

The information effect: constructive memory, testimony, and epistemic luck

Kourken Michaelian

Received: 2 May 2011 / Accepted: 9 August 2011 / Published online: 19 August 2011
© Springer Science+Business Media B.V. 2011

Abstract The incorporation of post-event testimonial information into an agent’s memory representation of the event via constructive memory processes gives rise to the misinformation effect, in which the incorporation of inaccurate testimonial information results in the formation of a false memory belief. While psychological research has focussed primarily on the incorporation of *inaccurate* information, the incorporation of *accurate* information raises a particularly interesting epistemological question: do the resulting memory beliefs qualify as knowledge? It is intuitively plausible that they do not, for they appear to be only luckily true. I argue, however, that, despite its intuitive plausibility, this view is mistaken: once we adopt an adequate (modal) conception of epistemic luck and an adequate (adaptive) general approach to memory, it becomes clear that memory beliefs resulting from the incorporation of accurate testimonial information are not in general luckily true. I conclude by sketching some implications of this argument for the psychology of memory, suggesting that the misinformation effect would better be investigated in the context of a broader “information effect”.

Keywords Memory · Eyewitness memory · Constructive memory · Testimony · Epistemic luck · Misinformation effect

The incorporation of post-event testimonial information into an agent’s memory representation of the event via constructive memory processes gives rise to the misinformation effect, in which the incorporation of inaccurate testimonial information results in the formation of a false memory belief. While psychological research has focussed primarily on the incorporation of inaccurate information, the incorporation of accurate

K. Michaelian (✉)
Department of Philosophy, Bilkent University, Bilkent, Ankara 06800, Turkey
e-mail: kirk.michaelian@gmail.com

information raises a particularly interesting epistemological question: do the resulting memory beliefs qualify as knowledge?

This puzzle has not so far been discussed in the literature¹; given the centrality of episodic memory as an epistemic source, however, its importance should be clear. It is intuitively plausible that the beliefs in question do not meet the standards for knowledge, for they appear to be only luckily true. I will argue, however, that, despite its intuitive plausibility, this view is mistaken: once we adopt an adequate (modal) conception of epistemic luck and an adequate (adaptive) general approach to memory, it becomes clear that memory beliefs resulting from incorporation of accurate testimonial information are not in general luckily true. I conclude by sketching some implications of this argument for the psychology of memory, suggesting that the misinformation effect would better be investigated in the context of a broader “information effect”.

The question with which I am concerned here should be distinguished from the question, discussed by [Martin and Deutscher \(1996\)](#) and, more recently, by [Bernecker \(2010\)](#), whether true prompting is compatible with memory. There are two differences here: first, incorporation is not the same phenomenon as prompting; second, the two questions belong to different domains. The question about true prompting is whether a suggestible agent who accepts a true prompt can be said to remember—whether the receipt of the prompt violates the causal condition on remembering. The question about helpful incorporation, in contrast, is about whether a memory belief resulting from helpful incorporation can qualify as knowledge—whether the belief can meet the conditions for knowledge. Note, first, that the relevant mechanism is different in the two cases. As explained in Sect. 1.4, incorporation refers to the incorporation of testimonial information into an agent’s stored memory representation of an event. Prompting, in contrast, refers to the impact of testimony received at the time of retrieval. Thus we cannot assume that an argument about one case will generalize to the other. Note, second, that while the question discussed under the heading of prompting is metaphysical, the question about incorporation discussed here is epistemological: a belief resulting from incorporation might meet the conditions for knowledge without meeting the causal condition on remembering; a belief resulting from prompting might meet the causal condition on remembering without meeting the conditions for knowledge. One could pose an epistemological question about prompting, and one could pose a metaphysical question about incorporation, but I deal with neither of those questions here (though I do briefly discuss the implications of incorporation for the causal theory of memory in Sect. 4.4).

1 The misinformation effect

I begin by setting out my basic conception of episodic memory beliefs and the constructive process that produces them.

¹ In a recent article, Shanton discusses a related question ([Shanton 2011](#)); I respond to her arguments in Sect. 4 below.

1.1 Episodic memory

There are various ways of drawing the episodic/semantic distinction, but I will assume here that episodic memories are those stored by the episodic system, and similarly for semantic memories; the episodic system, in turn, is devoted to representations of the agent's experience, while the semantic system is devoted to other information (Michaelian 2011a).² But note that, while the episodic system stores representations of the agent's experience, it does not store only experiential information (information deriving wholly from experience): experiential information can be transformed or combined with information deriving from other sources during encoding or consolidation before achieving stable long-term storage; moreover, during the reconsolidation process following retrieval, stored information can be modified under the impact of incoming information from various sources.

There are also various ways of defining memory belief, but I will assume that an (occurrent) memory belief is a belief formed by endorsing information produced by retrieval from memory; an episodic memory belief, in turn, is a belief formed by endorsing information produced by retrieval from episodic memory. Note that this characterization of memory belief assumes that memory stores representations or contents rather than beliefs.³ While epistemologists sometimes write as if memory stores (dispositional) beliefs, this involves a mistake. Memory does not store only representations that the agent has endorsed, and it does not store only representations that the agent would endorse upon retrieval (Michaelian 2011b). Nor is this an imperfection of memory—to the extent that the function of memory can be reduced to information storage, it is, roughly, the storage of relevant information, rather than the storage of endorsed information (Michaelian 2011c).⁴ When an agent retrieves a memory representation and forms a belief with that representation as its content, a metacognitive decision process intervenes between retrieval and belief-formation, even in cases where the agent believed the information when the memory was initially encoded (Koriat and Helstrup 2007).⁵ In other words: memory does not store beliefs, it literally produces them.

² In order to distinguish true episodic memory from the mere “episodic-like” or what-where-when (WWW) memory possessed by some animals (Suddendorf and Corballis 2007), a satisfactory definition will probably have to take into account the role of autoeotic consciousness, the distinctive phenomenology of episodic memory (Tulving 1983; Vandekerckhove and Panksepp 2009).

³ Vosgerau (2010) argues, more radically, that memory does not store even contents but rather templates from which determinate contents can be constructed.

⁴ See, e.g., work on the role of relevance in determining forgetting (Anderson 1991) or recent work on fitness-relevance on encoding (Nairne and Pandeirada 2008).

⁵ This decision process is in most cases more or less phenomenally invisible, since it is not typically accomplished by a reflective, deliberate decision by the agent: retrieved information is typically endorsed/rejected as a result of system 1 rather than system 2 processing (Frankish 2010) (e.g., automatic source monitoring processes Johnson 1997; Mitchell and Johnson 2000), i.e., it is unreflective and automatic.

1.2 (Re)construction in memory

The definition of memory belief in terms of endorsement of retrieved representations will be inadequate if we assume a naïve picture of episodic retrieval as a process in which a discrete stored record is retrieved unchanged; episodic retrieval is, rather, a reconstructive process (Tulving 1982; Schacter and Addis 2007; Matthen 2010; Robin 2010; Michaelian 2011d). The memory system does not simply make available a stored representation of an event. Before retrieval outputs a memory, inferences are made to interpret stored information. Various items of stored information are combined; these items stem not only from experience but also from other sources—non-experiential information might earlier have been incorporated into episodic memory, or activated semantic information might be incorporated at the time of retrieval. And information available in the context of retrieval can also be incorporated. The reconstructive character of the retrieval process means that retrieval typically produces a representation that was never stored in memory or even previously entertained by the agent—that retrieval, strictly speaking, is a matter of the production of a new representation. For this reason, too, memory can be said literally to produce new beliefs.

As noted above, memory is not only reconstructive but also constructive, that is, the stored representations which are recombined in retrieval are themselves products of a constructive encoding process. Indeed, information is not invulnerable to change while stored: stable storage is preceded by a lengthy process of consolidation, during which records are labile; and a similar process of reconsolidation, during which records are again malleable, follows retrieval (Michaelian 2011d; Dudai 2004; Alberini 2005). Moreover, if connectionist models of memory are right, the distributed nature of storage means that stored records are never invulnerable to change, as the addition of new information inevitably alters stored information to some extent (McClelland 1995, 2011; Sutton 1998). Constructive encoding will play little role in my discussion here: construction at encoding does give rise to an epistemological question analogous to the question about reconstructive retrieval on which I focus (see Sect. 1.3), but, though I suspect that the answers to these questions are roughly analogous, it would require a separate argument to establish this. On some explanations of the misinformation effect, reconsolidation plays the main role; on the explanation that I favour (see Sect. 1.4), the effect is to be understood largely in terms of reconstruction at retrieval.

1.3 Incorporation of testimonial information

The problem with which this paper is concerned arises due not to the constructive character of episodic memory in general but rather specifically to the incorporation, through constructive processes, of information about an event received via testimony after an agent's experience of the event—the problem arises, in other words, due to the sort of incorporation of non-experiential information into episodic memory responsible for the misinformation effect studied by Loftus and other eyewitness memory

researchers Loftus (1979/1996).⁶ Note that, while the misinformation effect occurs when inaccurate information is incorporated, obviously either accurate or inaccurate information can be incorporated into memory.⁷

While I focus on reconstructive retrieval, which allows the misinformation effect to occur, constructive encoding appears to give rise to a problem analogous to that discussed here, for it allows a reversed misinformation effect to occur (Lindsay and Johnson 1989; Rantzen and Markham 1992; Holliday and Hayes 2002). While the misinformation effect occurs when post-event information is incorporated into memory, the reversed misinformation effect occurs when pre-event information is incorporated. It is possible that the reversed misinformation effect can be explained in part in terms of reconstructive retrieval: information received prior to the experience of the relevant event is stored, does not affect the encoding of the experiential record, but then is combined with the experiential information about the event at retrieval. But it is likely that the reversed misinformation effect is to be explained at least in part in terms of constructive encoding: information received prior to the experience of the event affects the encoding of the experiential record itself. I suspect that my argument can be extended to cover the reversed misinformation effect, but I will not attempt to show this here.

The various available conceptions of testimony⁸ range from narrow to broad (Fricker 1995; Gelfert 2009; Michaelian 2010); I will adopt an unusually broad conception here. On my conception, testimony is a matter of any serious communication by an agent, either by means of an assertion or by other means. E.g., the content might be communicated not by means of an assertoric utterance but by means of a gesture or by means of the presupposition of a question. Adopting this unusually broad conception of testimony allows us to make contact smoothly with the eyewitness memory literature, since researchers in that area are concerned not only with misinformation that is explicitly asserted but also with misinformation that is conveyed by means of presuppositions of questions, etc.

1.4 Harmful incorporation

I will refer to the incorporation of inaccurate post-event testimonial information as “harmful incorporation”. In a case of harmful incorporation, the incorporation of testimonial information results in the retrieval of an inaccurate memory and thus, if the memory is endorsed, in the formation of a false memory belief; this gives rise to the misinformation effect, “the impairment of memory for the past that arises after exposure to misleading information” (Loftus 2005, p. 361).

⁶ I do not deal here with incorporation of information received from other external sources—my argument depends on claims about deception in direct human communication and thus cannot straightforwardly be extended to non-testimonial information.

⁷ There will also be cases in which some accurate and some inaccurate information is incorporated; for the sake of simplicity, I set such cases aside here, considering only cases in which either accurate or inaccurate information, but not both, is incorporated.

⁸ Note that I use the term here in the philosophical sense, rather than in a psychological or legal sense (“eyewitness testimony”, “expert testimony”, etc.).

In a typical experiment on the misinformation effect, the agent sees a complex event; next, she receives inaccurate information about some aspect of the event; finally, she is asked to remember the event; the misinformation effect occurs when the agent reports having seen what the misinformation describes—a yield sign is replaced with a stop sign, a screwdriver is replaced with a hammer, etc. (Loftus 2005).⁹ One plausible account (see Ayers and Reder 1998 for alternative explanations) of the mechanism responsible for harmful incorporation is provided by the source monitoring framework (Lindsay 1994)¹⁰: roughly, traces originating both in the experience of an event and in testimony about the event can coexist in memory; at retrieval, source monitoring processes can misclassify these traces as originating in experience; this results in the incorporation of testimonial information into memory and thus into the memory belief formed if the memory is endorsed (which it normally will be, since it seems to originate in experience). Though I will assume that something like the source monitoring explanation is right, the exact nature of the mechanism responsible for the misinformation effect does not matter for present purposes—it is consistent with my argument that, due to reconsolidation, the receipt of the testimonial information results directly in the modification or replacement of the stored experiential information.

Three additional points about incorporation. First: In order for incorporation of inaccurate information to result in an impairment of memory, the agent must originally have had either an accurate representation of the relevant aspect of the event or no representation of that aspect of the event. If incorporated inaccurate information replaces another inaccurate representation of the relevant aspect of the event, no impairment occurs—it makes no difference to the overall accuracy of memory. Second: In at least some cases, the accuracy of the agent’s memory itself is impaired, rather than only her verbal reports about what she remembers (Loftus 2005; Zaragoza and Lane 1994). I will focus on cases in which memory itself is impaired, that is, on cases in which harmful incorporation actually occurs. Third: In most cases, the agent is unaware that she is incorporating testimonial information (Loftus 2005). The misinformation effect can also occur when the agent deliberately chooses to accept testimonial information. But such cases seem importantly different from the central cases in which incorporation happens without awareness,¹¹ and I therefore set them aside here.

2 The contamination view

The overall epistemic status of beliefs produced by harmful incorporation is clear: since they are false, they are straightforwardly not knowledgeable. The interesting epistemological question to which the incorporation of testimonial information into memory via reconstructive retrieval gives rise concerns rather beliefs produced by

⁹ As Loftus points out, while it is easier to modify peripheral details, whole memories for events that did not occur can be implanted under certain conditions.

¹⁰ One virtue of the source monitoring explanation is that it has an easier time accounting for the reversed misinformation effect than does Loftus’ overwriting account (Holliday and Hayes 2002).

¹¹ In fact, they do not seem to be cases of memory belief at all; see the discussion of delayed formation of testimonial belief in Sect. 2 below.

what we can refer to as “helpful incorporation”, where helpful incorporation occurs in the same way as harmful incorporation, except that accurate rather than inaccurate information is incorporated, so that a true belief is formed if the resulting memory is endorsed. The question on which I focus here concerns the epistemic status of such beliefs: they are (I will argue) both true and, typically, justified, but are they knowledgeable?

In this section, I set out and motivate the natural response to this question.

2.1 Incorporation as contamination

The cases of helpful incorporation on which I will focus satisfy the restrictions given above for harmful incorporation: the testimonial information is really incorporated into memory (rather than merely affecting the agent’s verbal reports), and this normally happens without the agent’s awareness. Helpful incorporation can result in an improvement to memory only given that the agent stores either an inaccurate representation of the relevant aspect of the event or no representation of that aspect.

Helpful incorporation need be no more complicated or artificial than harmful incorporation: an agent can see a complex event but misperceive or fail to perceive an aspect of it; later, she might receive accurate information about that aspect of the event; if this post-event information is incorporated, retrieval will produce an accurate memory and thus, if the memory is endorsed, a true memory belief. (E.g.: the agent misperceives a stop sign as a yield sign but is later asked a question with the presupposition that the sign was a stop sign; when she later recalls the event, she remembers a stop sign.) Just as a reversed misinformation effect can occur due to the reception of inaccurate pre-event information, an analogous effect can occur due to the reception of accurate pre-event information; I have set the reversed misinformation effect aside, and I will not discuss this sort of effect here.

I take the natural view on the epistemic status of beliefs resulting from incorporation of testimonial information to be what I will refer to as the “contamination view”: the incorporation (without awareness) of testimonial information into a memory representation prevents the belief resulting from acceptance of the representation from qualifying as knowledge. The contamination view covers both harmful incorporation and helpful incorporation. With respect to harmful incorporation, the view is right but uninteresting: as noted above, beliefs resulting from harmful incorporation fail to qualify as knowledge simply because they are false. It is less obvious whether the contamination view is right with respect to helpful incorporation, but it is plausible that it is—something seems to go wrong, epistemically speaking, in cases of helpful incorporation.

I pause to point out that it is important not to confuse beliefs resulting from incorporation with testimonial beliefs, for the differences between testimonial beliefs and memory beliefs resulting from incorporation means that one natural way of arguing against the contamination view is ineffective. While agents clearly can acquire knowledge by accepting information received via testimony, reductionists and anti-reductionists differ over whether agents enjoy a default entitlement to accept such information (Michaelian 2008); if we confuse beliefs resulting from incorporation

with testimonial beliefs, it will appear that a standard anti-reductionist position in the epistemology of testimony implies that, in virtue of this default entitlement, incorporation is straightforwardly compatible with knowledge.

According to anti-reductionism, agents enjoy a default entitlement to accept information received via testimony. Consider, e.g., Burge's "acceptance principle": "A person is entitled to accept as true something presented as true and that is intelligible to him, unless there are stronger reasons not to do so" (Burge 1993, p. 467). Advocates of such views might be tempted to argue against the contamination view as follows: since the agent has a default entitlement to accept testimonial information, nothing need go wrong, epistemically speaking, when she accepts such information without being aware that she is doing so—she is simply exercising her entitlement; thus, as long as the information at issue in a case of helpful incorporation is accurate, incorporation can, contra the contamination view, give rise to knowledge.

In Sect. 1, I distinguished cases in which the agent consciously chooses to accept remembered testimonial information rather than remembered experiential information—this can occur because no relevant remembered experiential information is accessible or because she takes herself to have reason to trust testimonial information over accessible experiential information—from the cases of interest here, cases in which incorporation happens without the agent's awareness. Indeed, cases of the former sort appear not to be cases in which what is ultimately produced is a memory belief but, rather, cases in which memory plays an intermediary role in the delayed formation of a testimonial belief. In a case of incorporation without awareness, a testifier *S* tells the agent that *P*; the content later retrieved from memory is *P*; the agent accepts the retrieved content, forming a (memory) belief that *P*. Contrast a case in which the subject consciously chooses to accept remembered testimonial formation: a testifier *S* tells the agent that *P*; the content later retrieved from memory is *S* said that *P*; the agent accepts the retrieved content, forming a memory belief that *S* said that *P*; the agent then accepts the content of *S*'s testimony, forming the testimonial belief that *P*. While, from an experimental point of view, it will often be difficult to distinguish between cases of these two types, delayed formation of testimonial belief is clearly psychologically distinct from incorporation of testimonial information into memory without awareness; and this psychological difference grounds an epistemic difference.

Burge remarks that "[t]he justification of the Acceptance Principle says that one is entitled to accept intelligible contents "presented as true." We must perceive a speech act as involving a presentation-as-true in order to be justified under the principle" (Burge 1993, p. 481); this means, I take it, that, in order to be entitled to accept testimony, the agent must accept it because it is presented as true (though Burge makes it clear that she need not consciously judge that it is being presented as true). This condition is satisfied in ordinary cases of formation of testimonial belief, and it is satisfied in cases of delayed formation of testimonial belief, but it is not satisfied in cases of incorporation: the agent does not accept the testimonial content because it is presented as true by the speaker—when she accepts the content, no trace of the presentation-as-true by the speaker is accessed (and indeed, in many cases, no trace of the presentation by the speaker will remain in memory at all).

Nor can we argue that the endorsement of the testimonial content in an incorporation case is epistemically legitimate because the acceptance principle applies directly to contents delivered by the agent's own memory. The idea here would be that an agent has a default entitlement to accept something that is presented to her as true, regardless of whether the source is internal or external. The problem is that contents retrieved from memory are not presented as true—as noted above, a decision about the probable truth of a retrieved representation is always necessary, based on features of the representation itself or of the retrieval process.

2.2 The justification condition

The contamination view is compatible with the justification of beliefs resulting from incorporation—the claim is not that such beliefs must be unjustified. Nor is it plausible to attempt to account for the intuition that something goes wrong, epistemically speaking, in cases of incorporation by claiming that the justification condition is not met. Obviously, we should not say that beliefs resulting from incorporation are always justified, as there will be cases in which the agent has defeaters for the relevant belief, just as there are cases of incorporation-free memory beliefs in which the agent has defeaters; I set these cases aside here, to focus on typical cases, in which the agent has no reason to suspect her memory of misleading her.

One might suspect that, in cases of belief resulting from incorporation, the agent does have defeaters: suppose that *S* remembers that there was a stop sign at the scene of the accident and accepts this representation, forming the belief that there was a stop sign at the scene; and suppose that her memory that there was a stop sign at the scene resulted from incorporation; then it might seem that there is a true proposition such that, if *S* were to believe it, her justification for believing that there was a stop sign at the scene would be undermined, namely, that her memory that there was a stop sign at the scene resulted from incorporation. But if my argument below is right, this proposition does not in fact defeat *S*'s justification: if the argument is right, there is nothing wrong, epistemically speaking, with forming a memory belief by endorsing a representation resulting from incorporation. I will therefore take it that, in a typical incorporation case, the agent's justification is not defeated. Note, moreover, that, if the specific proposition that the relevant memory results from incorporation were to defeat the justification for a belief resulting from incorporation, so, presumably, would the more general proposition that the memory results from a constructive process; and since all episodic memories result from constructive processes, we should not take the latter proposition to defeat justification if we want to allow that there are any justified episodic memory beliefs.

Both externalists and internalists will agree that ordinary episodic memory beliefs are justified, though they will tell different stories about the nature of the justification for such belief. On either story, however, if we grant that there is justification in non-incorporation cases, we should also grant that there is justification in incorporation cases. (It is impossible to discuss even the main varieties of externalism and internalism here; I can only discuss representative theories.)

Externalism: According to reliabilism, the degree of justification of a belief is determined by the level of reliability of the process that produced it. Since incorporation is a routine part of the constructive retrieval process, we should classify the belief-producing process at work in incorporation cases as being the same as that at work in non-incorporation cases, that is, incorporation should be viewed an optional step in the ordinary process responsible for the production of episodic memory beliefs. It might be objected that we should distinguish between the ordinary retrieval process and the process at work in incorporation cases—that the former is reliable does not mean that the latter is also reliable. But this is an arbitrary distinction, for retrieval is always reconstructive (see Sect. 3), and reconstructive retrieval regularly involves incorporation—there is no more reason to distinguish between ordinary retrieval and retrieval with incorporation than there is to distinguish between ordinary retrieval and retrieval with any other effect of construction. Unless we are prepared to deny that we have any episodic memory knowledge, we must say that this reconstructive process is sufficiently reliable to confer a degree of justification sufficient for knowledge on the beliefs that it produces. Given reliabilism, then, we should say that beliefs resulting from helpful incorporation are typically justified.

Internalism: According to a bland sort of access internalism, the justificatory status of a belief is determined by how things stand with respect to it from the agent's point of view—it is determined wholly by internally accessible factors. On such a theory, if there is to be a difference, in terms of justificatory status, between beliefs resulting from incorporation and other episodic memory beliefs, there must be an internally accessible difference between the bases for these beliefs. But as far as internally accessible factors are concerned, incorporation cases are normally indistinguishable from non-incorporation cases: the agent is not aware that she has incorporated testimonial information; nor, in many cases, could she become aware that she has incorporated testimonial information, for, from the agent's point of view, the relevant testimonial information will typically be indistinguishable from experiential information. Again, unless we are prepared to deny that we ever have any episodic memory knowledge, we must say that the beliefs produced in non-incorporation cases are justified; and since nothing internally accessible distinguishes these beliefs from the beliefs produced in incorporation cases, we must also say that the latter beliefs are justified.

One might grant that, if memory beliefs not resulting from incorporation are justified, then so are memory beliefs resulting from incorporation, but object that I have failed to provide an account of memorial justification. However, given the structure of my argument here, which is meant to be neutral among theories of justification, I need not provide such an account. One counts memory beliefs as justified when they satisfy the general conditions for justification set out in one's preferred theory of justification. Regardless of one's preferred theory of justification, it should be the default position that memory beliefs produced without incorporation normally count as justified, since to assume otherwise is to assume, implausibly, that we have virtually no memory knowledge. The reliabilist, then, should assume that memory beliefs are typically produced by a reliable process; and similarly for partisans of other theories of justification.

Still, one might grant that the default assumption should be that memory beliefs produced without incorporation normally count as justified but object that I have not

dealt with the explanatory task of showing that such beliefs typically meet the standards for justification. Given that any reasonable set of conditions for justification will be such that most memory beliefs satisfy them (see Sect. 4.4), I need not show this here. Nor is there space to deal with this explanatory question: the question must be dealt with separately for each distinct theory of justification, and the answer to the question for a given theory will often be complex. Elsewhere [Michaelian \(2011b\)](#), I have dealt in detail with the version of the question that arises for a form of reliabilism, arguing that metamemory processes play a crucial role in ensuring the reliability of memory. Versions of the explanatory question for other forms of externalism, and for the main forms of internalism, will presumably require similarly lengthy treatments, so it is not feasible to deal with them here.

2.3 The truth condition

The contamination view is also compatible with the truth of the beliefs resulting from helpful incorporation—the claim is not that such beliefs must be false. Nor is it plausible to attempt to account for the intuition that something goes wrong in cases of helpful incorporation by claiming that the truth condition is not met. Certain views of episodic memory content do appear to imply that, when helpful incorporation occurs, it will normally produce a false belief. But there is independent reason to reject these views; and, even should one of them turn out to be right, the epistemological puzzle posed by helpful incorporation will reappear in a different form.

Various theorists have proposed views according to which the content of an episodic memory belief refers to the experience of the agent in which the remainder of the content originates; [Fernández \(2006\)](#) (see also [Fernández 2008a,b](#)), e.g., defends a view according to which the content of a memory that an agent *S* would express by saying that she remembers that *P* is an ordered pair <having had a veridical perception that *P*, *S*>, that is, e.g., something like **I saw that *P***. If something like this view is right, then my claim that the incorporation of accurate testimonial information into memory will normally result in a true belief is in trouble: given a reflexive view of episodic memory content, incorporation of accurate testimonial information will in fact normally result in a false memory belief. Suppose that an agent either failed to perceive the stop sign at the scene of an accident or misperceived it as a yield sign. Suppose that she later incorporates the information that there was a stop sign at the scene of the accident and accepts the resulting representation, thus forming a belief. If the reflexive view is right, this belief will have the content **I saw a stop sign at the scene**; this belief is false.¹²

There is, however, independent reason to reject the view. Episodic memories involve visual perspective, and, crucially, the point of view from which the agent remembers a scene is not always that from which she experienced it: while, in a field memory, the

¹² The reflexive view permits that helpful incorporation gives rise to true beliefs in unusual cases. Suppose that the agent saw the stop sign but encoded no memory for it or encoded a memory that was later lost. Suppose that she later incorporates the information that there was a stop sign. The resulting belief, with the content **I saw a stop sign at the scene**, is true. Such beliefs occur rarely enough that we can set them aside here.

agent remembers the scene from the point of view from which she originally experienced it, in an observer memory, she remembers the scene from the point of view of an external observer, so that she sees herself in the scene (Nigro and Neisser 1983; Sutton 2010). While views on which the content of an episodic memory has the form of **I saw a stop sign at the scene** rather than **there was a stop sign at the scene** are natural given that we have only field memories in mind, they are much less plausible with respect to observer memories. (Obviously, we do not want the content of an observer memory to come out having the form **I saw myself standing next to a stop sign!**) Perhaps a reflexive view capable of accommodating observer memories can be formulated, but I see no obvious way of doing so.¹³

Even if we grant that a version of the reflexive view is right, moreover, this simply introduces an additional complication into my argument. Even if episodic memory beliefs always have the form **I saw a stop sign at the scene**, so that beliefs resulting from incorporation are normally false, this does not mean that the embedded content **there was a stop sign at the scene** need be false or luckily true. In cases of helpful incorporation, the embedded content will be true, and there is thus a question about the epistemic status of this part of the belief; my argument in Sect. 4, if it succeeds, shows that the accuracy of this part of the belief is not in general due to epistemic luck.

2.4 The anti-luck condition

Even granted that memory beliefs resulting from helpful incorporation are true and justified, the claim that they are not knowledgeable remains intuitively plausible—there seems clearly to be something epistemically defective about such beliefs. I take it that this is due to the fact that they appear to be only luckily true. Virtually all epistemologists hold—and this accords well with our pretheoretic intuitions—that a justified belief fails to qualify as knowledge if it is only luckily true. If beliefs resulting from helpful incorporation appear to be only luckily true, this will account for the plausibility of the contamination view. Indeed, I will take the fact that beliefs resulting from helpful incorporation appear to be luckily true to be the main source of support for the contamination view; I will therefore take it that, if the appearance can be shown to be false, we have no reason remaining to endorse the contamination view.

Note that the claim that the contamination view can be motivated by the appearance that helpful incorporation involves epistemic luck does not presuppose the epistemic theory of memory. According to the epistemic theory (Ryle 1949; Malcolm 1977), memory is retained knowledge. One might argue on the basis of this theory that helpful incorporation is incompatible with memory: helpful incorporation involves epistemic luck; knowledge is incompatible with epistemic luck; memory is retained knowledge; so memory is incompatible with helpful incorporation. But the issue here is whether helpful incorporation is compatible with knowledge, not whether it is com-

¹³ The reflexive theorist might argue that a field memory and an observer memory of the same scene can have the same content—that they simply represent different ways of remembering the same experience of the scene. But this move is implausible, given that a field memory and a related observer memory seem clearly to have different (though overlapping) contents—they contain different information.

patible with memory: on the contamination view, helpful incorporation, because it involves epistemic luck, is incompatible with knowledge, whether or not it is compatible with memory. Depending on one's theory of memory, one might accept my argument against the contamination view while admitting that beliefs resulting from helpful incorporation are not memory beliefs (that their contents are not memories): e.g., as I point out below (Sect. 4.4), a proponent of a traditional version of the causal theory of memory (Martin and Deutscher 1996) will have to deny that beliefs resulting from helpful incorporation are memory beliefs but can still accept that they are knowledgeable.¹⁴

Note that, by the same token, we cannot decide for or against the contamination view by charging that it draws its intuitive appeal from a naïve conception of memory as non-constructive or non-inferential. While it is true that the conception of memory as purely preservative which predominates in philosophy is empirically inadequate (Michaelian 2011d), replacing the naïve conception with an empirically adequate conception of memory as involving construction does not by itself rule out the contamination view. The view concerns knowledge, rather than memory, and it might be that, while memory is compatible with the constructive processes giving rise to helpful incorporation, knowledge is not.

Before developing the suggestion that the appearance that helpful incorporation involves epistemic luck accounts for the appeal of the contamination view, I deal with a remaining dialectical difficulty. One might suspect that there is a simpler way of accounting for the intuition that there is something defective, epistemically speaking, about beliefs resulting from helpful incorporation than to suppose that it is driven by the suspicion that such beliefs are lucky. In many contexts (including legal contexts) in which we are prepared to rely on the memory of another agent, we are prepared to do so on the condition that her memory is “really hers”—that is, that she really witnessed what she claims to remember; if we learn that she is merely repeating what she has been told by another agent, whether or not she is doing so unknowingly, we withdraw or significantly decrease our trust in her report. Thus one might suspect that our reluctance to evaluate beliefs resulting from helpful incorporation as knowledgeable simply reflects the privileged status assigned in our ordinary practices to memories stemming from the agent's own experiences.

The problem with the suggestion that we can account for the intuitive plausibility of the contamination view by appealing to our ordinary practices, a suggestion reflected in the rhetoric of eyewitness memory researchers, is that the contamination view is meant precisely to reflect these ordinary practices; pointing out that, in practice, we are reluctant to trust memories resulting from incorporation thus does not explain the appeal of the contamination view. The suggestion that the appeal of the view stems from the normally unarticulated suspicion that beliefs resulting from helpful incorporation are lucky, in contrast, explains the appeal of the view; and it thus provides us with a way of deciding whether the view is to be retained or rejected.

¹⁴ While construction, including incorporation, does pose problems for the causal theory of memory, there is reason to prefer the causal theory over the epistemic theory of memory: in addition to the problems pointed out by Bernecker (2007), there is the basic problem that a definition of memory in normative terms fails to make contact with the descriptive conception of memory used in psychology.

Beliefs resulting from helpful incorporation do indeed appear to be luckily true. In a case of helpful incorporation, the testimonial information that the agent incorporates happens to be accurate. But she does not incorporate that information because she knows (or even believes) that it is accurate—indeed, she is not even aware that she is incorporating the information. Due to this lack of awareness, if she were to have received inaccurate information rather than the accurate information that she in fact received, she would likely have incorporated that information instead—if she had not received accurate testimony, she might easily have ended up with a false belief rather than a true belief. It thus seems to be due to luck that she ends up with a true belief rather than a false belief.

Indeed (though nothing in my argument turns on this), beliefs resulting from helpful incorporation appear to be Gettiered.¹⁵ Consider Russell's example of the stopped clock (Russell 1948): The agent looks at a clock. She knows that the clock is normally reliable, has no reason to doubt its accuracy on this occasion, etc. The clock says that it is 9:00; she accepts this information and so comes to believe that it is 9:00. The process that she used was reliable, nothing is amiss in internalist terms—everything is in order, epistemically speaking. Or nearly everything: the clock is in fact stopped, and it is not 9:00; due to bad epistemic luck, the agent has formed a false belief. Now modify the case so that the agent happens to look at the clock at 9:00 (which changes nothing in terms of her justification). In the modified case, good epistemic luck counteracts bad epistemic luck, and she forms a true belief. But because the belief is true only due to luck, it does not qualify as knowledge.¹⁶

Cases of helpful incorporation appear to have a similar structure: The agent retrieves a representation from her episodic memory. She knows that her memory is normally reliable, has no reason to doubt its accuracy on this occasion, etc. According to the representation, there was a stop sign at the scene of the accident that she witnessed; she accepts this information and so comes to believe that there was a stop sign at the scene. The process that she used was reliable, nothing is amiss in internalist terms—everything is in order, epistemically speaking. Or nearly everything: in fact, the retrieved record is the result of incorporating false testimonial information (there was no stop sign at the scene); due to bad epistemic luck, the agent has formed a false belief. Now modify the case so that there was a stop sign at the scene, i.e., so that the incorporated information is accurate (which changes nothing in terms of her justification). In the modified case, good epistemic luck counteracts bad epistemic luck, and the agent forms a true belief. But because the belief is true only due to luck, it does not qualify as knowledge.

3 Sceptical implications of the contamination view

Having described and motivated the contamination view in Sect. 2, I will, in this section, show that, despite its intuitive plausibility, the view has a highly implausible

¹⁵ I draw here on Zagzebski's analysis of Gettier cases (Zagzebski 1994).

¹⁶ Note that reliability is compatible with epistemic luck: a belief might be formed by a process that is reliable (in the sense that it tends to produce mostly true beliefs) but nevertheless be true only due to luck, since, absent luck, the process would have produced a false belief on the occasion in question.

sceptical consequence: it implies that we have far less episodic memory knowledge than we ordinarily take ourselves to have, that we lack such knowledge in many cases in which we would ordinarily take ourselves to have it.

Given that the contamination view implies the sceptical conclusion, we must either accept the conclusion or reject the view. Though it is possible in principle to accept the conclusion, this move should be viewed as a last resort: while it is not impossible that the conclusion is true, its implausibility means that it is to be resisted; we should be highly reluctant to accept it. It is therefore preferable to reject the contamination view, despite its intuitive plausibility: the cost of rejecting the contamination view, assuming that we have good grounds for doing so, is lower than that of accepting the sceptical consequence. In Sect. 4, I show that we have independent reason to reject the contamination view; thus, since doing so allows us to preserve a plausible position on the extent of episodic memory knowledge, we should reject it.

3.1 Reconstruction and reliance on testimony

The contamination view does not imply the sceptical conclusion on its own but rather in conjunction with two highly general empirical observations, one about the extent of reconstruction in episodic memory, and one about the extent of our reliance on testimony.

The first observation is simply that reconstruction in episodic memory is extremely pervasive: retrieval from episodic memory does not merely occasionally involve reconstruction; retrieval rather always involves reconstruction and thus routinely draws on information from sources other than experience. Though I cannot offer a serious defence of this claim here, I pause to offer an example (drawing on [Sutton 2010](#)) of the sort of evidence that supports it; this evidence should suggest to epistemologists who are used to conceiving of memory as essentially preservative that they are relying on an inadequate conception of memory.

As noted in Sect. 2, an agent can remember a scene from either field or observer perspective. It is tempting to think of field memories as privileged, even to think of field memories alone as being genuine memories, for, in an obvious way, observer memories fail to be true to the agent's original experience: since observer memories depict the agent herself, they are obvious reconstructions, departing significantly from the agent's original experience; it is natural to view field memories, on the other hand, as being more or less well-preserved copies of the agent's original experience. But this temptation is to be resisted, for, while there is indeed a greater divergence between the original experience and the retrieved memory in the case of observer memories, this does not mean that field memories are any less the products of reconstruction. That field memories are, like observer memories, products of reconstruction is demonstrated by the ability of agents to switch between field and observer perspectives within a single episode of remembering ([Rice and Rubin 2009](#)):

Field perspectives are in the same boat as those involving observer perspectives with regard to the involvement of active, constructive processes, and can of course also turn out to be false. Many or most personal memories are accessible from either field or observer perspective, and ...the perspective can often be

switched: this suggests that the difference in perspective is one of form rather than content, and that the same underlying (complex and distributed) representations can animate occurrent memories involving either perspective. (Sutton 2010)

The possibility of field/observer perspective switching demonstrates that stored information does not determine perspective. The perspective of the memory is rather a product of reconstructive retrieval: just as, in the case of an observer memory, the system makes its best guess as to what the scene must have looked like from the relevant observer perspective, in the case of a field memory, the system makes its best guess as to what the scene must have looked like from the agent's perspective as she experienced it. The upshot is that even memories which appear merely to preserve the original experience are themselves reconstructed.

The second empirical observation is mundane: agents constantly receive testimony from other agents about an extremely wide range of topics, including events that they themselves have experienced; and it is not rare that this testimony is accurate. I simply note here that I see no plausible way of rejecting this observation.

The pervasiveness of construction, together with the frequency with which testimony is received, suggests that incorporation cannot be rare: given that retrieval normally involves drawing on whatever stored information is accessible in order to produce a best guess about the relevant event, and given that much testimonial information is stored, the representations produced by retrieval will often incorporate testimonial information—many perfectly ordinary cases of episodic remembering involve incorporation. Thus, if the contamination view is right, in many ordinary cases of episodic remembering, cases in which we would ordinarily take ourselves to have knowledge, we lack knowledge. Assuming, in particular, that received testimony is fairly often accurate, in many ordinary cases of episodic remembering, we lack knowledge due to helpful incorporation.

While this implication does not threaten all episodic memory knowledge, it does reduce the extent of episodic memory knowledge significantly. Obviously, if our episodic memory beliefs turn out, due to harmful incorporation, to be false more often than we ordinarily take them to be, we have no choice but to accept this reduction in the extent of our knowledge. The interesting question here concerns cases of helpful incorporation—here, there is room to manoeuvre, since the relevant beliefs are true.

3.2 Attempts to avoid the sceptical conclusion

One way of avoiding the sceptical conclusion is to modify the contamination view so that it refers directly to the luckiness of beliefs resulting from incorporation rather than to their knowledgeability, while adopting an unorthodox view on the compatibility of knowledge and epistemic luck: first, rather than saying that true beliefs resulting from incorporation are epistemically defective in that they are not knowledgeable, say that they are defective in that they are only luckily true; second, adopt something like Hetherington's unorthodox view that knowledge is compatible with epistemic luck (Hetherington 1999).

Making these two moves might allow the defender of the contamination view to avoid the sceptical conclusion: she can maintain that, while beliefs resulting from helpful incorporation are epistemically defective in that they are only luckily true, such beliefs are, since they are justified, nevertheless knowledgeable and that thus the contamination view does not threaten to significantly reduce the extent of our episodic memory knowledge, despite the fact that helpful incorporation is frequent. The problem with this strategy is simply that, while Hetherington's view on the compatibility of knowledge and epistemic luck is coherent, it represents a radical departure from the consensus among epistemologists, and it is highly counterintuitive; the strategy should thus appeal only to those already convinced by Hetherington's argument, and I will not consider it further here.¹⁷

A more promising strategy is to reject the move from the observations that construction is pervasive and that agents frequently receive accurate testimony to the conclusion that helpful incorporation, in the relevant sense, occurs frequently. The core idea here is that, while the two observations do indeed imply that incorporation of accurate testimonial information cannot be rare, they do not imply that helpful incorporation in the relevant sense is frequent: though the ready availability of accurate testimonial information to reconstructive remembering means that incorporation of accurate testimonial information into memory beliefs will often occur, this does not yet show that it will occur without control by the agent; and ultimately it is control, rather than awareness, that matters here. The idea is that while agents are normally not aware that they are incorporating testimonial information, nonetheless the incorporation of accurate testimonial information is in most cases in an important sense under the control of the agent, for agents are implicitly sensitive to the honesty and competence of testifiers and thus are implicitly sensitive to the accuracy of received testimonial information.¹⁸

If this strategy works, then the contamination need not imply the sceptical conclusion: most cases of helpful incorporation will involve control, and the contamination view can plausibly be modified to refer to incorporation without control rather than incorporation without awareness. But the strategy is empirically untenable: agents are not sensitive to the honesty/dishonesty of testifiers,¹⁹ and thus they are not sensitive to the accuracy of received testimony. I have dealt with this question in detail elsewhere (Michaelian 2011c); here, I simply review my argument. As the work of Park and Levine and their colleagues (Levine et al. 1999, 2006; Park and Levine 2001; Vrij 2008), among others, shows, agents have no real ability to discriminate honest from dishonest testimony (that is, they are sensitive neither to honesty nor to dishonesty). It can appear that we are sensitive to honesty (though not dishonesty), since there is a high probability that, given that an item of testimony is honest, the recipient will judge that it is honest. But this is simply an effect of the "truth bias" which operates

¹⁷ The strategy also renders the contamination view significantly less interesting, since it requires that we lower our standards for knowledge.

¹⁸ This strategy is suggested by Fricker's position in the epistemology of testimony, according to which agents are entitled to accept testimony because they monitor testifiers for signs of dishonesty and incompetence (Fricker 1995).

¹⁹ I discuss sensitivity to competence below (Sect. 4).

in normal circumstances: we tend to judge that honest testimony is honest because we tend to judge that any given item of testimony is honest, regardless of actual honesty. I thus conclude that this strategy, too, does not allow the defender of the contamination view to avoid the sceptical conclusion.

4 Against the contamination view: the information effect

Since I see no way for the defender of the contamination view to avoid the sceptical conclusion, I will take it that the view implies the conclusion. In this final section of the paper, I argue that, fortunately, we have independent reason to reject the contamination view: once we adopt an adequate conception of epistemic luck, and once we adopt an adequate approach to constructive memory, it becomes clear that we should reject the view.

4.1 The modal conception of epistemic luck

While I have been discussing the status of beliefs resulting from helpful incorporation as epistemically lucky, I have so far said nothing explicit about the nature of epistemic luck—and hence nothing about how to determine whether a given belief is lucky or not, beyond relying on our intuitions.

In motivating the contamination view, however, I relied on an implicit understanding of luck in terms of control; as Pritchard (Pritchard and Smith 2004; Pritchard 2005) points out, this is a natural conception of luck (defended, e.g., by Statman 1991). In the epistemic case, this comes to something like the following: good (bad) epistemic luck occurs when an agent forms a true (false) belief, the truth (falsity) of the belief being beyond the agent's control; that is, roughly, the agent does not see to it that she forms a true belief.

While the control conception is natural, it is, as Pritchard argues, inadequate for at least two reasons. First, many events that are beyond an agent's control (e.g., the sun rising Latus 2000) clearly are not lucky, so that the control conception of luck in general is inadequate. Second, and more importantly, belief formation, in particular, is typically not (and possibly never) under the agent's control (McHugh 2011), so that the control conception of epistemic luck is inadequate, as it implies, implausibly, that virtually all true beliefs, and certainly all basic perceptual beliefs, are lucky. Since we want a unified conception of luck (so that if the control conception is inadequate with respect to perceptual beliefs, it is simply inadequate), this is not necessary for my argument here, but note that an analogous point can be made with respect to memory beliefs: formation of many memory beliefs (though perhaps not those that result from a deliberate decision by an agent to endorse the content of a retrieved memory representation) is beyond the agent's control, for much the same reason as formation of basic perceptual beliefs is beyond the agent's control—she finds herself seeming to remember that *P*, and this is already enough to produce in her a belief that *P*, just as seeming to see that *P* is sufficient to produce in her a belief that *P*.

In place of the control conception, Pritchard proposes a more plausible modal conception of luck²⁰:

- L1** If an outcome is lucky, then it is an outcome which occurs in the actual world but which does not occur in most of the nearest possible worlds to the actual world (worlds which most resemble the actual world).
- L2** If an outcome is lucky, then it is significant to the agent concerned.

Pritchard points out that, given this conception, there are a number of varieties of epistemic luck, many of them benign; the malignant variety with which we are concerned here, a variety incompatible with knowledge, is veritic epistemic luck, the sort of luck at work when “[i]t is a matter of luck that the agent’s belief is true” (Pritchard 2005, p. 146):

- L1-E** If a belief is luckily true, then the agent forms it in the actual world but instead forms a false belief in most nearby possible worlds.
- L2-E** If a belief is luckily true, then its truth is significant to the agent concerned.

I will assume that, at least in the case of memory belief, if an agent forms a belief, the truth of the belief is almost always significant to her; thus I will be concerned only with L1-E in what follows.

I will not review Pritchard’s argument for the modal conception of luck here, since I have nothing to add to it; I will simply assume that the conception is roughly right. Of course, a variety of different modal conditions on knowledge (safety, sensitivity, etc.) have been proposed by different theorists (see Comesana 2007 for an overview), and it might be worried that the focus on Pritchard’s particular anti-luck condition is unmotivated or arbitrary, that the plausibility of the contamination view can be explained in terms of the apparent violation by beliefs resulting from helpful incorporation of some other modal condition, so that I should deal here with at least a representative range of the available conditions. The worry is legitimate, but unfortunately, given the subtle differences among the conditions, space simply does not permit exploring the implications of a variety of conditions for the epistemic status of beliefs resulting from helpful incorporation. I do, however, briefly discuss Sosa’s particular proposal below (Sect. 4.4), and what I say there goes for modal conditions in general: as modal conditions become increasingly demanding, they tend to imply that helpful incorporation is incompatible with knowledge, but this should count as evidence against the more demanding modal conditions rather than as evidence against the thesis that incorporation is compatible with knowledge.²¹

Though I will not argue for the modal conception here, I note that an important virtue of the conception is that it explains the intuitive plausibility of the control conception: “if one has control over a certain event, such that one is able to (typically)

²⁰ I rely here on the version proposed by Pritchard and Smith (2004).

²¹ A general problem that affects all modal conditions on knowledge, including the condition that I employ here, is that there seems to be no context-independent way of determining closeness of possible worlds; this introduces an element of context-sensitivity to any claim that a given condition is or is not satisfied by a given belief. As there is currently no real alternative to employing modal conditions to capture the anti-luck dimension of knowledge, I bracket this problem here.

determine that a certain outcome obtains, then that is naturally understood as implying that in a wide class of relevant nearby possible worlds that outcome is realized and therefore not lucky” (Pritchard 2005, p. 130). In other words, though control is not necessary for the absence of luck, it is normally sufficient for it. This association between absence of control and luck explains why it is intuitive to suppose that beliefs resulting from helpful incorporation are luckily true (and hence not knowledgeable): in cases of helpful incorporation, the agent does not see to it that she forms a true belief, since the truth of the belief depends on the accuracy of the testimonial information that she unknowingly incorporates, that is, the truth of the belief is beyond her control; given the close association between absence of control and luck, this makes it natural to suppose that her belief is veritically lucky. The association also helps us to understand how the intuition that helpful incorporation involves epistemic luck could be mistaken: because control, though sufficient for the absence of luck, is not necessary for it, the truth of the agent’s belief might not be due to luck despite her lack of control over it. If that a belief resulting from helpful incorporation is not luckily true, the explanation for the absence of luck, since the agent lacks control, should be sought not in the agent herself but rather in her environment or in the agent-environment interaction.

Given the modal conception of epistemic luck, a given belief resulting from helpful incorporation will be luckily true only if the agent does not also form a true belief in most relevant nearby possible worlds. I will assume that the relevant nearby possible worlds are those in which the agent receives and incorporates testimony. (Obviously, I will not assume that only worlds in which the agent receives accurate testimony are relevant.) Thus whether the agent forms a true belief in most nearby worlds depends on whether she receives true testimony in most of those worlds. This point is crucial: given the modal conception, it is possible that, despite the agent’s lack of control over the formation of a true belief in incorporation cases, beliefs resulting from helpful incorporation are not luckily true; if it is not a matter of luck that the agent receives true testimony (i.e., if she receives true testimony in most nearby worlds), then the beliefs in question are not luckily true.

Recall the point made above that an absence of luck might be due not to features of the agent but rather to features of her environment or of her interaction with her environment. My suggestion is that, once we focus on the nature of the interaction between agents and their environments, including other agents, in cases of incorporation, it becomes clear that beliefs resulting from helpful incorporation are not luckily true.

4.2 Adaptive memory

This suggestion is guided by the general methodological assumption that we should adopt an adaptive perspective on memory (as on cognition in general Anderson 1990), including on the reconstructive character of retrieval. I will also assume that, in general, true belief is adaptive (McKay and Dennett 2009). While folk psychology tends to see memory as being extremely prone to failures and errors of various kinds, as a deeply flawed system in need of constant surveillance and correction, on the adaptive approach, memory (like other cognitive systems) is viewed as a system evolved to

serve the needs of the organism of which it is a component. Given this approach, when we encounter a feature of memory that initially appears to be a defect of the system, we should, rather than immediately concluding that it really is a defect, first look to see whether it functions adaptively. In particular, given that true belief is in general adaptive, when we encounter a feature of memory that initially appears to be a defect of the system, we should first look to see whether it serves the production of true beliefs by the system.

When it comes to the constructive character of episodic memory, it is important to note that, while it is often useful to study distortions resulting from construction, as these can reveal the ordinarily hidden workings of memory (Roediger 1996), construction need not be a source of inaccuracy in memory (Michaelian 2011d). Consider the case of boundary extension (Intraub et al. 1992): though boundary extension involves the agent remembering more of a scene than she actually saw, it is consistent with the accuracy of memory if the system’s “guesses” about the unseen portions of a scene are accurate. Though this point is simple and should not be controversial, it is often overlooked; it is likely a failure fully to acknowledge the point that has led researchers to focus on the misinformation effect without considering the possibility that it is a consequence of the adaptive functioning of memory.

An important insight into the adaptive function of the incorporation of testimonial information into episodic memory is provided by research on deception and deception detection. I noted in Sect. 3 that research on human deception detection shows that we are truth-biased, that we tend to evaluate testimony as being honest, regardless of its actual honesty. This truth bias (or something very like it) is already described by Reid (in a passage that has often been cited in recent work on the epistemology of testimony). According to Reid, this first bias “tallies” with another, complementary bias:

The wise and beneficent Author of nature, who intended that we should be social creatures, and that we should receive the greatest and most important part of our knowledge by the information of others, hath, for these purposes implanted in our natures two principles that tally with each other ... [We have] a propensity to speak the truth, and to use the signs of language, so as to convey our real sentiments. ... [We have] a disposition to confide in the veracity of others, and to believe what they tell us. (Reid 1764/1970, pp. 238–240)

In other words: the truth bias, in virtue of which we tend to evaluate most received testimony as honest, does not operate in isolation; we also have an “honesty bias”, in virtue of which we tend to give mostly honest testimony. Just as work on deception detection has confirmed Reid’s suggestion that we are biased in favour of accepting testimony, work on lying in everyday life is beginning to confirm his suggestion that we do not lie easily (Michaelian 2010; DePaulo et al. 1996; Levine et al. 2010).^{22, 23}

²² There is also brain-imaging work that seems to confirm Reid’s suggestion. Summarizing this work, Verschuere et al. (2010) argue that it suggests that “truth constitutes the default response of the brain, and ... lying involves intentional suppression of the predominant truth response”.

²³ Sperber (2001) and Sperber et al. (2010), appears to argue on evolutionary grounds against the existence of these biases: dishonesty is often advantageous to testifiers, and thus testimony must be effectively monitored for deception by recipients in order for communication to have evolved. But if the strategy of dishonesty results in sufficiently infrequent dishonest testimony, we cannot infer that the default strategy must

Granted the existence of the honesty bias, there emerges a straightforward view of the incorporation of testimonial information by reconstructive retrieval processes as adaptive: given the operation of the honesty bias, most (though not all—obviously, the bias can be overridden by other factors) of the testimony received by an agent will be true; and if most of the testimony received by an agent is true, then, just as the blind acceptance of testimony is a reliable process for the formation of testimonial beliefs (Michaelian 2011c), the unknowing incorporation of testimonial information into episodic memory is a reliable process for the formation of memory beliefs.

Note that this assumes that agents are largely competent about the topics about which they testify, or at least that we have some capacity to filter out incompetent testimony (since honest but incompetent testimony will normally be false). There is room to contest this assumption, clearly, and I cannot hope to offer a full defence of it here. I am optimistic, however, that it can be defended successfully, for it seems to be supported by three considerations. First, general adaptive considerations suggest that our belief-forming processes should be largely reliable—Quine’s point that “creatures inveterately wrong in their inductive tendencies have a pathetic but praiseworthy tendency to die off” applies to belief-forming processes in general. Second, there is evidence that we tend to allow testimony to influence our memories when we believe that the relevant agents are likely to be more accurate than we are (Gabbert et al. 2007). Finally, there is evidence (Mascaro and Sperber 2009) that we have some degree of sensitivity to testifiers’ competence.

Given the operation of the honesty bias, moreover, it becomes apparent that beliefs resulting from helpful incorporation are typically not lucky, despite the relevant agents’ lack of control. Given that a testifier is honesty-biased on a particular occasion of testimony, she will also testify honestly in most nearby possible worlds on that occasion—it is not a matter of luck that the testifier testifies honestly. Thus it will usually be the case that, when an agent forms a true belief due to the incorporation of accurate testimonial information, she also forms that true belief in most nearby possible worlds—it is not a matter of luck that her belief is true, despite her lack of control. The harmonious interaction between the agent’s tendency to incorporate testimonial information and an environment populated by honesty-biased testifiers rules out luck in most cases.

To emphasize this point, I return to the example of the stopped clock (discussed in Sect. 2). Given the modal conception of epistemic luck, the agent’s belief in the stopped clock case that it is 9:00 is luckily true because in many nearby possible worlds, the agent forms a false belief instead: if the clock stops at a slightly different time, the agent ends up with a false belief—it is a matter of luck that the clock tells the correct time, so it is a matter of luck that the agent forms a true belief. Typical cases of helpful incorporation are not like this: the testifier tries to tell the truth, and, since she is competent, usually succeeds—it is not a matter of luck that the testifier tells the truth, so it is not a matter of luck that the agent forms a true belief. Typical cases of helpful incorporation are, rather, like cases in which an agent forms a true belief by relying on a properly functioning clock: just as it is not normally a matter

Footnote 23 continued

be one of vigilance towards deception; and, for general reasons, communication must be predominantly honest in order for it to have evolved (Michaelian 2008; Searcy and Nowicki 2005).

of luck that the agent relies on a properly functioning clock when forming a belief about the time, it is not normally a matter of luck that the agent's memory incorporates testimonial information when forming a belief about a witnessed event; and just as it is not a matter of luck that a properly functioning clock tells the right time, it is not a matter of luck that an honesty-biased testifier provides accurate testimony.

If beliefs resulting from helpful incorporation are not usually epistemically lucky, the source of the appeal of the contamination view is undermined; I conclude that the view should be rejected.

4.3 The information effect

Once we adopt an adaptive approach to constructive memory, including the incorporation of testimonial information, the existing focus in eyewitness memory research on the misinformation effect begins to look inappropriate. If the argument given above from the operation of the honesty bias succeeds, the misinformation effect should occur relatively rarely in natural settings: the incorporation of testimonial information will typically improve memory for the past, rather than impairing it.²⁴

This suggests that the misinformation effect is better investigated in the context of a broader “information effect”, where this refers to the effect on memory of the incorporation of testimonial information, whether accurate or inaccurate. A positive information effect occurs when the incorporation of accurate testimonial information results in the production of a true memory belief that the agent otherwise would not have had, either because she would instead have formed a false belief or because she would instead have formed no belief. If accurate testimonial information is incorporated but the agent already would have formed a true belief, no effect (in terms of overall accuracy) occurs. A negative information effect occurs when the incorporation of inaccurate testimonial information results in the production of a false memory belief that the agent otherwise would not have had, either because she would instead have formed a true belief or because she would instead have formed no belief. If inaccurate testimonial information is incorporated but the agent already would have formed a false belief, no effect in terms of overall accuracy occurs.

Whether the information effect is, on the whole, epistemically beneficial depends on the relative frequency with which the positive and negative effects occur. As Fig. 1 illustrates, if the negative information effect usually occurs, then the incorporation of testimony diminishes the reliability, or both the reliability and the power, of memory; if, on the other hand, the positive information effect usually occurs, as is suggested by my argument in this section, then the incorporation of testimony improves both the reliability and the power of memory.²⁵

²⁴ There is no guarantee of this, however, since incorporation will be adaptive even if the misinformation effect occurs more frequently, as long as the payoff for helpful incorporation is sufficiently high when it does occur.

²⁵ I use “reliability” and “power” here in the epistemologists’ sense: the reliability of a belief-producing process refers to the ratio of true beliefs produced by the process to total beliefs produced by it, whereas power refers to the ratio of true beliefs to total beliefs plus failures to form a belief.

	accurate information stored	inaccurate information stored	no information stored
accurate information incorporated	no effect	positive information effect (increase in power and reliability)	positive information effect (increase in power and reliability)
inaccurate information incorporated	negative information effect (decrease in reliability and power)	no effect	negative information effect (decrease in reliability)

Fig. 1 The positive and negative information effects. Psychologists have so far studied mainly the negative information effect (the misinformation effect)

4.4 Objections

I conclude by considering two potential difficulties.

First, it might be objected that the claim that helpful incorporation can result in memory knowledge is incompatible with the causal theory of memory (Martin and Deutscher 1996)—I take the causal theory to be our best theory of memory—on which remembering requires a continuous causal connection (via a memory trace) between the original experience and the putative memory. Elsewhere, I have proposed a modified version of the causal theory, designed to take the constructive character of remembering into account (Michaelian 2011d). But even the modified causal theory requires a continuous causal connection, though it does permit the addition of new content through constructive processes, whereas in at least some cases of helpful incorporation, the relevant causal connection is absent (e.g., where the agent simply did not experience the relevant aspect of the original event); thus the claim that helpful incorporation can give rise to memory knowledge appears to be incompatible even with the causal theory of constructive memory.

There are two points to be made in response to this objection. First: My focus here has been on whether beliefs resulting from helpful incorporation are knowledge, not on whether they are memory knowledge. Thus I can in principle grant that, while beliefs resulting from helpful incorporation meet the standards for knowledge, they fail to qualify as genuine memory beliefs. Second: Another option is to modify the theory of memory even further, so that it allows remembering without a continuous causal connection via a memory trace. I will not try here to determine whether such a theory can be developed without adopting an implausible epistemic theory of memory, but it would be worthwhile to explore the possibility: helpful incorporation is an aspect of the ordinary functioning of normal episodic memory systems; there is thus good reason to count beliefs resulting from helpful incorporation as genuine memory beliefs.

Second, there is the apparently contradictory view defended in a recent article by Shanton (2011). Drawing on the literature on constructive memory (including eyewitness memory), Shanton argues that episodic memory beliefs fail to meet the standards for knowledge given by Sosa's virtue epistemology (Sosa 2007). According to Sosa's theory, a belief counts as knowledge if it is the product of an exercise of an epistemic competence, where competence is a matter of both reliability in the process sense and

security, where the security condition is satisfied on a given occasion if no false beliefs is produced in a nearby possible world on that occasion. Shanton claims that, typically, neither condition is satisfied by episodic memory beliefs: formation of beliefs by episodic retrieval is an unreliable process; and many episodic beliefs fail to satisfy the security condition. Rather than concluding that we have no episodic memory knowledge, she recommends rejecting the reliability and security conditions.

I have been assuming that episodic retrieval is a reliable belief-producing process, and I need this claim if my argument is to work given externalism, so there is a genuine tension between Shanton's view and mine here. But note that Shanton does not actually provide much reason for thinking that episodic retrieval is an unreliable belief-producing process: her argument amounts to pointing out that, due to its constructive character, episodic memory is fallible in some surprising ways, and then asking whether it strikes us as reliable. Our tendency to answer "no" this question is largely a product of our reaction to the divergence between the reality of the nature of memory, as revealed by scientific psychology, and our folk-psychological "filing cabinet" picture of memory (to use Shanton's term).²⁶ The argument depends on the sort of prejudice described above, according to which construction can only be a source of inaccuracy. Pointing out that, due to its constructive character, episodic retrieval sometimes produces inaccurate beliefs simply does not establish that episodic retrieval is an unreliable process, any more, say, than giving examples of visual illusions establishes that vision is an unreliable process. An imperfect process can be a reliable process, and there are general theoretical reasons to take formation of beliefs by episodic retrieval to be reliable. The assumption that episodic retrieval is reliable is therefore the safer bet.

With respect to security, there is no real tension between my view and Shanton's. I have said nothing so far about security but have focussed on the absence of epistemic luck. Whereas security means no false belief in a nearby possible world, absence of luck means no false beliefs in very many nearby possible worlds, a much less demanding standard. If memory beliefs fail to satisfy the security requirement (and I will not try to determine here whether they indeed do), I suggest, with Shanton, that the requirement should be abandoned. But while Shanton's argument, if successful, shows that many episodic memory beliefs, including beliefs resulting from helpful incorporation, fail to satisfy the security requirement, it does not show that they fail to satisfy the anti-luck condition. I have argued here that beliefs resulting from helpful incorporation do typically satisfy the anti-luck condition, and I suspect that the same goes for many of the other cases of construction discussed by Shanton.

Though I will not try to show this here, I appeal, again, to the view of constructive memory as adaptive. Memory is fallible; due to its constructive character, in particular, it is susceptible to various sorts of manipulation which can cause it to produce false beliefs. But that, on many occasions, the memory system could have been manipulated into producing a false belief does not mean that, on those occasions, it is a matter of luck that it produces a true belief; given the structure of the agent's environment, it

²⁶ I noted above that folk psychology tends to view memory as error-prone; this tendency coexists uneasily in folk psychology with the filing cabinet picture.

might be the case that there is no luck involved. If memory is adaptive, we have reason to think that it relates to its environment in precisely this way.

Acknowledgments Thanks for comments and discussion to two referees for *Synthese*, Luca Barlassina, Conor McHugh, Fabio Del Prete, Joëlle Proust, John Sutton, and audiences at the 2010 European Society for Philosophy and Psychology at the Ruhr-Universität Bochum, the 2010 Italian Society for Analytic Philosophy at the Università degli Studi di Padova, the 2011 Communication and Cognition conference at the Université de Neuchâtel, a meeting of the KnowJust seminar at the Institut Jean-Nicod, a meeting of the Research Colloquium in Philosophy and Cognitive Science at the Ruhr-Universität Bochum, and the philosophy department at University of Illinois (Urbana-Champaign). The preparation of this article was supported by the Agence Nationale de la Recherche, under the contract ANR-08.BLAN-0205-01.

References

- Alberini, C. M. (2005). Mechanisms of memory stabilization: are consolidation and reconsolidation similar or distinct processes?. *Trends in Neurosciences*, 28(1), 51–56.
- Anderson, J. R. (1990). *The adaptive character of thought*. Hillsdale: Erlbaum.
- Anderson, J. R. (1991). Is human cognition adaptive?. *Behavioral and Brain Sciences*, 14(3), 471–517.
- Ayers, M., & Reder, L. (1998). A theoretical review of the misinformation effect: Predictions from an activation-based memory model. *Psychonomic Bulletin & Review*, 5(1), 1–21.
- Bernecker, S. (2007). Remembering without knowing. *Australasian Journal of Philosophy*, 85(1), 137–156.
- Bernecker, S. (2010). *Memory: A philosophical study*. Oxford: Oxford University Press.
- Burge, T. (1993). Content preservation. *Philosophical Review*, 102(4), 457–488.
- Comesana, J. (2007). Knowledge and subjunctive conditionals. *Philosophy Compass*, 2(6), 781–791.
- DePaulo, B. M., Kashy, D. A., Kirkendol, S. E., Wyer, M. M., & Epstein, J. A. (1996). Lying in everyday life. *Journal of Personality and Social Psychology*, 70(5), 979–995.
- Dudai, Y. (2004). The neurobiology of consolidations, or, how stable is the engram?. *Annual Review of Psychology*, 55(1), 51–86.
- Fernández, J. (2006). The intentionality of memory. *Australasian Journal of Philosophy*, 84(1), 39–57.
- Fernández, J. (2008a). Memory and time. *Philosophical Studies*, 141(3), 333–356.
- Fernández, J. (2008b). Memory, past and self. *Synthese*, 160(1), 103–121.
- Frankish, K. (2010). Dual-process and dual-system theories of reasoning. *Philosophy Compass*, 5(10), 914–926.
- Fricker, E. (1995). Telling and trusting: Reductionism and anti-reductionism in the epistemology of testimony. *Mind*, 104(414), 393–411.
- Gabbert, F., Memon, A., & Wright, D. (2007). I saw it for longer than you: The relationship between perceived encoding duration and memory conformity. *Acta Psychologica*, 124(3), 319–331.
- Gelfert, A. (2009). Indefensible middle ground for local reductionism about testimony. *Ratio*, 22(2), 170–190.
- Hetherington, S. (1999). Knowing failably. *Journal of Philosophy*, 96(11), 565–587.
- Holliday, R. E., & Hayes, B. K. (2002). Automatic and intentional processes in children's recognition memory: The reversed misinformation effect. *Applied Cognitive Psychology*, 16(1), 1–16.
- Intraub, H., Bender, R. S., & Mangels, J. A. (1992). Looking at pictures but remembering scenes. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 18(1), 180–191.
- Johnson, M. K. (1997). Source monitoring and memory distortion. *Philosophical Transactions of the Royal Society of London Series B, Biological Sciences*, 352(1362), 1733–1745.
- Koriat, A., & Helstrup, T. (2007). Metacognitive aspects of memory. In T. Helstrup & S. Magnussen (Eds.), *Everyday memory* (pp. 251–274). London: Psychology Press.
- Latus, A. (2000). Moral and epistemic luck. *Journal of Philosophical Research*, 25, 72–149.
- Levine, T. R., Kim, R. K., & Hamel, L. M. (2010). People lie for a reason: Three experiments documenting the principle of veracity. *Communication Research Reports*, 27(4), 271–285.
- Levine, T. R., Kim, R. K., Park, H. S., & Hughes, M. (2006). Deception detection accuracy is a predictable linear function of message veracity base-rate: A formal test of park and levine's probability model. *Communication Monographs*, 73(3), 243–260.

- Levine, T. R., Park, H. S., & McCornack, S. A. (1999). Accuracy in detecting truths and lies: Documenting the “veracity effect”. *Communication Monographs*, 66(2), 125–144.
- Lindsay, D. S. (1994). Memory source monitoring and eyewitness testimony. In D. F. Ross, J. D. Read, & M. P. Toglia (Eds.), *Adult eyewitness testimony: Current trends and developments*. New York: Cambridge University Press.
- Lindsay, S. D., & Johnson, M. K. (1989). The reversed eyewitness suggestibility effect. *Bulletin of the Psychonomic Society*, 27(2), 111–113.
- Loftus, E. R. (1979/1996). *Eyewitness testimony*. Cambridge: Harvard University Press.
- Loftus, E. F. (2005). Planting misinformation in the human mind: A 30-year investigation of the malleability of memory. *Learning & Memory (Cold Spring Harbor, N.Y.)*, 12(4), 361–366.
- Malcolm, N. (1977). *Memory and mind*. Ithaca: Cornell University Press.
- Martin, C. B., & Deutscher, M. (1966). Remembering. *Philosophical Review*, 75(2), 161–196.
- Mascaro, O., & Sperber, D. (2009). The moral, epistemic, and mindreading components of children’s vigilance towards deception. *Cognition*, 112(3), 367–380.
- Matthen, M. (2010). Is memory preservation?. *Philosophical Studies*, 148(1), 3–14.
- McClelland, J. L. (1995). Constructive memory and memory distortions: A parallel-distributed processing approach. In D. L. Schacter (Ed.), *Memory distortion* (pp. 69–90). Cambridge, MA: Harvard.
- McClelland, J. L. (2011). Memory as a constructive process: The parallel-distributed processing approach. In S. Nalbantian, P. Matthews, & J. L. McClelland (Eds.), *The memory process: Neuroscientific and humanistic perspectives* (pp. 129–151). Cambridge, MA: MIT Press.
- McHugh, C. (2011). Judging as a non-voluntary action. *Philosophical Studies*, 152(2), 245–269.
- McKay, R. T., & Dennett, D. C. (2009). The evolution of disbelief. *Behavioral and Brain Sciences*, 32(06), 493–510.
- Michaelian, K. (2008). Testimony as a natural kind. *Episteme: A Journal of Social Epistemology*, 5(2), 180–202.
- Michaelian, K. (2010). In defence of gullibility: the epistemology of testimony and the psychology of deception detection. *Synthese*, 176(3), 399–427.
- Michaelian, K. (2011a). Is memory a natural kind?. *Memory Studies*, 4(2), 170–189.
- Michaelian, K. (2011b) Metacognition and endorsement. *Mind & Language* (Forthcoming).
- Michaelian, K. (2011c). The epistemology of forgetting. *Erkenntnis*, 74(3), 399–424.
- Michaelian, K. (2011d). Generative memory. *Philosophical Psychology*, 24(3), 323–342.
- 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. 95–175). Oxford: Oxford University Press.
- Nairne, J. S., & Pandeirada, J. N. S. (2008). Adaptive memory: Remembering with a stone-age brain. *Current Directions in Psychological Science*, 17(4), 239–243.
- Nigro, G., & Neisser, U. (1983). Point of view in personal memories*1. *Cognitive Psychology*, 15(4), 467–482.
- Park, H. S., & Levine, T. (2001). A probability model of accuracy in deception detection experiments. *Communication Monographs*, 68(2), 201–210.
- Pritchard, D. (2005). *Epistemic luck*. Oxford: Oxford University Press.
- Pritchard, D., & Smith, M. (2004). The psychology and philosophy of luck*1. *New Ideas in Psychology*, 22(1), 1–28.
- Rantzen, A., & Markham, R. (1992). The reversed eyewitness testimony design: More evidence for source monitoring. *Journal of General Psychology*, 119(1), 37–43.
- Reid, T. (1764/1970). *An inquiry into the human mind*. Chicago: University of Chicago Press.
- Rice, H. J., & Rubin, D. C. (2009). I can see it both ways: First- and third-person visual perspectives at retrieval. *Consciousness and Cognition*, 18(4), 877–890.
- Robin, F. (2010). Imagery and memory illusions. *Phenomenology and the Cognitive Sciences*, 9(2), 253–262.
- Roediger, H. L. (1996). Memory illusions. *Journal of Memory and Language*, 35(2), 76–100.
- Russell, B. (1948). *Human knowledge: Its scope and limits*. London: Allen & Unwin.
- Ryle, G. (1949). *The concept of mind*. London: Hutchinson.
- 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 B Biological Sciences*, 362(1481), 773–786.

- Searcy, W. R., & Nowicki, S. (2005). *The evolution of animal communication: Reliability and deception in signalling systems*. Princeton: Princeton University Press.
- Shanton, K. (2011). Memory, knowledge and epistemic competence. *Review of Philosophy and Psychology*, 2(1), 89–104.
- Sosa, E. (2007). *A virtue epistemology: Apt belief and reflective knowledge*. Oxford: Oxford University Press.
- Sperber, D. (2001). An evolutionary perspective on testimony and argumentation. *Philosophical Topics*, 29, 13–401.
- Sperber, D., Clément, F., Heintz, C., Mascaro, O., Mercier, H., & Origgi, G., et al. (2010). Epistemic vigilance. *Mind & Language*, 25(4), 359–393.
- Statman, D. (1991). Moral and epistemic luck. *Ratio*, 4(2), 146–156.
- Suddendorf, T., & Corballis, M. C. (2007). The evolution of foresight: What is mental time travel, and is it unique to humans?. *Behavioral and Brain Sciences*, 30(03), 299–313.
- Sutton, J. (1998). *Philosophy and memory traces*. Cambridge: Cambridge University Press.
- Sutton, J. (2010). Observer perspective and acented memory: Some puzzles about point of view in personal memory. *Philosophical Studies*, 148(1), 27–37.
- Tulving, E. (1982). Synergistic ephory in recall and recognition. *Canadian Journal of Psychology/Revue Canadienne de Psychologie*, 36(2), 130–147.
- Tulving, E. (1983). *Elements of episodic memory*. Oxford: Oxford University Press.
- Vandekerckhove, M., & Panksepp, J. (2009). The flow of anoetic to noetic and autoeotic consciousness: A vision of unknowing (anoetic) and knowing (noetic) consciousness in the remembrance of things past and imagined futures. *Consciousness and Cognition*, 18(4), 1018–1028.
- Verschuere, B., Spruyt, A., Meijer, E. H., & Otgaar, H. (2010). *The ease of lying*. Consciousness and Cognition, November 2010.
- Vosgerau, G. (2010). Memory and content?. *Consciousness and Cognition*, 19(3), 838–846.
- Vrij, A. (2008). *Detecting lies and deceit: Pitfalls and opportunities* (2nd ed.). New York: Wiley.
- Zagzebski, L. (1994). The inescapability of gettier problems. *The Philosophical Quarterly*, 44(174), 65–73.
- Zaragoza, M. S., & Lane, S. M. (1994). Source misattributions and the suggestibility of eyewitness memory. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 20(4), 934–945.