitchell & Johnson, 2000, p. 180)

On the source monitoring framework, the reliability of memory is secured not despite the reconstructive nature of retrieval but (in part) because of it: because they rely on reliable indicators of the origins of traces, reconstructive retrieval processes tend to come to the right conclusions about those origins. Because the indicators are imperfectly reliable, these inferences will sometimes go wrong. (They can go wrong also if prior inferences involved in encoding went wrong.) Such failures of source monitoring can account for many of the memory distortions reviewed above (Mitchell & Johnson, 2000, p. 181): in confabulation, e.g., the retrieval mechanism misattributes a memory to experience on the basis of its level of detail—a reliable process produces an inaccurate result.

The general story about the reliability of constructive memory follows the same pattern. Memory involves both construction and reconstruction. Content generation can occur at either point. This means that there is in general no literal record of experience: an initial representation is transformed first at encoding

(and perhaps also during consolidation); this already-transformed representation is subject to additional transformations at retrieval. But these transformations are in no way arbitrary: they are designed to ensure that most of the beliefs eventually produced by retrieval are accurate.[7]

It is important to note that, though monitoring is a form of metamemory, it is nonetheless a component of memory itself: it need not be that a subject first retrieves a content and then (in a distinct cognitive act) forms a judgement as to the source of the content; the formation of the judgement about the source of the content can itself be a step on the way to forming the belief ultimately "retrieved" from memory. Thus monitoring can generate some of the content of the eventual memory belief: the subject stores content deriving from her experience and then, during retrieval, generates additional content to the effect that the stored content derives from experience; the additional content is then combined with the stored content to form the content of the belief eventually formed by the retrieval process, a belief that the subject had an experience with a certain content. The content of the retrieved representation goes beyond the content originally stored. Nor does the additional content derive from a simple inference from premises to conclusions. A retrieval process goes to work on a stored content and generates additional content. The additional content is then combined with the stored content. The belief finally produced is a belief with this combined content.

2. Updating the Causal Theory of Memory

Having described construction and reconstruction in memory, having clarified the relation of (re)construction to content generation, and having explained how the reliability of constructive memory is secured, we are in a position to update the classical causal theory of memory so that it is adequate with respect to constructive memory.

Martin and Deutscher were, in developing the causal theory,[8] reacting to the deficiencies of the various empiricist theories of memory popular at the time (1966). Empiricist theories differ in their details, but they have in common that they deny the need for a causal condition in the theory of memory: they deny that a causal connection between the initial representation and the later belief is necessary for genuine remembering to occur. As Martin and Deutscher (1966) show, such theories are inadequate. They describe a counterexample with the following structure: a subject observes an event and forms a belief that such-and-such an event occurred; the belief is stored in her memory store; later, the memory is lost (due to damage to her brain); after having lost the memory, the subject is hypnotized; the hypnotist implants in her a belief that such-and-such an event occurred; the implanted belief is indistinguishable from the original one (Martin & Deutscher, 1966, p. 174). An empiricist theory will imply that the subject remembers that the event occurred. But clearly she does not. The plausible suggestion made by Martin and Deutscher is that the empiricist theory delivers the wrong verdict about this case precisely because it

does not include a causal condition: what is missing in the case is a causal chain connecting the initial observation to the subject's later representation.

As Martin and Deutscher recognize, it is not enough simply to require that there be some causal connection or other between the observation and the representation, for many causal chains are simply not of the right sort (they are deviant). They describe a case with the following structure: a subject observes an event and forms a belief that such-and-such an event occurred; the belief is stored in her memory store; later, the memory is lost (due to damage to her brain); but in the meantime, the subject has told another about the event; after the subject has lost the memory, the other repeats her earlier description of the event back to her; time goes by, and the subject forgets having heard the repetition of the description—she now only believes that such-and-such an event occurred (Martin & Deutscher, 1966, p. 180). In this case, a causal chain does connect the observation of the event to the subject's later representation—the problem is that the causal chain is not of the right sort.[9] The first (but not the only) restriction on the class of admissible causal chains introduced by the causal theory is that in order for remembering to occur, the causal chain connecting the initial representation with the later representation must go via a memory trace.

The proposal to incorporate a reference to memory traces into the theory of memory has been surprisingly controversial (Sutton, 1998). One standard objection presupposes that the theory of memory is a conceptual analysis: the core idea is that we should not build a story about memory traces into the theory of memory, for the theory would then imply that certain neuroscientific assumptions are built into our concept of remembering (which, obviously, they are not). Such objections can perhaps be met, but I am under no obligation to do so here: the new causal theory is offered not as an analysis of the concept of memory, but rather as an account of memory itself.[10]

Though the causal theory of constructive memory is therefore unaffected by objections from adequacy conditions on conceptual analyses, a superficially similar objection is relevant to the theory. Zemach (1983) argues against building a reference to memory traces into the philosophical theory of memory on the ground that the theory of memory should not dictate to empirical science what it must discover about the workings of memory (pp. 32–33). The objection is puzzling, given that there has been no suggestion from empirical science that memory might function without memory traces. The objection is puzzling, that is, until we realize that Zemach does not count "holographic" traces as traces: since Martin and Deutscher explicitly describe memory traces as "structural analogues" of experiences—the model here is the grooves of a record—their version of the causal theory runs the risk of dictating to empirical science.

Two responses to the objection are available. We might, first, point out that since there has been no suggestion from empirical science that memory might function without memory traces of some sort, it is safe for philosophical theorists of memory to build into their theories a claim that causation of rememberings goes via memory traces of some sort (without specifying any further details). This might have been the

better option when Zemach wrote (in 1983). But our empirical understanding of memory has developed greatly since then, and we are now entitled to make a bolder move: there is overwhelming evidence from empirical investigations of memory that memory involves traces of precisely the holographic (or distributed or super-positional [Sutton, 1998]) sort mentioned by Zemach. We might therefore revise the causal theory so that it refers to memory traces described not as structural analogues but rather as holographic. Such an approach (which I endorse) does not in any interesting sense dictate to empirical science what it must discover about the workings of memory—the revision precisely allows the philosophical theory of memory to learn from what empirical science has already discovered about the workings of memory.[11]

I said earlier that it is still not sufficient to guarantee that a causal connection is appropriate to require that it goes via a memory trace. Recall the second of Martin and Deutscher's cases described above: in this case, though the trace is destroyed by the time the false remembering occurs, the causation does initially go via a trace. Martin and Deutscher are aware of this problem, and they therefore add a requirement to the effect (approximately) that the memory trace must be causally active at the time of remembering, that it must contribute directly to the remembering. Their discussion of this requirement is somewhat obscure,[12] but the basic idea is clear enough: the memory trace has to exist all along; and it has to be doing causal work at the time of the remembering. The addition of this requirement is a step in the right direction, but it is still not quite sufficient. As it stands, the theory remains vulnerable to the following sort of counterexample:

> A past experience of mine, *e*, produced a physical memory trace in my brain. I do not remember *e* at all, but, owning an autocerebroscope, I can inspect the anatomy of my brain. What I do then is read the said trace from time to time, much as I read the inscriptions in my diary. (Zemach, 1983, pp. 37–38)

It is important not to be distracted by the fanciful details of the case: Zemach is onto something here, viz., that there are multiple ways in which a trace might be doing causal work at the time of (putative) remembering. In order to avoid this sort of counterexample, we might simply add the requirement that the trace cause the representation in virtue of having been retrieved. But this specific requirement is rendered redundant by the next modification that I will propose.

Consider a case in which a subject has a badly damaged episodic memory system. Normally, we have seen, the episodic system is reliable: though it is constructive, its reliability is ensured by the properties of the constructive processes it employs, e.g., by the use of reliable heuristics in source monitoring. Suppose that the memory system has, in virtue of being damaged, become highly unreliable—perhaps the heuristics used in source monitoring are wildly inappropriate, with the consequence that imagined scenes are often classified as experienced, experienced scenes are often classified as dreamed, etc. The memory system might, despite its unreliability, retrieve an occasional accurate representation (e.g., an experienced scene might be represented as experienced). The classical causal theory implies that on those

occasions the subject remembers. But clearly she does not: the case is not one of a subject with a malfunctioning system which happens occasionally to function properly—in such a case, the subject clearly manages occasionally to remember—but one of a subject with a permanently damaged system which, by chance, occasionally produces an accurate representation. I propose that, in order to avoid this sort of counterexample, an additional requirement be incorporated into the causal theory: the causal chain must go not only via a memory trace but through a properly functioning (that is, a reliable) memory system.[13] The proposal is, in short, that the classical causal theory of memory should be replaced with a causal reliabilist theory of memory.[14]

A second, more dramatic modification turns the causal reliabilist theory into the new causal theory. The old theory is silent about the relations between the content of the initial representation, that of the memory trace, and that of the eventually retrieved representation. But implicit in discussions of the theory are something like the following assumptions:

(1) The trace provides the full content of the later representation.
(2) The initial representation provides the full content of the trace (even if some of the content available in the representation is not recorded by the trace).

These conditions together rule out content-generating construction and reconstruction. The problem is that content generation regularly occurs at various points along the path from the initial representation to the representation eventually retrieved: the trace need not provide the full content of the representation, for reconstructive processes involved in retrieval can generate new content; and the initial representation need not provide the full content of the trace, for constructive processes involved in encoding and consolidation can generate new content. The implicit assumptions about the relations between the content of the initial representation, that of the memory trace, and that of the later representation are, in short, inadequate in light of the involvement of content-generating constructive and reconstructive processes in memory.[15]

My proposal is that we replace these assumptions with the following conditions:

(1) The content of the later representation does not go too far beyond that of the trace.
(2) The content of the trace does not go too far beyond that of the initial representation.

To see the point of these formulations, consider a case involving the phenomenon of boundary extension.[16] Someone might see part of a scene and encode a trace representing more of the scene than she saw. Later, she retrieves the trace and represents the larger scene. Does she remember the scene? The urge to answer that obviously she does not remember the scene, since (after all) she did not see the whole of it, should be resisted, for to answer that way is to commit oneself to ruling out very many perfectly ordinary cases of remembering. Boundary extension is not the exception to the rule; it is the rule. For reasons of method, studies of constructive memory phenomena tend to focus on cases in which construction goes wrong, cases,

that is, in which there is a mismatch between the eventual representation and the initial experience. There is a danger that this will lead us to think that construction typically results in such mismatches, to think of construction as unreliable and hence of constructed memories as not being bona fide memories. But construction, we have seen, is a feature of many ordinary cases of remembering. Construction does not typically result in mismatches—constructive memory is reliable. Thus we need not and should not say that the subject does not genuinely remember the scene because she did not see the whole of it; we should instead say that she remembers the scene even though she did not see the whole of it. This will sound almost paradoxical—but only as long as we are in the grip of a naive picture of memory. Memory is no more a matter of passive transmission than perception is of passive reception: epistemologists have already been able to some extent to reconcile themselves to thinking of perception as an active, constructive source of knowledge; we now must get used to conceiving of memory as a similarly active and constructive source of knowledge.

Given that content generation is compatible with remembering, why have I suggested that there are limits on how much content generation is compatible with remembering? The answer to this question has to do with the way in which the respective roles of perception, inference, and memory in the cognitive economy of a subject are to be distinguished. Perception is a belief-independent process which generates content. Inference is a belief-dependent process which generates content. Memory, I have argued, is a process which is sometimes belief-independent and sometimes belief-dependent and which generates content. There is thus a question about how to characterize the unique role played by memory in the subject's cognitive life: how can memory be distinguished (at this level of abstraction) from perception and inference? My proposal is that the distinction should be drawn in terms of the quantity of content typically generated by these belief-producing processes. Perception involves the generation of relatively little new content: the content of a perceptual belief normally does not go much beyond the content of the experience which leads to it. Inference—I refer here to full-blown inferences, rather than to the inference-like constructive processes involved in perception or memory—can involve the generation of significant new content. The proposal is that memory is situated somewhere between these two extremes: like inference, it can generate significant new content; like perception, its generation of new content is subject to certain rigid constraints.[17]

One might worry that the proposed similarity conditions are an ad hoc means of ensuring that memory is reliable despite being constructive. But the worry is misdirected: we want the theory to account for the reliability of memory; and the new causal theory already builds in reliability by requiring causation via a reliable memory system. One might nevertheless worry that once we have taken on board the lesson that memory is thoroughly constructive, we should drop even the weakened similarity conditions proposed here. It would be interesting to explore the theory that would result from this modification, but I take it that the psychology of constructive memory does not force us to drop the similarity requirement entirely: it tells us that remembering generally involves much less similarity than we would intuitively

suppose; but it does not suggest that a retrieved memory need not be at all similar to any representation encoded earlier.

The proposed conditions on the adequacy of a trace are vague as they stand. I suspect that this vagueness is ineliminable: there is a difference between remembering and merely seeming to remember; but there is no reason to expect that we can draw the line with much precision. If a subject ''remembers'' something that is not the case, we know that she merely seems to remember. If she ''remembers'' something that is the case, and if little content generation occurs along the relevant causal chain, we know that she genuinely remembers. If she ''remembers'' something that is the case, and if massive content generation occurs along the chain, we know that she merely seems to remember. We cannot say more than this.

The following is my official formulation of the new causal theory:

> **The Causal Theory of Constructive Memory:** *S* remembers *P* iff
>
> (a) *P*;
> (b) *S* represents *P*;
> (c) there is a causal chain running back from *S*'s current representation of *P* to an earlier representation of hers;
> (d) the causal chain goes continuously via a (distributed) memory trace with the content *P* (or something sufficiently close to *P*);
> (e) the causal chain goes continuously via a reliable memory system (responsible for the (re)construction of the trace and the current representation);
> (f) *S*'s earlier representation had the content *P* (or something sufficiently close to the content of her memory trace); and
> (g) there is an appropriate relation between *P* and *S*'s earlier representation.

A few remarks about the formulation. Clause (a): the theory is a theory of factive memory; that *S* remembers *P* entails that *P* is true. Clause (b): a subject can have a representation with a given content without having a propositional attitude with the content. The theory thus covers, in addition to cases in which a belief is stored and retrieved, cases in which a subject stores a content without having formed a belief with the content and cases in which a subject retrieves a content without forming a belief in the content. Clause (g): various relations are ''appropriate'' in the relevant sense: the subject might perceive *P*, infer *P*, etc. The notion of appropriateness at work here is intended to be a generalization of the notion of reliability to cover not only belief-producing processes but also representation-producing processes in general. Finally: the theory is not a theory of knowing from memory; *S* might remember *P* and yet not know *P*—this will happen if she has a defeater for her belief *P*.

Lest it be thought that the new causal theory departs only in insignificant ways from the classical causal theory (of which it is indeed a recognizable descendant), I pause to emphasize its novel features. In order to accommodate the constructive nature of memory, the new causal theory incorporates a very relaxed similarity requirement. In order to explain how memory can be reliable, despite this relaxed requirement, it incorporates a reference to properly functioning memory systems. These two modifications produce a theory which, unlike the classical causal

theory, is consistent with the empirical psychology of memory. And, unlike the classical theory, the new theory has counterintuitive implications: unlike the classical causal theory, e.g., it permits that one can successfully remember more than one learned. The new causal theory thus goes significantly beyond the classical causal theory.

The new causal theory also differs from existing accounts of generative memory.[18] While he maintains that semantic ("factual") memory is purely preservative, Dokic (2001) argues that episodic memory is not. His focus, however, is on a very specific kind of generation: in order to account for the fact that episodic memory provides the subject with "a reason to believe that the information carried by it . . . comes directly from the subject's own past life" (unlike factual memory, which only provides the subject with a reason to believe that an event happened; Dokic, 2001, p. 216), Dokic proposes an account of episodic memory on which "the fact that [it comes directly from the subject's own experience] is presented in the memory experience itself" (2001, p. 228); while this is an account of episodic memory as generative of a specific sort of content, it does not cover the sort of generative transformation of content on which I focused in section 1 and which are covered by the new causal theory. (The new causal theory covers the sort of generation on which Dokic focuses; see the discussion of source monitoring above.)

The new causal theory also differs from Bernecker's (2008) theory, on which only certain limited sorts of transformation of content are compatible with remembering. Bernecker's theory permits that remembering is compatible with transformations that merely keep information current (e.g., tense updating) and with transformations that reduce content (e.g., existential generalization), but it explicitly forbids transformations which generate additional (supplemental) content (2008, p. 164). While on Bernecker's theory memory is perhaps generative in a weak sense— according to him, the content produced as output by memory need not be literally the same as that it took as input—his account of construction, unlike that incorporated into the new causal theory, does not cover transformations in which information not present in the input is incorporated into memory. Thus his theory is empirically inadequate, as it misclassifies cases of genuine but supplemental memory of the sort discussed in section 1 as cases of merely apparent memory.

Finally, the new causal theory differs from Matthen's (2010) recent account, which does not deal with the nature of the causal connection required for remembering, the question of how much generation of content is compatible with remembering, or the problem of how the reliability of memory is ensured despite its constructive character. Matthen's focus is, rather, epistemological: how can memory justify a belief, given that it does not simply preserve content?[19] But as Matthen argues, in the course of dealing with this question, that what is preserved in memory is "a trace from which it is possible to reconstruct an image or belief" (2010, p. 14), and as his discussion makes clear that he acknowledges that the trace can differ from the experience from which it stems, his account is compatible with the new causal theory.

3. Epistemological Implications of the New Causal Theory

The default view about memorial justification is that memory is capable of preserving but not of generating justification.[20] Lackey has recently challenged this ''preservationist'' view, arguing for a moderate form of ''generationism'' (she does not use the term), a view according to which memory can generate justification by generating a new belief with a previously stored content. The gist of the argument is the following: when a subject initially acquires some information, it can be stored in her memory without her first forming a belief having it as its content; later retrieval of the resulting record can result in the formation of such a belief; assuming that the relevant cognitive processes are sufficiently reliable, that belief will be justified (Lackey, 2005, pp. 650–651).[21] Moderate generationism already departs significantly from the default preservationist view, but if the argument of sections 1 and 2 above is right, we should, at least given one standard theory of epistemic justification, endorse a more radical form of generationism, a view according to which memory can generate justification both by generating a new belief with a previously stored content and by generating a new belief along with its very content.[22]

According to reliabilism about epistemic justification (Goldman, 1992), a belief is justified just in case it is produced by a reliable process, i.e., a process that tends to generate mostly true beliefs. Given this theory of justification, radical generationism follows directly from the new causal theory of memory. The new causal theory permits that memory is sometimes a simple belief-dependent process, a process that takes a belief as input and delivers the same belief as output; in such cases, it can only be conditionally reliable and so can confer no justification on the belief that it produces—in these cases, memory only preserves justification. According to moderate generationism, memory can generate justification by generating a new belief with a previously stored content. The new causal theory permits that memory can generate a new belief with a previously stored content; when it does so, it functions as a belief-independent process. Since the theory says that memory functions reliably in these cases, it implies (together with reliabilism) that memory then generates justification. Acording to radical generationism, memory can generate justification by generating a new content, along with a belief with that content. The new causal theory permits that memory can generate a new content, along with a belief with that content; when it does so, it functions either as a belief-independent process or as a belief-dependent process with some non-doxastic inputs. Since the theory says that memory functions reliably in these cases, it implies (together with reliabilism) that memory then generates justification: cases of the former sort are straightforward; in cases of the latter sort, since the process has some non-doxastic starting-points, at least some of the justification for the belief that it produces might have been generated by the process itself rather than simply inherited from the doxastic starting-points of the process.[23]

Epistemologists often assume that the epistemology of memory is relatively straightforward (at least as far as the question of justification is concerned): since they assume that memory only preserves beliefs, they conclude that memory can only

preserve justification; the thought is that it is a straightforward matter to work out the epistemology of memory, given the correct theory of memory. Those who make this assumption are half right and half wrong: it is indeed a relatively straightforward matter to work out the epistemology of memory, once the correct theory of memory is in hand; it is just that epistemologists have so far typically relied on a natural but inadequate theory of memory—with an empirically adequate theory of memory as constructive in hand, memory appears as clearly generative of beliefs and therefore of justification.

Acknowledgements

Notes

[1]   Alba and Hasher discuss constructive and reconstructive processes in the context of schema theory; and they object to schema theory on the ground that memory representations are richer and more detailed than the theory suggests. But even if we acknowledge that the particular theory criticized by Alba and Hasher overexaggerates the extent to which memory is schematic, their concepts of selection, abstraction, interpretation, integration, and reconstruction remain available; their argument is not meant to show that there is no role for constructive and reconstructive processes in memory.

[2]   I return to metamemory below in the context of a discussion of the means by which the reliability of constructive memory is secured.

[3]   This response to the worry that the constructive nature of encoding in human memory need not be taken into account by the philosophical theory of memory assumes (in addition to the view that declarative memory is a natural kind [Michaelian, forthcoming b]) certain standard (but not uncontroversial) views about natural kinds and necessity. But even those who reject the views in question should dismiss the worry. Construction is a solution to the problem of finite storage capacity (Cherniak, 1983). We should thus suspect that all actual (and most interesting possible) creatures have constructive memories. Whatever we think about natural kinds and necessity, then, we should be prepared to acknowledge that construction is an important feature of memory, one that ought to be taken into account by the philosophical theory of memory. Note that this does not imply that construction by itself is a complete solution to the problem of finite storage capacity (on the role of forgetting in solving the problem of finite storage capacity, see Michaelian, forthcoming a).

[4]   Memory distortions, again, are investigated not because memory is supposed to be unreliable but rather because it is supposed to be reliable: memory distortions are cases in which a process produces an inaccurate memory despite its reliability.

[5]   Note that Mitchell and Johnson motivate the source monitoring framework precisely by pointing out that memory distortions raise the question how the reliability of constructive memory is achieved (2000, p. 180). For additional discussion of monitoring, see Michaelian (2010).

[6]   Quoted in Johnson and Raye (2000, p. 37).

[7]   This is not all that they are designed to achieve, of course: the transformations are designed also to ensure that the beliefs eventually produced by retrieval are useful.

[8]   Though reference is often made to the causal theory of memory, relatively few papers developing the theory have so far been written. (Compare the number of papers devoted to developing the analogous causal theory of perception.) I therefore rely heavily on Martin and Deutscher (1966), the locus classicus of the causal theory. See also Deutscher (1989).

[9]   A natural thought at this point in the development of the theory is that the causal chain is of the wrong sort in virtue of passing outside the subject's body. While natural, the thought universalizes a contingent feature of human memory: Martin and Deutscher point out that there could be creatures whose memories are stored in removable devices, devices which are clearly not part of their bodies (1966, p. 181). This sort of possibility is especially salient in these days of active externalism (Clark & Chalmers, 1998).

[10] Though Martin and Deutscher's approach is largely a priori, they, too, are concerned to theorize memory itself, rather than the concept of memory, as Sutton and Windhorst (2009) emphasize in their recent reappraisal of the paper.

[11] There remains a slight risk that empirical science will in the future arrive at results incompatible with the revised theory; but to point this out is just to say that empirical science might in the future arrive at results incompatible with its own present view of memory.

[12] It turns on a distinction (which I will not attempt to summarize here) between something's being "operative in a circumstance" and something's being "operative for a circumstance."

[13] The description of a memory system appropriate in this context is the most general description of the system that we can formulate. See Michaelian (forthcoming b).

[14] I draw inspiration here from Pendlebury's (1994) causal reliabilist theory of perception.

[15] I write for simplicity as if there is only one initial representation and one trace. This is for convenience only: as emphasized above, in fact construction in memory is often a matter of the integration of content from a variety of sources.

[16] We can easily describe analogous cases involving reconstruction.

[17] Note that the proposal is not that the differences among the quantities of content generated by perception, memory, and inference provide a criterion by means of which to distinguish among these sources; I want here simply to understand how the sources differ with respect to the generation of content.

[18] Though Sutton's (1998) discussion of the reconstructive character of memory informs my approach here, his focus is on the nature of memory traces; he thus does not aim to develop the sort of general account of memory at issue in the debates over the causal and empiricist theories of memory.

[19] The context here is Burge's discussion of the "acceptance principle," according to which "a person is entitled to accept as true something that is presented as true [including by memory] and is intelligible to him, unless there are stronger reasons not to do so" (Burge, 1993, p. 467); Burge's defence of the principle assumes that there is a purely preservative form of memory.

[20] Lackey (2005) cites Audi (1995), Dummett (1994), Owens (2000), and Plantinga (1993) as endorsing some form of preservationism about memorial justification.

[21] In addition to this argument, she gives arguments for generationism which start from the fact that a subject's relation to defeaters (beliefs that the subject has or beliefs that she ought to have) for a belief can change over time while the belief is preserved in memory. As these arguments are thoroughly discussed in Lackey (2007) and Senor (2007), I will not take them up here.

[22] Moderate generationism and radical generationism are discussed in my PhD thesis, *On memory and testimony* (Michaelian, 2009). They should not be confused with the

different positions in the debate around preservationism recently labelled "moderate generativism" and "radical generativism" by Bernecker (2010).

[23] Note that the argument for radical generationism does not depend on this claim about the capacity of "mixed" processes to generate justification: if the causal theory of constructive memory is right, content generation sometimes occurs in the context of a mixed process but sometimes also occurs in the context of a belief-independent process; the claim thus matters only for determining how frequently the generation of justification occurs via content generation.

References

Alba, J. W., & Hasher, L. (1983). Is memory schematic? *Psychological Bulletin*, *93*, 203–231.

Ambrogi Lorenzini, C. G., Baldi, E., Bucherelli, C, Sacchetti, B., & Tassoni, G. (1999). Neural topography and chronology of memory consolidation: A review of functional inactivation findings. *Neurobiology of Learning and Memory*, *71*, 1–18.

Audi, R. (1995). Memorial justification. *Philosophical Topics*, *23*, 31–45.

Bahrick, H. P., Hall, L. K., & Berger, S. A. (1996). Accuracy and distortion in memory for high school grades. *Psychological Science*, *7*, 265–271.

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

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

Brainerd, C. J., & Reyna, V. F. (2002). Fuzzy-trace theory and false memory. *Current Directions in Psychological Science*, *11*, 164–169.

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

Burge, T. (1993). Content preservation. *Philosophical Review*, *102*, 457–488.

Cherniak, C. (1983). Rationality and the structure of human memory. *Synthese*, *57*, 163–186.

Clark, A., & Chalmers, D. (1998). The extended mind. *Analysis*, *58*, 7–19.

Deutscher, M. (1989). Remembering 'Remembering'. In J. Heil (Ed.), *Cause, mind, and reality: Essays honoring C.B. Martin* (pp. 53–72). Dordrecht: Kluwer.

Dokic, J. (2001). Is memory purely preservative? In C. Hoerl & T. McCormack (Eds.), *Time and memory* (pp. 213–232). Oxford: Oxford University Pres.

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

Dummett, M. (1994). Memory and testimony. In B. K. Matilal & A. Chakrabarty (Eds.), *Knowing from words* (pp. 251–272). Dordrecht: Kluwer.

Fischhoff, B., & Beyth, R. (1975). I knew it would happen: Remembered probabilities of once-future things. *Organizational Behavior and Human Performance*, *13*, 1–16.

Goldman, A. (1992). *Liaisons: Philosophy meets the cognitive and social sciences*. Cambridge, MA: MIT Press.

Hyman, I. E., Husband, T. H, & Billings, F. J. (1995). False memories of childhood experiences. *Applied Cognitive Psychology*, *9*, 181–197.

Hyman, I. E., & Pentland, J. (1996). The role of mental imagery in the creation of false childhood memories. *Journal of Memory and Language*, *35*, 101–117.

Intraub, H., Bender, R. S., & Mangels, J. A. (1992). Looking at pictures but remembering scenes. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *18*, 180–191.

Johnson, M. K. (1988). Discriminating the origin of information. In T. F. Oltmanns & B. A. Maher (Eds.), *Delusional beliefs* (pp. 34–65). New York: Wiley.

Johnson, M. K., Hashtroudi, S., & Lindsay, D. S. (1993). Source monitoring. *Psychological Bulletin*, *114*, 3–28.

Johnson, M. K., & Raye, C. L. (1981). Reality monitoring. *Psychological Review*, *88*, 67–85.

Johnson, M. K., & Raye, C. L. (2000). Cognitive and brain mechanisms of false memories and beliefs. In D. L. Schacter & E. Scarry (Eds.), *Memory, brain, and belief* (pp. 35–86). Cambridge, MA: Harvard University Press.

Kelley, C. M., & Jacoby, L. L. (1998). Subjective reports and process dissociation: Fluency, knowing, and feeling. *Acta Psychologica*, *98*, 127–140.

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

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

Lackey, J. (2007). Why memory really *is* a generative epistemic source: A reply to Senor. *Philosophy and Phenomenological Research*, *74*, 209–219.

Levine, L. J. (1997). Reconstructing memory for emotions. *Journal of Experimental Psychology: General*, *126*, 165–177.

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

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

McClelland, J. L., McNaughton, B. L., & O'Reilly, R. C. (1995). Why there are complementary learning systems in the hippocampus and neocortex: Insights from successes and failures of connectionist models of learning and memory. *Psychological Review*, *102*, 419–457.

McNamara, T. P., & Diwadkar, V. A. (1997). Symmetry and asymmetry of human spatial memory. *Cognitive Psychology*, *18*, 87–121.

Michaelian, K. (2009). *On memory and testimony*. Unpublished doctoral dissertation. University of Massachusetts, Amherst.

Michaelian, K. (2010). In defence of gullibility: The epistemology of testimony and the psychology of deception detection. *Synthese*, *176*, 399–427.

Michaelian, K. (forthcoming a). The epistemology of forgetting. *Erkenntnis*.

Michaelian, K. (forthcoming b). Is memory a natural kind? *Memory Studies*.

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

Owens, D. (2000). *Reason without freedon: The problem of epistemic normativity*. London: Routledge.

Park, S., Intraub, H., Yi, D.-J., Widders, D., & Chun, M. M. (2007). Beyond the edges of a view: Boundary extension in human scene-selective visual cortex. *Neuron*, *54*, 335–342.

Pendlebury, M. (1994). Content and causation in perception. *Philosophy and Phenomenological Research*, *54*, 767–785.

Plantinga, A. (1993). *Warrant and proper function*. Oxford: Oxford University Press.

Rhodes, G. (1996). *Superportraits: Caricatures and recognition*. Hove: Psychology Press.

Roediger, H. L. (1996). Memory illusions. *Journal of Memory and Language*, *35*, 76–100.

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

Sara, S. J. (2000). Retrieval and reconsolidation: Toward a neurobiology of remembering. *Learning and Memory*, *7*, 73–84.

Schacter, D. L., & Addis, D. R. (2007). The cognitive neuroscience of constructive memory: Remembering the past and imagining the future. *Philosophical Transactions of the Royal Society of London B*, *362*, 773–786.

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

Senor, T. D. (2007). Preserving preservationism: A reply to Lackey. *Philosophy and Phenomenological Research*, *74*, 199–208.

Smith, J. D., Shields, W. E., & Washburn, D. A. (2003). The comparative psychology of uncertainty monitoring and metacognition. *Behavioral and Brain Sciences*, *26*, 317–373.

Squire, L. R., & Zola-Morgan, S. (1991). The medial temporal lobe memory system. *Science*, *253*, 1380–1386.

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

Sutton, J., & Windhorst, C. (2009). Extended and constructive remembering: Two notes on Martin and Deutscher. *Crossroads*, *4*, 79–91.

Zaragoza, M. S., & Mitchell, K. J. (1996). Repeated exposure to suggestion and the creation of false memories. *Psychological Science*, *7*, 294–300.

Zemach, E. M. (1983). Memory: What it is, and what it cannot possibly be. *Philosophy and Phenomenological Research*, *44*, 31–44.