



Memory in two dimensions

Jordi Fernández¹ 

Received: 22 June 2023 / Accepted: 13 December 2023 / Published online: 22 January 2024
© The Author(s) 2024

Abstract

Memories can be accurate or inaccurate. They have, then, accuracy conditions. A reasonable picture of the accuracy conditions of a memory is that a memory is accurate just in case the reference of a memory satisfies the information provided by the memory. But how are the references of our memories determined exactly? And what are the accuracy conditions of memories, given their references? In this paper, I argue that the notion of accuracy conditions for memories is ambiguous. There are two types of conditions which can be plausibly construed as accuracy conditions for memories. I motivate this idea by using some resources from two-dimensional semantics. The outcome of applying two-dimensionalism to memory is that memories have two kinds of accuracy conditions. In both cases, causal relations play an important role in the framing of those conditions. But the role is quite different in each case. For one type of accuracy conditions, the causal relations which produce a memory play the role of fixing the reference of that memory. For the other type of accuracy conditions, the causal relations which produce a memory become part of the information which needs to be satisfied by the reference of the memory for it to be accurate. However, in both cases, the picture according to which a memory is accurate just in case the reference of a memory satisfies the information provided by the memory reemerges as being correct, though for interestingly different reasons.

Keywords Memory · Reference · Truth-conditions · Two-dimensionalism

✉ Jordi Fernández
jorge.fernandez@adelaide.edu.au

¹ Philosophy Department, University of Adelaide, Adelaide, SA 5005, Australia

1 Introduction

Memory provides us with information about the past in different ways. I may form a certain belief in the past, such as the belief that there is a red apple in the kitchen, and that belief may be preserved up to the present time by memory. This may happen even if, in the present, I do not have any memory experience of a red apple being in the kitchen.¹ Alternatively, I may undergo, in the present, a memory experience of a red apple being in the kitchen. And this may happen even if, in the past, I never formed the belief that there is a red apple in the kitchen. In both cases, memory is providing me with information about the presence of a red apple in the kitchen. And yet, it seems that the type of memory involved is quite different in each case. This paper is concerned with the latter type of memory; the type of memory which allows us to have memory experiences of past states of affairs. Our discussion will be focused, specifically, on the accuracy of memory experiences.²

Our memories can be evaluated as either accurate or inaccurate. This suggests that our memories have accuracy conditions, that is, conditions relative to which they are accurate, and conditions relative to which they are inaccurate.³ Consider, now, the objects which constitute the subject matter of our memories; the objects with which they are concerned. We may call those objects, the ‘references’ of our memories. It seems reasonable to think that the accuracy conditions of our memories must involve, in some way, their references. After all, one would think that if a memory is accurate, it is because the information which is provided by the memory applies to the object which constitutes the subject matter of that memory. To be sure, this is only a rough picture of the accuracy conditions of our memories. But, rough as it is, this picture seems to be, *prima facie*, a promising starting point. And yet, specifying the details of this picture is remarkably challenging, for two reasons. First of all, it is not obvious what relation determines, for every memory, the reference of that memory. And, relatedly, it is also unclear what kind of accuracy conditions should be attributed to a memory given what its reference is, and given how that reference has been determined. The aim of this paper is to address both of these issues, thus clarifying the intentionality of memory.

My main tenet in this paper will be that the notion of accuracy conditions for memories is ambiguous. My contention will be that there are two kinds of conditions which can be reasonably construed as accuracy conditions for our memories. In order to draw the relevant distinction, I will borrow some resources from a certain semantics framework, namely, ‘two-dimensional semantics’ or ‘two-dimensionalism’. Two-

¹ While discussing both memory and perception, I will talk about experiences being ‘of’ some states of affairs. This kind of talk is only meant to abbreviate that, if prompted, the subject would report their experience by saying that they remember, or that they perceive, the relevant state of affairs.

² For the sake of brevity, I will refer to those experiences which are generated by memory as ‘memories’. Hopefully the use of the term ‘memory’ to refer both to a faculty and to the deliverances of that faculty will cause no confusion.

³ In what follows, I will use the umbrella term ‘satisfaction conditions’ to refer to the conditions relative to which a mental state is satisfied, and the narrower terms ‘accuracy conditions’ to refer to the satisfaction conditions of perceptual experiences and memories, and ‘truth conditions’ to refer to the satisfaction conditions of beliefs.

dimensionalism was originally devised to deal with the semantics of indexicals.⁴ The main idea in two-dimensionalist frameworks is that we should pull apart, for any utterance or belief, the context where it occurs from the context where it is to be evaluated. In discussions regarding the semantics of indexicals, several distinctions about linguistic content have been drawn as a result of separating those two contexts. Thus, Robert Stalnaker makes a distinction between two propositions associated with an utterance; the proposition expressed by it, and the ‘diagonal proposition’ which corresponds to it.⁵ Along similar lines, John Perry makes a distinction between the proposition expressed by an utterance, and the proposition ‘created’ by it.⁶ David Chalmers also applies a distinction of this type to mental (as opposed to linguistic) content, and separates what he calls the ‘subjunctive’ intension of a thought from its ‘epistemic’ intension.⁷ I will argue that a similar distinction between two contexts must be drawn when we evaluate, not an utterance or a belief, but a memory. Once the relevant distinction is drawn, two types of accuracy conditions for memories arise. As we will see, the two types of conditions are similar in that, in both cases, the causal relations which produce a memory play a prominent role in the accuracy conditions of that memory. But the two types of conditions are also different, I will suggest, in that the relevant role is, in one case, that of securing the reference of the memory whereas, in the other case, it is that of providing the information which needs to be satisfied by the reference of the memory (for the memory to be accurate, that is). I will argue that the distinction between the two kinds of accuracy conditions for memories is not artificial, or unnecessarily technical. For both conditions which result from the application of two-dimensionalism to memory seem to satisfy some of the desiderata that we have for an explanatorily useful notion of accuracy conditions.

I will proceed as follows. In Sect. 2, I will highlight some facts about memory which, arguably, should be illuminated by an account of accuracy conditions for memories. I will distinguish three such facts: The specificity of our memories, the communicability of our memories, and the capacity of our memories to guide our behaviour. In Sect. 3, I will introduce a two-dimensionalist framework, and use it to disambiguate the notion of accuracy conditions for memories. The resulting outcome will be a distinction between what I will call the ‘objective content’ of a memory and its ‘subjective content’. In Sect. 4, I will discuss the sources of the accuracy conditions of our memories. Perhaps unsurprisingly, the sources of the objective contents of our memories will turn out to be different from those of their subjective contents. Finally, in Sect. 5, I will argue that the specificity of our memories and their communicability are best explained by using the notion of objective content, whereas the capacity of our memories to guide our behaviour is best explained by using the notion of subjective content. Accordingly, I will conclude that the two types of conditions

⁴ See (Kaplan, 1979) for an influential piece in this line of research. See (Chalmers, 1995) and (Davies & Humberstone, 1980) for two-dimensionalist approaches to modality.

⁵ In (Stalnaker, 1978) and (Stalnaker, 1981).

⁶ See (Perry, 1988).

⁷ In (Chalmers, 2002). Hereafter, the content of a mental state, such as a memory, should be understood, very sparsely, as the satisfaction conditions of that mental state.

which result from applying two-dimensionalism to memory have a legitimate claim to being the accuracy conditions of our memories.

2 The significance of accuracy conditions

What exactly hinges on the nature of accuracy conditions for memories, and on the issue of how those conditions are determined? It seems that, ideally, we would like the notion of accuracy conditions for memories to play a helpful role in the explanation of, at least, three different facts about memory. Consider, first of all, the fact that our memories are about specific objects.⁸ Imagine that, in my kitchen, there are two apples; apple A and apple B. They are identical, except for the fact that apple A is red and apple B is green. Imagine that, at different times in the past, I have perceptually experienced A and I have perceptually experienced B. Suppose, now, that I have a memory of a red apple being in the kitchen; a memory which, as a matter of fact, causally originates in my having perceptually experienced A. In this scenario, my memory seems to concern, not just an apple which was in the kitchen in the past, but a particular apple that I perceptually experienced, and whose perceptual experience is responsible for my memory. In other words, it seems to concern apple A. To bolster this intuition, suppose that, after leaving my kitchen, and while having a memory of a red apple being there, I try to refer, demonstratively, to the apple that I seem to remember. I could do this, for example, by forming, on the basis of my memory, a belief that I would normally express with an utterance of the kind ‘that apple was red.’ Alternatively, I could do this by making an actual utterance of the kind ‘that apple was red’ while attending to my memory. Intuitively enough, in those two scenarios, my belief and my utterance are not empty, or nonsensical. On the contrary, they seem to be the kind of belief and the kind of utterance which can be evaluated as true or false. This suggests that the (either linguistic or conceptual) component in them which corresponds to ‘that apple’ has succeeded in singling out a particular object; an object with which my belief and my utterance are concerned. And it seems that the object in question is apple A. After all, both my belief and my utterance appear to be correct. And they would not be correct if they were concerned with apple B.

Our memories seem to have, then, the property of enabling us to refer, in speech and in thought, to specific objects in our past.⁹ One would think that a natural way of explaining this property of memories is by appealing to their accuracy conditions. After all, it seems reasonable to assume that if our memories allow us to form beliefs, and to make utterances, about particular objects in the past, it is because those

⁸ To be clear, I do not mean to imply that our memories are exclusively about objects, that is, that they are about objects without being about either states of affairs or events which involve those objects. For a discussion of the view that memories can sometimes be about objects in that fundamental way, see (Openshaw, 2022). For the purposes of this discussion, I will assume that our memories are always about states of affairs. I will also assume that our memories are about particulars, such as objects and events, and about properties, to the extent that they are about states of affairs which consist in the relevant particulars having the relevant properties. As far as I can see, nothing in our discussion of the accuracy conditions of memories which follows depends on those assumptions.

⁹ For a discussion of the role of memory in demonstrative thought, see (Campbell, 2002, 177–193).

memories themselves are, in some sense, about particular objects in the past. And the notion of accuracy conditions provides us with a helpful way of spelling out what the relevant sense of ‘about’ is. Contributing to an account of the specificity of memories which allows for demonstrative reference to the past seems to be, then, a desideratum that we should keep in mind while we try to clarify the nature of accuracy conditions for memories, and the issue of how those conditions are determined.

Furthermore, consider the fact that our memories can be communicated to other subjects. Suppose, for example, that you and I are wondering whether an apple that we both saw in my kitchen (which happens to be apple A) was red or it was green. I have a memory of a red apple being in the kitchen which originates in my having perceptually experienced A and, to settle the issue, I utter, based on my memory, ‘that apple was red.’ Intuitively enough, it seems that this utterance is expressing what I seem to remember. And it also seems that, when you hear my utterance, you will be able to understand the utterance (assuming, that is, that you are a competent English speaker). In other words, it seems that, by making a certain utterance, I have the capacity to communicate to you what I seem to remember in virtue of having one of my memories. One would think that an account of the communicability of memories will need to appeal, too, to their accuracy conditions. After all, it seems reasonable to suppose that if, when I make an utterance of the kind ‘that apple was red’, communication between us requires that you understand my utterance, then, as a result of hearing my utterance, you must come to think, or entertain, what I am saying by making that utterance. And it also seems reasonable to suppose that if, by making an utterance of the kind ‘that apple was red’, I am expressing one of my memories, then what I am saying, by making that utterance, must be what I seem to remember in virtue of having the memory in question. It seems, then, that an explanation of what other subjects come to think, or entertain, based on my utterances when I communicate my memories to them is going to depend on an account of what I seem to remember by having those memories. And the notion of accuracy conditions provides us with a helpful way of spelling out the idea of ‘what I seem to remember’ in virtue of having a memory. This suggests that there is another desideratum that we should bear in mind while we try to clarify the nature, and the sources, of the accuracy conditions of our memories, namely, contributing to an explanation of the fact that we can communicate what we seem to remember, in virtue of having memories, to other people by making certain utterances.

Finally, consider the fact that our memories can guide our behaviour. Imagine, for example, that, in the recent past, I have perceptually experienced the presence of apple A in the kitchen, and I now have a memory which originates in that perceptual experience. Suppose, furthermore, that I have the desire to eat a red apple. (I may believe, for example, that red apples are less sour than green apples, and I may have a preference for that kind of taste at the moment.) So I head towards the kitchen. Intuitively enough, it seems that the reason why I head towards the kitchen, and not in some other direction, has something to do with the memory that I am having at that moment. It appears, in other words, that what I seem to remember is part of what causally explains my behaviour at that time. In that sense, my memory seems to be guiding my behaviour. Now, it seems natural to think that the accuracy conditions of our memories should also shed some light on this feature of our memories. For,

on the face of it, it would seem that our memories can only inform our behaviour while trying to find, for example, some food that we have perceived in the past, if our memories are about that food, and about where that food was perceived to be located in the past. And, once again, memories having certain accuracy conditions seems to be a reasonable way of spelling out what the relevant type of ‘aboutness’ consists in. It seems, then, that contributing to an account of the capacity of memory to guide our behaviour is yet another desideratum which is worth keeping in mind while we try to clarify what the accuracy conditions of memories are, and how they are determined.

These three desiderata for a notion of accuracy conditions for memories provide us with a helpful guide while we try to think about the nature, and the sources, of those conditions. Ideally, a notion of accuracy conditions for memories should satisfy all three of the desiderata that we have discussed. And, accordingly, competing notions of accuracy conditions should be evaluated, at least in part, based on how many of these desiderata they can satisfy. What are, then, the accuracy conditions of memories? And how should we think about the sources of those conditions for the purposes of illuminating the three facts about memory highlighted above? Let us turn to this issue next.

3 Two dimensions of accuracy

The basic idea about memories having accuracy conditions is a normative idea. It is the idea that, for each memory, there are conditions relative to which it is evaluated as being accurate, and conditions relative to which it is evaluated as being inaccurate. It seems reasonable to think, then, that if you want to know what the accuracy conditions of a memory are, you should ask yourself what it would take for the memory to be accurate. But, as we are about to see, the expression ‘what it would take for a memory to be accurate’ can be read in two different ways. This ambiguity gives rise to an ambiguity in the very notion of accuracy conditions for memories.

In order to draw a distinction between two ways of understanding the idea of what it would take for a memory to be accurate, it will be helpful to construe the accuracy conditions of memories as propositions.¹⁰ And, furthermore, it will be helpful to construe propositions as sets of possible worlds.¹¹ Consider one of my memories; the memory of a red apple being in the kitchen. Suppose that this memory originates in my having had a perceptual experience of a red apple being in the kitchen. The apple

¹⁰ When I say that the truth conditions of a memory can be construed as some proposition *p*, I simply mean that we can express the accuracy conditions of the memory by saying that the memory in question is accurate just in case *p*. (Why not ‘just in case *p* is true’? I am thinking here of propositions as facts, and not as the primary bearers of truth-value.)

¹¹ In what follows, then, I will assume a possible-worlds conception of propositions. Thus, for any proposition *p*, the proposition that *p* will be conceived as the set of possible worlds in which *p* (Lewis, 1986). This conception, however, is not essential to the proposal about two types of accuracy conditions for memories offered in this section. The distinction between the objective content and the subjective content of a memory can be easily reformulated in terms of, for example, structured propositions (Salmon, 1986; Soames, 1985); see note 15 below.

in question happens to be apple A. Let us consider, now, what it would take for my memory to be accurate.

One of the things that we may be enquiring when we ask ‘what it would take’ for my memory to be accurate is the following: What are the possible worlds such that if I had had my memory there, it would have been accurate, or true, *in* those worlds? What we are asking, then, is what possible worlds are such that, given the conditions in which my memory is occurring in those worlds, my memory turns out to be true in those worlds? (In this case, we are ignoring the conditions under which my memory has actually happened. We are focusing on the conditions under which my memory could have happened. We are asking ‘would my memory have been true if it had occurred under those other conditions?’) These seem to be the worlds in which my memory originates in my having perceived a red apple in the kitchen, whether that red apple was apple A or not. We may individuate, then, one kind of accuracy conditions for my memory by using the proposition which contains only those worlds. Call this kind of accuracy conditions, the ‘subjective content’ of my memory.

Alternatively, by asking what it would take for my memory to be accurate, we may be enquiring the following: What are the possible worlds such that, given the circumstances in which my memory has actually occurred, my memory is accurate, or true, *of* those worlds? Or, perhaps more naturally: What are the possible worlds which are accurately represented by my memory, given the circumstances in which it has actually occurred? (In this case, we are ignoring the alternative conditions under which my memory could have occurred. We are focusing on the conditions under which my memory actually happened. We are fixing those, and asking which possible worlds are well represented by my memory given that it occurred under those conditions.) These seem to be the worlds in which apple A is red and A is in the kitchen, whether my memory originates in a perception of A or not. We may individuate, then, another kind of accuracy conditions for my memory by using the proposition which contains only those worlds. Call this kind of accuracy conditions, the ‘objective content’ of my memory.

The distinction between evaluating memories as being true in possible worlds, as opposed to true of possible worlds, is different from the distinction between evaluating memories as being true, as opposed to authentic. Sven Bernecker introduced the latter distinction by pointing out that, sometimes, one may want to evaluate a memory because one is interested in whether the memory corresponds to, or is in accord with, the past state of the world. And, sometimes, one may want to evaluate a memory because one is interested in whether the memory corresponds to, or is in accord with, the subject’s initial experience of the world. To mark this distinction, one could use, as Bernecker suggests, the term ‘truth’ to refer to the relation of accordance between the subject’s memory and the world, and the term ‘authenticity’ to refer to the relation of accordance between the subject’s memory and their initial experience of the world.¹² It is important to notice that the distinction between evaluating memories as being true in possible worlds, as opposed to true of possible worlds, does not map onto the distinction between evaluating them as being authentic, as opposed to being

¹² See (Bernecker, 2009, 38–40). For some interesting uses of this distinction with regards to the phenomenon of observer memory, see (McCarroll, 2018) and (Michaelian and Sant’Anna 2022).

true. Whether one evaluates a memory as being true or authentic, in Bernecker's terms, one is evaluating the memory relative to the actual situation, the situation in which it occurs. One may choose to focus on the correspondence between the subject's memory and the world, or between the subject's memory and their initial experience of the world. But, in either case, one is focusing on the situation in which the memory actually occurs. (One is either focusing on the past state of the world in that situation, or on the subject's initial experience of it in that situation.) By contrast, if one evaluates a memory as being true of a certain possible world, one may or may not be evaluating the memory relative to a situation in which the memory occurs.

Instead, the distinction between evaluating memories as being true in possible worlds, as opposed to true of possible worlds, is analogous to a distinction between two ways of evaluating an utterance of a sentence containing an indexical expression; a sentence such as 'it is cold here.' Suppose that an utterance of 'it is cold here' is made in some spatial location X. In one sense, the utterance is true just in case that token of 'here' has been uttered at a place where it is cold, whether that place is X or not. In another sense, the utterance is true just in case it is cold at X, whether a token of 'here' is uttered at X or not. Since 'it is cold here' could be uttered at a cold place other than X, and X could be cold even if 'here' is not uttered at X, these turn out to be two different kinds of truth-conditions for the relevant utterance of 'it is cold here.' Similarly, the objective and subjective content of my memory of a red apple being in the kitchen turn out to be different as well. To illustrate the difference between them, consider three possible situations.

Let us start with W_0 . In W_0 , there are two apples in my kitchen, apple A and apple B. In W_0 , A is red and B is green. At some point, I enter my kitchen, I look at apple A and I perceptually experience A (call the perceptual experience that I have when I look at apple A in the kitchen, ' P_A '). Also, at some point, I enter my kitchen, I look at apple B and I perceptually experience B (call the perceptual experience that I have when I look at apple B in the kitchen, ' P_B '). Furthermore, in W_0 , P_A is a perceptual experience of a red apple being in the kitchen, and P_B is a perceptual experience of a green apple being in the kitchen. Later on, and as a result of having had P_A , I have a memory of a red apple being in the kitchen. Call this memory 'M'. Basically, W_0 is the possible situation that we have been considering thus far.

Consider, now, a different possible situation in which, even though I have M, and even though M is a memory of a red apple being in the kitchen, something has gone wrong in the production of M. In W_1 , M originates in P_A , but P_A is a perceptual experience of a green apple being in the kitchen. However, my having P_A is not an episode of misperception in W_1 because, as it happens, apple A is green in W_1 . What about apple B in this alternative situation? Let us stipulate that, in W_1 , B is green, I have P_B , and P_B is a perceptual experience of a green apple being in the kitchen.

Finally, consider a different possible situation W_2 . Situation W_2 is, in a way, the mirror image of W_0 . In what way exactly? In W_2 , apple A is green and apple B is red. I enter my kitchen, I look at apple A, and I have perceptual experience P_A . Perceptual experience P_A is, in W_2 , a perceptual experience of a green apple being in the kitchen. I look at apple B and I have perceptual experience P_B . Perceptual experience P_B is, in W_2 , a perceptual experience of a red apple being in the kitchen. Later on, I have

memory M , and M is still a memory of a red apple being in the kitchen. But memory M , in this situation, does not originate in P_A , but in P_B .

Let us now ask ourselves: What are the accuracy conditions of memory M occurring in W_0 ? What would it take for M to be accurate? If, with these questions, we are enquiring about the subjective content of my memory, then the relevant conditions are that M originates in a red apple in the kitchen which is perceived by me as such (whether that apple is A or not).¹³ The accuracy conditions of my memory M , occurring in W_0 , can then be captured by a proposition which includes W_0 and W_2 , but excludes W_1 . If, by contrast, we are enquiring about the objective content of my memory, then the relevant conditions are that apple A is red and that apple A is located in the kitchen (whether a perception of those facts gives rise to memory M or not). The accuracy conditions of my memory M , occurring in W_0 , can then be captured by a proposition which includes W_0 , and excludes W_1 and W_2 . We can represent graphically the two propositions which constitute the two types of accuracy conditions for M if we borrow Robert Stalnaker's 'diagonalisation' matrix.¹⁴

	W_0	W_1	W_2
W_0	T	F	F
W_1	T	F	F
W_2	F	F	T

The horizontal rows of the matrix represent the different propositions that constitute the objective contents of memory M in each of the possible situations that we have considered (in each context of occurrence, as it were). The top row, then, represents the objective content of memory M occurring in W_0 . By contrast, the subjective content of memory M , occurring in W_0 , is constituted by the proposition represented by the diagonal which runs from the top left corner of the matrix to its bottom right corner. Since the subjective content of my memory M , occurring in W_0 , includes W_2 , but its objective content does not, the subjective content of my memory turns out to be different from its objective content.¹⁵

¹³ One might think that the relevant conditions should be, more specifically, that M originates in a red apple in the kitchen which is perceived by me in virtue of having P_A . But this cannot be right if we are enquiring about the subjective content of M in W_0 . After all, W_2 seems to be one of the situations such that, if M had happened in that situation, it would have been true in that situation. And yet, in W_2 , M does not originate in a red apple in the kitchen which is perceived by me in virtue of having P_A .

¹⁴ This way of representing the distinction is slightly misleading in that, in (Stalnaker, 1978) and (Stalnaker, 1981), Stalnaker is not trying to capture a difference between two types of mental content with this matrix. He is only concerned with linguistic content.

¹⁵ What would the objective and subjective content of M , occurring in W_0 , look like if we construed propositions in a Russellian way, as ordered pairs of properties and objects, and not as sets of possible worlds? The objective content of M in W_0 would be the ordered pair of the property of being a red apple in the kitchen and object A , <being a red apple in the kitchen, A > (as opposed to the set of possible worlds in which A is a red apple in the kitchen, $\{W: \text{In } W, A \text{ is a red apple in the kitchen}\}$). The subjective content of M in W_0 would be the ordered pair of the property of having been caused by a perception of a red apple

4 The sources of accuracy conditions

How are the two types of accuracy conditions for memories determined? Take the case of memory M happening in W_0 , for instance. What makes it the case that memory M has certain accuracy conditions? This question needs to be answered differently, depending on whether the relevant accuracy conditions are those which constitute the objective content of M , or the subjective content of M .

The objective content of M seems to be determined by the causal relations which are responsible for the production of that memory. The fact that, in W_0 , my memory originates in P_A , and the fact that P_A has been caused by the presence of apple A in the kitchen, seem to be the reasons why the conditions under which M is accurate, in the objective sense, are that apple A , and not some other apple such as apple B , is red and is in the kitchen. The relation which fixes the reference of my memory M , occurring in W_0 , seems to be, then, the causal process which has led to M in W_0 . And, in W_0 , that process goes through P_A , and originates in the presence of apple A in the kitchen. That is why the reference of M , when it occurs in W_0 , is apple A . And that is why what matters for the accuracy of M , in the objective sense, is whether apple A in the kitchen was red or not.

To highlight the role that causation plays in determining the reference of memories, as far as objective content is concerned, consider what would have happened if my memory M had originated in the presence of a different apple in front of me. What would have happened is that its objective content would have been different. This is because the contribution that the reference of M makes to the relevant type of accuracy conditions would have been different. Situation W_2 illustrates this point well. Since, in W_2 , the causal process which has led to M goes through P_B and originates in apple B , this causal process picks up apple B , and not apple A , as the reference of my memory M in W_2 . As a result, what matters for the accuracy of M in W_2 , in the objective sense, is whether apple B in the kitchen was red or not. Hence, the occurrence of M in W_2 has an objective content which includes W_2 , and excludes W_1 and W_0 , whereas the occurrence of M in W_0 has an objective content which includes W_0 , and excludes W_1 and W_2 .

Let us go back, at this point, to the broad picture of accuracy conditions with which we started our discussion. That broad picture can now be spelled out in an informative way when it comes to the objective dimension of accuracy. The picture, let us recall, is that a memory is accurate if the information which is provided by the memory applies to the object which constitutes the reference of that memory. The outcome of our discussion thus far is that this intuitive picture can now be reaffirmed as being correct. And the reason why it is correct concerns causation. But, when the relevant sense of 'accurate' is the objective sense, the role that causation plays in this picture needs to be understood in a particular way. The thought is not that a memory is accurate if the information provided by the memory, which includes information about its own causal origin, is satisfied by the reference of that memory. Causal rela-

in the kitchen, and memory M , <having been caused by a perception of a red apple in the kitchen, M > (as opposed to the set of possible worlds in which M has been caused by a perception of a red apple in the kitchen, $\{W: \text{In } W, M \text{ has been caused by a perception of a red apple in the kitchen}\}$).

tions are not part of the information which is provided by our memories within this picture of the accuracy conditions of our memories; not when the relevant conditions amount to objective content. Instead, causal relations only play the role of selecting the objects which constitute the references of our memories.

Thus, it seems that, in the end, our original picture of the accuracy conditions of memories needs to be spelled out as follows, when the relevant conditions amount to objective content: A memory is accurate just in case the object which constitutes the reference of a memory (the object at the end of the causal history of the memory) satisfies the information provided by the memory. And the information provided by the memory is the information that the object in question had certain properties in the past (in the case that we have been discussing, for example, those of being a red apple and being located in the kitchen), and not that the memory itself has any particular property.

This picture of the sources of objective content is reminiscent of the ‘causal theory of reference’ for linguistic expressions such as proper names and indexical terms.¹⁶ It is also friendly to the so-called ‘causal theory of memory.’ According to the causal theory of memory, a mental state qualifies as a memory of some fact if it causally originates in the subject’s perceptual experience of that fact.¹⁷ This is a metaphysical view, a view about what it is for a mental state to qualify as a memory. The picture sketched above, by contrast, is a semantic view. But it is friendly to the causal theory of memory in that, if the causal theory of memory is correct, then it is guaranteed that there is such a thing as the perceptual experience in which my memory *M*, occurring in W_0 , causally originates. (Otherwise, *M* would not be a memory.) And, on the proposed view, the relation which fixes the reference of my memory *M*, as far as objective content is concerned, is the causal process which has led to *M* in W_0 . Thus, the proposed way of understanding our original picture of the accuracy conditions of memories, when those amount to objective content, seems quite appealing if the metaphysics of memory is that proposed by the causal theory of memory.

Things look quite different on other conceptions of the metaphysics of memory. Consider, for example, the view according to which memory is ‘reconstructive,’ in the sense that our memories are formed by integrating information from different sources, such as our past perceptual experience, inference, testimony and the imagination.¹⁸ On this view, my memory *M*, in W_0 , does not need to causally originate in a perceptual experience of a red apple in order for it to qualify as a memory. Instead, *M* may be a mental image which has been formed by assembling together, let us say, a causally unrelated perception of an apple at some point in the past, the recent testimony that there is a red piece of fruit in the house, and my inference that the kitchen would be the most likely place for fruit to be found in the house. On the reconstructive view of memory, there is no reason to rule out *M* as a memory. But, on

¹⁶ One can find elements of this theory, for example, in (Kaplan, 1979) and (Donnellan, 1972), and a defence of it in (Kripke, 1980).

¹⁷ On the causal theory of memory, see (Martin & Deutscher, 1966).

¹⁸ This picture of the metaphysics of memory has become dominant in psychology (Dale et al., 1978; Loftus & Palmer, 1974; Loftus & Pickrell, 1995) and is becoming increasingly popular in philosophy as well (Michaelian, 2016; De Brigard, 2014; Sutton, 2007).

the reconstructive view of memory, there is no such thing as the causal process which has led to my having M, or at least there does not need to be. There may be multiple causal processes which have led to my having M. In the just-mentioned example, for instance, there is a causal chain of events which leads from my having perceived an apple in the past to my having M, a partly overlapping chain of events which leads from my having received testimony about the presence of red fruit in the kitchen to my having M, and a partly overlapping chain of events which leads from my having made a certain inference to the best explanation to my having M. The causal process which leads to my having M, then, looks more like a tree with multiple branches than a straight line. How congenial is this picture of the metaphysics of memory to the proposed way of understanding our original picture about the accuracy conditions of memories, when those amount to objective content?

Consider what a constructivist could say about our original picture of the accuracy conditions of memories, as far as objective content is concerned. This is the picture according to which, if a memory is accurate, it is because the information which is provided by the memory applies to the object which constitutes the reference of that memory. As far as I can see, the constructivist could say one of three things. They may, first of all, adopt the picture as being correct. They may grant, that is, that there is such a thing as the accuracy conditions of memories, and that the accuracy of memories depends on their references. As noted above, the constructivist cannot appeal to a single causal process as a reference-fixing mechanism for memories. Thus, the picture of the sources of objective content proposed above is, in this scenario, incompatible with constructivism about memories. The constructivist will need some other reference-fixing mechanism to select the object to which the information provided by the memory is supposed to apply (in the M example, a reference-fixing mechanism which singles out the object which needs to have been a red apple located in the kitchen in order for M to be accurate).¹⁹

Alternatively, the constructivist may take the original picture according to which, if a memory is accurate, it is because the information which is provided by the memory applies to the object which constitutes the reference of that memory, as being broadly, but not completely, correct. They may grant that there is such a thing as the accuracy conditions of memories, and that the accuracy of memories depends on their references. But they may also claim that reference is not an all-or-nothing affair. They may claim that there is partial reference of a memory to multiple objects. Accordingly, they may suggest, accuracy comes in degrees as well, and a memory should be conceived as being false in some respects and true in other respects, depending on which partial reference of that memory to an object we are focusing on.²⁰ The picture of the sources of objective content proposed above is, in this scenario, compatible with constructivism about memories. One would simply need to introduce the notion of ‘partial reference’ of a memory to a number of objects, depending on how many

¹⁹ Kourken Michaelian, for example, accepts that memories are accurate or inaccurate within a constructivist framework in his (2013). It is unclear to me, however, whether the accuracy of memories can, on his view, be explained in the terms of the references of those memories, and what properties those references need to satisfy.

²⁰ Some constructivists speak of a ‘degree of falsity’ in memory (Bernstein and Loftus 2009, 373). That kind of talk is consistent with the constructivist position that I am envisaging here.

causal processes lead from those objects to the memory at issue, and a corresponding notion of ‘degree of accuracy’, where the degree of accuracy of a memory will be a function of the degree to which the memory refers to an object which satisfies the information provided by it.²¹

Finally, the constructivist may take the original picture according to which, if a memory is accurate, it is because the information which is provided by the memory applies to the object which constitutes the reference of that memory, as being incorrect. They may reject that there is such a thing as accuracy in memory and, for that reason, they may reject the need to explain the accuracy conditions of memories in terms of their references.²² This constructivist position allows for memories to be successful in a practical sense. It allows for memories being useful, or serving the subject’s purposes well. But it does not allow for memories to be successful in a representational sense. For that reason, this position is incompatible, not only with the picture of the sources of objective content proposed above, but with any picture of the accuracy conditions of our memories, whether those accuracy conditions amount to objective content or not.

Let us now turn to the subjective content of memory M in W_0 . The subjective content of M seems to be determined by its phenomenology. When I have M in W_0 , the phenomenology that I experience does not appear to be neutral on the question of what I happened to perceive in the past, and it does not appear to be neutral on the question of where the memory that I am having comes from either.²³ On the contrary, the phenomenology of having M seems to provide me with answers to those two questions, rightly or wrongly. For example, it seems to be part of what it is like for me to have M that, in virtue of having M , it thereby seems to me that I have perceived (not perceptually experienced, but veridically perceived) a red apple being in the kitchen. Similarly, it seems to be part of what it is like for me to have M that, in virtue of having M , it thereby seems to me that the very memory that I am having originates in a perception of a red apple being in the kitchen. More broadly, it seems to be part of what it is like for me to have M that, in virtue of having M , I am presented with M itself, and with what appears to be its causal history. This complex phenomenology seems to be the reason why the conditions under which M is accurate, in the

²¹ The causal theory of names provides us, once again, with examples for how to do this. See, for instance, Michael Devitt’s development of the causal theory of names in terms of partial designation and degrees of truth in (Devitt, 1974).

²² Stanley Klein, for example, claims that ‘there is no principled reason for episodic recollection to adhere to any particular degree of fidelity to the past’ and that ‘the memory content served up to consciousness need not entail “precision of match” to past events as a criterion of success’ (Klein, 2014, 438–439).

²³ The use of ‘in the past’ here could lead to confusion. There is a view according to which the phenomenology of our memories involves a ‘feeling of pastness’ (Russell, 1921, 161–162). On that view, it is part of what it is like for me to have M that, in virtue of having M , it thereby seems to be that, in the past, there was a red apple in the kitchen. If, by ‘in the past’, we mean, in the terminology from (Lewis, 1976), ‘in the objective past,’ then I am sceptical that the phenomenology of memory involves a feeling of pastness. Consider a situation in which, after seeing the apple in the kitchen, I travel back in time to a moment before the red apple was in the kitchen. Suppose that, then, I have M . We would not want to say that the phenomenology involved in my having M is misleading in any way. And yet, the red apple is not in the kitchen in the past; not in the objective sense. If, by ‘in the past,’ we mean ‘in my personal past,’ then I agree that it is part of the phenomenology of having M that I have the feeling that there was a red apple in the kitchen in the past. This, I take it, follows from the second aspect of the phenomenology of having M mentioned above.

subjective sense, are that M has been caused by my having perceived a red apple in the kitchen. In this case, the relation which fixes the reference of the occurrence of memory M in W_0 seems to be the relation of being presented to the subject of M. In W_0 , my having memory M brings with it a certain phenomenology; a phenomenology in virtue of which M itself is presented to me. It is presented to me as having a certain causal history. That is why the reference of M in W_0 , when it comes to subjective content, is memory M itself. And that is why what matters for the accuracy of M, in the subjective sense, is whether M has a certain causal history or not.

To highlight the role that phenomenology plays in determining the reference of memories, as far as subjective content is concerned, consider what would have happened if my memory M had originated in the presence of an apple in the kitchen which is different, but qualitatively indistinguishable, from apple A. What would have happened is that its subjective content would not have been different. And the reason why it would not have been different is that the phenomenology of having M would still have presented M to me. In fact, it would have presented M to me as having the same causal history. Thus, the contribution that the reference of M would have made to the relevant type of accuracy conditions would have been the same. Once again, situation W_2 illustrates this point well. In W_2 , my having memory M brings with it a certain phenomenology which involves the memory that I am having. And the phenomenology at issue is the same as the phenomenology of having M in W_0 , even if the apple in which my memory originates is a different apple.²⁴ In both cases, part of what it is like for me to have my memory is to be presented with the fact that my memory originates in a perception of a red apple in the kitchen. And, in W_2 , the memory that I am having when I undergo this phenomenology is, still, memory M. Thus, as far as subjective content is concerned, my act of remembering picks up memory M itself as the reference of M in W_2 . As a result, what matters for the accuracy of M in W_2 , in the subjective sense, is whether M originates in a perception of a red apple in the kitchen or not. Hence, the occurrence of M in W_2 has a subjective content which includes W_0 and W_2 , and excludes W_1 , just like the occurrence of M in W_0 does.

We can now flesh out the broad picture of accuracy conditions for memories at the start of this discussion with regards to the subjective dimension of accuracy as well. The picture according to which a memory is accurate if the information provided by the memory applies to the object which constitutes the reference of that memory remains a cogent picture. And the reason why it remains a cogent picture has to do, once again, with causation. But, when the relevant sense of ‘accurate’ is the subjec-

²⁴ This claim is not true on some notions of ‘phenomenology’ and ‘experience.’ According to Dorothea Debus, for instance, a subject’s conscious experience while remembering an object O_A is different from their experience while remembering a different but qualitatively indistinguishable object O_B , just in virtue of the fact that O_A and O_B are different objects (2008, 421–422). I, by contrast, am individuating the subject’s phenomenology, or experience, while having a memory less finely than Debus is. I am assuming that if the subject themselves would not be able to discriminate between their experience while recollecting O_A and their experience while recollecting O_B , then the subject’s phenomenology of recollection is the same in both cases. The claim that the subjective content of a memory is grounded on the subject’s experience while having that memory turns out not to be true if we individuate experiences as finely as Debus does. But so does the claim that, while having a memory, a subject has privileged access to what their own experience is like.

tive sense, the role that causation plays in this picture needs to be understood in a different way. The thought is now that a memory is accurate if the information provided by the memory, which includes information about its own causal history, is satisfied by the reference of that memory. Causal relations become part of the information which is provided by our memories in this picture of their accuracy conditions. They no longer play the role of selecting the object which constitutes the reference of a memory. That role is reserved for the phenomenology of remembering.

Thus, it seems that, in the end, our original picture of the accuracy conditions of memories needs to be spelled out as follows, when the relevant conditions amount to subjective content: A memory is accurate just in case the object which constitutes the reference of a memory (that is, the memory itself) satisfies the information provided by the memory. And the information provided by the memory is the information that the memory in question has a certain causal history (in the case that we have been discussing, that of originating in a perception of a red apple located in the kitchen), and not that the specific object which is, in fact, at the end of that causal process has any particular property. It seems, then, that the details of the picture according to which a memory is accurate just in case its reference satisfies the information provided by the memory need to be understood differently, depending on whether the relevant sort of accuracy is objective or subjective. Either way, though, the intuitive picture of the accuracy of memories with which we started our discussion comes out as being vindicated.

5 The reality of objective and subjective content

A possible concern about the distinction between objective and subjective content for memories needs to be addressed now. The concern is that, just because we have isolated two types of propositions that we can associate with a memory, this does not show that we have thereby isolated two semantic properties of memories. To ensure that we have not singled out the two types of propositions arbitrarily, we need to see that the types of propositions that we have isolated can be explanatorily useful. Let us turn our attention, then, back to the specificity and the communicability of our memories, as well as to their capacity to guide our behaviour.

Consider, first of all, the fact that our memories seem to have the property of enabling us to refer, in speech and in thought, to specific objects in the past. Suppose that, when I have M in W_0 , I form a belief that I would express with an utterance of the form ‘that apple was red’, or suppose that I make an utterance of that sort. It seems that my belief and my utterance will be about, specifically, apple A . If the reason why our memories allow us to form beliefs and make utterances about particular objects in the past is that our memories are themselves about particular objects in the past, then what kind of accuracy conditions for memories capture the sense in which our memories are about particular objects in the past? Is this feature of memories best captured by their having a certain objective content, or by their having a certain subjective content?

It seems that the type of accuracy conditions for memories which best captures the idea that our memories are, in some sense, about particular objects in the past is not

going to be the subjective content of our memories. Take my memory M in W_0 , for instance. The subjective content of M includes those worlds in which M originates in my having perceived a red apple in the kitchen. Which particular apple happens to be at the origin of the causal history of M does not seem to be relevant, then, for the accuracy of M ; not in the subjective sense. This is why, when I have M in W_2 , where the apple at the origin of the causal history of my memory is apple B, the subjective content of M is the same as when I have M in W_0 . And yet, if I formed, in W_2 , a belief that I would express with an utterance of the form ‘that apple was red’, or made such an utterance, my belief and my utterance would intuitively be about a different apple, namely, apple B. How can the subjective content of M be, then, what enables me to refer to different objects in W_0 and in W_2 if the subjective content of M is, in both situations, the same? The subjective content of M does not seem to be, as it were, fine-grained enough to perform that role.

The objective content of memories, by contrast, seems to be a much more plausible candidate as the feature of our memories which allows us to refer to specific objects in the past by forming beliefs and by making utterances on the basis of those memories. After all, the objective content of M in W_0 includes those worlds in which apple A, and not just any apple, is red and is in the kitchen. And the objective content of M in W_2 includes those worlds in which apple B, and not just any apple, is red and is in the kitchen. The fact that the objective contents of M in W_0 and in W_2 concern, respectively, apple A and apple B, seems to be what explains the fact that, in W_0 , I am referring to apple A when I form, on the basis of my memory, a belief that I would express with an utterance of the form ‘that apple was red’, or when I make such an utterance, whereas, in W_2 , I am referring to apple B when I form the corresponding belief, or make the corresponding utterance.

Consider now, the fact that we can communicate our memories to other subjects by making certain utterances. Take, for instance, the fact that, by making an utterance of the form ‘that apple was red’ on the basis of M , I have the capacity to communicate to you what I seem to remember in virtue of having M (assuming, that is, that you understand that utterance). Let us stipulate that the situation in which I am having M , and I am making my utterance on the basis of M , is still W_0 . It seems natural to think that the reason why I can communicate to you what I seem to remember in virtue of having M is, firstly, that what I seem to remember is what I am saying with my utterance and, secondly, that what I am saying with my utterance is what you come to think, or entertain, when you hear my utterance and understand it. But how should we construe the notion of ‘what I seem to remember’ for this explanation of the communicability of our memories to work? What kind of accuracy conditions for memories provide us with the operative notion of ‘what I seem to remember’ in this explanation?

It seems that, once more, the type of accuracy conditions for memories which best captures the idea that our memories can be communicated to other subjects by making certain utterances is not going to be the subjective content of our memories. Take my memory M in W_0 , for instance. The subjective content of M is the proposition which includes the possible worlds in which M has a certain causal history. But what I am saying, when I utter ‘that apple was red’, does not seem to concern my memory M or, for that matter, any other memory. After all, if I had uttered ‘that apple was red’

in a situation in which I do not remember the kitchen, but I am seeing a photograph of apple A in the kitchen taken in the past, then, intuitively enough, I would have been saying the same thing as what I am saying when, in W_0 , I utter ‘that apple was red’ on the basis of M. And yet, in that other situation, I would not have been saying anything about memory M, since M would not have occurred in the first place. It is hard to see, then, how what I am saying, when I utter ‘that apple was red’ on the basis of M in W_0 , can be the subjective content of M. It is also difficult to see how what you think, when you hear and understand my utterance, can be the subjective content of M. After all, in order to entertain the subjective content of a mental state, that mental state needs to be presented to its subject as part of the phenomenology of having the mental state in question. It follows, therefore, that the only subject who can entertain the subjective content of a mental state is the bearer of that mental state. And memory M is my memory, not yours. How would you be able to entertain, then, the subjective content of M when you hear and understand my utterance of the form ‘that apple was red’?

This does not mean that, when I express my memory M by uttering ‘that apple was red’, and you hear and understand my utterance, you cannot form a belief which is, in a sense, about memory M and its causal history. You may tell yourself, for example, ‘the memory on which the utterance that I am hearing is based comes from a perception of a red apple’. In that case, you are indeed referring to my memory M. You are referring to M thanks to a causal process which involves my uttering ‘that apple was red’ as a result of my having M, and involves your hearing and understanding my utterance as a result of my making the utterance. But, in that case, the truth conditions of your belief which, admittedly, concern memory M amount to the objective, and not the subjective, content of your belief. The fact remains, therefore, that the subjective content of my memory M cannot be shared between the two of us.

The objective content of memories, by contrast, seems to be a much more plausible candidate as the feature of our memories which allows us to communicate our memories to other subjects by making certain utterances on the basis of those memories. After all, the objective content of M in W_0 includes those worlds in which a particular apple in the kitchen, namely, apple A, is red. And, as we saw in our discussion of the specificity of memories, it seems quite natural to think that, when I utter ‘that apple was red’ on the basis of M, what I am saying is indeed that apple A in the kitchen was red. Thus, the intuition that what I am saying with my utterance, based on M, is what I seem to remember when I have M seems to make sense if we understand ‘what I seem to remember’ as the objective content of M. Furthermore, the intuition that what you come to think, or entertain, when you hear and understand the utterance that I am making based on M, is what I am saying with my utterance also seems to make sense if what I am saying with my utterance is the objective content of M. Since the objective content of M does not concern my memory M at all, there is no impediment for you to form a thought with that objective content, that is, a thought whose objective content is the proposition that contains those worlds in which a particular apple in the kitchen, apple A, was red.

Consider, finally, the fact that our memories can guide our behaviour. Suppose that, when I have M in W_0 , I have the desire to eat a red apple and head towards the kitchen. Intuitively enough, it seems that what I seem to remember is, at least in part, the reason why I head towards the kitchen. But how should we understand ‘what I

seem to remember' for this intuition to make sense? Is what I seem to remember in virtue of having M the objective content of M or its subjective content? It seems that the objective content of M is not going to be the feature of my memory which explains its capacity to guide my behaviour. Notice that if I had been in W_2 , as opposed to W_0 , when I have my memory M and my desire to eat a red apple, my behaviour would have been exactly the same. I would have headed towards the kitchen. And yet, as we have seen, the objective content of M is, in W_0 , the proposition that includes those worlds in which apple A in the kitchen is red, and the objective content of M is, in W_2 , the proposition that includes those worlds in which apple B in the kitchen is red. Those two propositions are different. How can the objective content of M explain, then, my behaviour in W_0 if, in W_2 , that behaviour is the same, and yet the objective content of M is different? The objective content of M seems to be, so to speak, too fine-grained for it to perform that role.

It seems that the type of accuracy conditions for memories which best captures the idea that our memories can guide our behaviour is the subjective content of our memories. The subjective content of M in W_0 is the proposition that includes the possible worlds in which M originates in a perception of a red apple being in the kitchen. Furthermore, the reason why this proposition constitutes the subjective content of M is that, in W_0 , my having M has a certain phenomenology. In virtue of having M, I am presented with the fact that M itself originates in a perception of a red apple being in the kitchen. This means that the subjective content of my memory captures the way in which things appear to me, from the first-person perspective, when I have memory M. And that seems to be the relevant perspective for the purposes of explaining my behaviour: In W_0 and W_2 , I behave in the same way when I have M and I want to eat a red apple because, quite simply, things appear the same way to me in both possible situations. In both cases, it appears to me that M originates in a perception of a red apple being in the kitchen. It is no wonder, then, that I am disposed to head towards the kitchen when I want to eat a red apple in both possible situations.

6 Conclusion

Let us take stock. We have seen that if we construe the accuracy conditions of memories as propositions, then we can distinguish two types of accuracy conditions for memories, since there are two types of propositions that we can systematically associate with our memories. We have seen, too, that the distinction is not arbitrary, since both types of conditions seem to play a helpful role in the explanation of some interesting facts about memory. What is common to both types of accuracy conditions for our memories is the prominence of causal relations in the framing of those conditions. But the role of causal relations is quite different in each case. When it comes to one type of accuracy conditions, causal relations fix the references of our memories, and our memories are not, in any sense, about themselves. But, when it comes to the other type of accuracy conditions, causal relations become part of the information which needs to be satisfied by the references of our memories in order for those memories to be accurate. And, in that case, the references of our memories

are fixed, through non-causal means, as those very memories. When it comes to this sort of accuracy, then, memories are very much about themselves.²⁵

Funding Open Access funding enabled and organized by CAUL and its Member Institutions

Declarations

Conflict of interest I hereby declare that I have no conflict of interest submitting this paper to the collection ‘Reference and Remembering’ of the journal Synthese.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article’s Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article’s Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Bernecker, S. (2009). *Memory: A philosophical study*. Oxford University Press.
- Bernstein, D. M., & Loftus, E. F. (2009). How to tell if a Particular memory is true or false. *Perspectives on Psychological Science*, 4, 370–374.
- Campbell, J. (2002). *Reference and consciousness*. Oxford University Press.
- Chalmers, D. J. (1995). *The conscious mind*. Oxford University Press.
- Chalmers, D. J. (2002). The components of content. In D. Chalmers (Ed.), *Philosophy of mind: Classical and contemporary readings* (pp. 608–633). Oxford University Press.
- Dale, P., Loftus, E., & Rathburn, L. (1978). The influence of the form of the question on the eyewitness testimony of preschool children. *Journal of Psycholinguistic Research*, 7, 269–277.
- Davies, M., & Humberstone, I. L. (1980). Two notions of necessity. *Philosophical Studies*, 38, 1–30.
- De Brigard, F. (2014). Is memory for remembering? Recollection as a form of episodic hypothetical thinking. *Synthese*, 191, 155–185.
- Debus, D. (2008). Experiencing the past: A relational account of recollective memory. *Dialectica*, 62, 405–432.
- Devitt, M. (1974). Singular terms. *Journal of Philosophy*, 74, 183–205.
- Donnellan, K. (1972). Proper names and identifying descriptions. In D. Davidson, & G. Harman (Eds.), *Semantics of Natural Language* (pp. 356–379). D. Reidel.
- Kaplan, D. (1979). Demonstratives. In J. Almog, J. Perry, & H. Wettstein (Eds.), *Themes from Kaplan* (pp. 481–563). Oxford University Press.
- Klein, S. B. (2014). Autonoesis and Belief in a personal past: An evolutionary theory of episodic memory indices. *Review of Philosophy and Psychology*, 5, 427–447.
- Kripke, S. (1980). *Naming and necessity*. Harvard University Press.
- Lewis, D. (1976). The paradoxes of time travel. *American Philosophical Quarterly*, 13, 145–152.
- Lewis, D. (1986). *On the plurality of worlds*. Blackwell Publishers.
- Loftus, E., & Palmer, J. (1974). Reconstruction of automobile destruction: An example of the interaction between language and memory. *Journal of Verbal Learning and Verbal Behavior*, 13, 585–589.
- Loftus, E., & Pickrell, J. (1995). The formation of false memories. *Psychiatric Annals*, 25, 720–725.
- Martin, C. B., & Deutscher, M. (1966). Remembering. *Philosophical Review*, 75, 161–196.

²⁵ I am very grateful to two anonymous referees for this journal for their helpful feedback on a previous version of this paper.

- McCarroll, C. (2018). *Remembering from the outside: Personal memory and the Perspectival mind*. Oxford University Press.
- Michaelian, K. (2013). The information effect: Constructive memory, testimony, and epistemic luck. *Synthese*, 190, 2429–2456.
- Michaelian, K. (2016). *Mental Time Travel: Episodic memory and our knowledge of the personal past*. MIT Press.
- Michaelian, K., & Sant’Anna, A. (2022). From authenticity to alethism: Against McCarroll on observer memory. *Phenomenology and the Cognitive Sciences*, 21, 835–856.
- Openshaw, J. (2022). Remembering objects. *Philosophers’ Imprint*, 22, 1–20.
- Perry, J. (1988). Cognitive significance and new theories of reference. *Nou’s*, 22, 1–18.
- Russell, B. (1921). *The analysis of mind*. Allen and Unwin.
- Salmon, N. (1986). *Frege’s puzzle*. MIT Press.
- Soames, S. (1985). Lost innocence. *Linguistics and Philosophy*, 8, 59–71.
- Stalnaker, R. (1978). Assertion. *Syntax and Semantics*, 9, 315–332.
- Stalnaker, R. (1981). Indexical Belief. *Synthese*, 49, 129–151.
- Sutton, J. (2007). *Philosophy and memory traces: Descartes to connectionism*. Cambridge University Press.

Publisher’s Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.