Against discontinuism: Mental time travel and our knowledge of past and future events*

Kourken Michaelian
kourken.michaelian@otago.ac.nz

Abstract

Continuists maintain that, aside from their distinct temporal orientations, episodic memory and future-oriented mental time travel (FMTT) are qualitatively continuous. Discontinuists deny this, arguing that, in addition to their distinct temporal orientations, there are qualitative metaphysical or epistemological differences between episodic memory and FMTT. This chapter defends continuism by responding both to arguments for metaphysical discontinuism, based on alleged discontinuities between episodic memory and FMTT at the causal, intentional, and phenomenological levels, and to arguments for epistemological discontinuism, based on alleged discontinuities with respect to the epistemic openness of the past and future, the directness or indirectness of our knowledge of past and future, and immunity to error through misidentification. The chapter concludes by sketching a positive argument for continuism.

Keywords: episodic memory; future-oriented mental time travel; causal theory; immunity to error through misidentification.

1 Continuism and discontinuism about mental time travel

Over the last several decades, the psychology of memory has undergone a major reorientation away from earlier conceptions of episodic memory as a specialized system dedicated to the storage and retrieval of information about the "what, when, and where" of past events and towards a conception of episodic remembering as a form of constructive mental time travel (MTT). Reinforced by impressive brain imaging evidence and extensive research on representational and phenomenological overlap between remembering the past and imagining the future (for recent reviews, see Klein 2013; Schacter

*Thanks to Dorothea Debus, Denis Perrin, André Sant'Anna, and audiences at the University of Auckland and the University of Waikato for comments.

et al. 2012; Szpunar 2010), this new conception emphasizes the deep similarities between episodic memory, future-oriented mental time travel (FMTT), and, increasingly, processes such as episodic counterfactual thinking (De Brigard 2014, De Brigard and Gessell this volume, Pezzulo this volume), in which the subject imagines alternatives to past events. Taking the new conception to its logical conclusion, many have argued that, rather than distinct faculties of episodic remembering and imagining, what we in fact have is "a single general faculty of mental time travel" (Suddendorf and Corballis, 2007). On this view, remembering the past and imagining the future are strictly continuous.

Philosophers have recently begun to take note of the turn towards MTT in psychology and, in line with older and more general philosophical suspicions about the ability of imagination to give us knowledge of future events, to push back against the associated "continuist" view. "Discontinuism", as we may refer to this emerging position, maintains that, notwithstanding the strength of the empirical support for continuism, there are fundamental metaphysical and/or epistemological differences between episodic memory and FMTT, differences in virtue of which it would be a mistake to view our capacities to remember the past and imagine the future as being merely special cases of a general capacity for MTT. This chapter aims to meet discontinuists on their own ground, providing a philosophical defence of continuism against discontinuist arguments. The chapter argues that, while treating episodic memory as a form of MTT not distinct in kind from FMTT requires us to abandon certain common philosophical views about memory and imagination, there turn out to be good philosophical reasons to abandon these views and therefore to reject the view that there are fundamental discontinuities between episodic memory and future-oriented MTT. The picture that emerges of a fully symmetrical ability to mentally travel backwards and forwards in time suggests that philosophers would do well to follow psychologists' lead, attending more closely than they have historically done to our capacity for future-directed episodic thought and to the ways in which that capacity gives rise to knowledge of future events.

1.1 Forms of continuism and discontinuism

Let us begin by fixing terms. *Discontinuists* maintain that, in addition to their distinct temporal orientations, there are further important differences between episodic memory and FMTT and that, in virtue of these further differences, episodic memory and FMTT amount to qualitatively different processes or states. As Debus puts it, "[episodic memories] of past events and [episodic imaginations] of future events are ultimately mental occurrences of two different kinds" (2014, 337).[1] On the discontinuist view, it is a mistake

[1]Debus formulates her view in terms of "recollective memory" and "sensory imagination of future events", but these terms map on to "episodic memory" and "future-oriented

to see MTT as a single, unified capacity which can be directed at either the past or the future. *Continuists*, in contrast, maintain that MTT is precisely such a unified capacity. On the continuist view, temporal orientation is the only important difference between episodic memory and FMTT.

Three further distinctions — between extreme and moderate, narrow and broad, and metaphysical and epistemological forms of (dis)continuism — will help to structure the following discussion.

1.1.1 Extreme and moderate discontinuism

Perrin (who introduced the terms "continuism" and "discontinuism" in his chapter in this volume) distinguishes between extreme and moderate forms of (dis)continuism. *Extreme continuism* denies the existence of any difference, beyond subjective temporal orientation, between episodic memory and FMTT. *Moderate continuism*, in turn, acknowledges that there may be limited differences between them but denies that these are sufficiently important to give rise to a difference in kind between the processes. *Extreme discontinuism* maintains that the differences between episodic memory and FMTT are so dramatic that they represent completely different kinds of process, incapable of being subsumed under a common superordinate kind. *Moderate discontinuism*, in turn, acknowledges that they may have enough features in common for us to classify them as belonging to a common kind but insists that the differences between them are sufficiently important for them to constitute distinct kinds of process within that broader kind.

We are thus faced with a spectrum of possible positions, ranging from extreme continuism to extreme discontinuism. Which of these positions should we take seriously? As Perrin suggests, someone who holds a simple "replay" view of memory, combined with a view of FMTT as a constructive or simulational process, might endorse extreme discontinuism. It is unlikely that anyone now holds such a combination of views, and we may thus set extreme discontinuism aside. Similarly, disregarding occasional rhetorical excesses, it is unlikely that anyone subscribes to extreme continuism, since the fact that there are a variety of differences, at the neural, representational, and phenomenological levels, between remembering the past and imagining the future is now well established; we may thus set extreme continuism aside. The real debate, then, is between moderate continuism and moderate discontinuism: everyone will agree that there are some important similarities and some important dissimilarities between episodic memory and FMTT; the question is whether the dissimilarities are merely a matter of degree, or whether they reflect a qualitative difference between the two processes.

---

mental time travel".

1.1.2 Narrow and broad discontinuism

The debate between continuists and discontinuists focuses on the contrast between FMTT and episodic memory, but episodic memory is only one form of past-oriented mental time travel (PMTT), which also includes processes such as episodic counterfactual thought. Thus we can distinguish between narrow and broad forms of (moderate) discontinuism. *Narrow discontinuism* maintains that, in addition to their distinct temporal orientations, there are further qualitative differences between FMTT and episodic memory, remaining silent about other forms of PMTT. *Broad discontinuism* maintains that, in addition to their distinct temporal orientations, there are further qualitative differences between FMTT and PMTT as a whole, including, but not limited to, episodic memory. (Continuism is always broad: given that we admit that episodic memory is continuous with FMTT, there would seem to be no reason to maintain that other forms of PMTT are not continuous with FMTT.)

The explicitly discontinuist arguments that have been offered in the literature are best understood as arguments for narrow discontinuism. As we will see, this is an important weak point shared by a number of otherwise distinct arguments. Discontinuists single out episodic memory and argue for a distinction in kind between it and FMTT. But what are we to say about other forms of PMTT? The discontinuist is faced with a choice. Either he groups other forms of PMTT with episodic memory or he groups them with FMTT. If they are grouped with episodic memory, narrow discontinuism collapses into broad discontinuism. If they are grouped with FMTT, narrow discontinuism maintains not only that that there are qualitative discontinuities between episodic memory and FMTT but also that there are qualitative discontinuities within PMTT, between episodic memory and other forms of PMTT. Neither of these positions is easy to maintain. In the case of the latter position, while it is somewhat plausible to contrast episodic memory with FMTT, it is much less plausible to contrast episodic memory with episodic counterfactual thought and other forms of PMTT; narrow discontinuism tends to lead to broad discontinuism. The former position (broad continuism), however, is itself an unstable position, as it is not particularly plausible to contrast PMTT as a whole with FMTT; grouping other forms of PMTT with episodic memory thus tends to lead us simply to adopt full-blown (moderate) continuism.

1.1.3 Metaphysical and epistemological discontinuism

A final distinction is between metaphysical and epistemological forms of (dis)continuism. The distinction between metaphysics and epistemology aligns roughly with the difference between descriptive questions about the mind (studied by philosophers of mind) and normative questions about

knowledge (studied by epistemologists). Thus *metaphysical discontinuism* accepts (and metaphysical continuism rejects) the existence of descriptive differences in kind between episodic memory and FMTT. For example, a view on which episodic memory has concrete past events as its objects, whereas FMTT has abstract types of possible future events as its objects, would amount to a form of metaphysical discontinuism (see section 2.2). Correspondingly, *epistemological discontinuism* accepts (and epistemological continuism rejects) the existence of differences in kind between the knowledge of past events that is provided by episodic memory and the knowledge of future events that is provided by FMTT. For example, a view on which certain kinds of misrepresentation are possible in remembering the past but not in imagining the future would amount to a form of epistemological discontinuism (see section 3.3).

Section 2 of this chapter discusses metaphysical continuism; section 3 discusses epistemological continuism. Which combinations of these metaphysical and epistemological views are coherent? Epistemological discontinuism presupposes metaphysical discontinuism, in the sense that, if metaphysical discontinuism turns out to be incorrect, it is unclear what motivation might remain for epistemological discontinuism. Thus, if we could be certain that the arguments given in section 2 were successful, we could in principle omit any discussion of the arguments of section 3. But given that the former arguments are not decisive, it is worth exploring independent arguments for epistemological discontinuism. While epistemological discontinuism seems to presuppose metaphysical discontinuism, metaphysical discontinuism does not entail epistemological discontinuism — there might in principle be qualitative differences between episodic memory and FMTT at the descriptive level that fail to translate into differences at the normative level. Thus onemight in principle endorse metaphysical discontinuism while also endorsing epistemological continuism.

Before proceeding, a note on where this chapter fits into the wider literature. The strategy of the chapter is primarily negative, in that it defends continuism by responding to arguments for discontinuism. The positive empirical case for metaphysical continuism has been made in the psychological literature discussed above; a positive philosophical case for metaphysical continuism is made in Michaelian (forthcoming). Supposing that the main arguments of this chapter — negative arguments against metaphysical and epistemological discontinuism — are successful, what is missing at this stage is a positive case for epistemological continuism; the final section of the chapter briefly sketches such a case, but this remains an important area for future work.

1.2 Empirical evidence and a priori arguments

There is a danger that philosophers and psychologists debating the continuity or discontinuity of episodic memory and FMTT will talk past each other. Psychologists endorse continuism primarily on the basis of empirical evidence for similarities between episodic memory and other forms of MTT. Philosophers endorse discontinuism primarily on the basis of arguments of a largely conceptual character. Can empirical evidence decide between continuism and discontinuism? Are purely conceptual arguments relevant to a debate over distinctions among kinds of mental processes?

1.2.1 The relevance of empirical evidence

The general thrust of MTT research — which has demonstrated extensive similarities between episodic memory and FMTT — clearly favours continuism. As Perrin points out, however, the MTT literature also includes a number of empirical findings concerning differences between episodic memory and FMTT at both the neural and the representational levels.[2]

A discontinuist might be tempted to invoke these findings in support of his view, but there are two difficulties involved in doing so. The first difficulty is that, while, if the observed differences were to suggest deep or extensive asymmetries between episodic memory and FMTT, they might in principle support discontinuism, in practice they amount to differences of degree and are therefore compatible with continuism. At the neural level, the discontinuist might cite the fact that brain imaging has revealed the involvement in FMTT of brain regions that are not involved, or are involved to a lesser extent, in episodic memory (see, e.g., Szpunar et al. 2007). Philosophers with no special expertise on brain imaging ought to defer to psychologists and neuroscientists on the interpretation of this sort of evidence, and this chapter will not discuss these findings in any detail. It is sufficient here to note that the consensus among psychologists and neuroscientists (e.g., Schacter and Addis 2007) is that the findings in question suggest differences of degree, rather than differences in kind, between episodic memory and FMTT, perhaps reflecting the greater processing demands of the latter. They thus provide little support for discontinuism.

At the representational level, one typical finding — echoing Hume's (1739) argument that memory is distinguished from imagination by its greater vivacity — is that episodic remembering produces representations containing greater sensory and contextual detail than is contained by those produced by FMTT (e.g., D'Argembeau and Van der Linden 2004). These representational-level differences might seem to provide a more promising

[2]Similarities and differences at the phenomenological level have also been important in MTT research; possible phenomenological discontinuities between episodic memory and FMTT are discussed at length in section 2.3.

basis for drawing a qualitative distinction between episodic memory and FMTT, but here, too, the differences ultimately seem to be a matter of degree, rather than kind. Reid (1764) — responding to Hume — remarked that he "would gladly know ... how one degree of vivacity fixes the existence of the object to the present moment; another carries it back to time past; a third, taking a contrary direction, carries it into futurity; and a fourth carries it out of existence altogether". Translated into contemporary terms,[3] Reid's point is, roughly, that the level of detail of representations of both future and past events is highly variable. Both episodic memory and FMTT may produce more or less detailed representations, and, while the representations produced by episodic memory may, on average, be more detailed than those produced by FMTT, the difference holds only on average. This means that there can be no clear cut-off point above which representations are sufficiently detailed to count as memory representations and below which they are too generic to count as memory representations. Thus the observed differences between episodic memory and FMTT at the representational level provide little support for discontinuism.

The second difficulty is posed by the fact that episodic memory is not the whole of PMTT, which also includes processes such as episodic counterfactual thought. Episodic counterfactual thought, like FMTT, calls on slightly different brain regions than does episodic memory (e.g., De Brigard et al. 2013) and might display differences at the representational level as well. If differences of this sort were to support a qualitative distinction between episodic memory and FMTT, they would presumably likewise support a similar distinction between episodic memory and other forms of PMTT. The discontinuist, then, can either treat the observed differences among episodic memory, FMTT, and episodic counterfactual thought as irrelevant or opt for a form of narrow discontinuism which explicitly claims that episodic memory is qualitatively distinct not only from FMTT but also from other forms of PMTT, grouping episodic counterfactual thought with FMTT under the broader heading of episodic imagination. If the discontinuist takes the former route, he must provide alternative reasons in support of discontinuism. If he takes the latter route, he owes us a reason for thinking that episodic memory, in particular, should be treated as special — since the various forms of episodic cognition all call on slightly different brain regions, we need some additional reason for singling out episodic memory while treating the other processes in question as being merely different forms of episodic imagination. Either way, the discontinuist must fall back on conceptual arguments.

[3] And ignoring Reid's direct realism; see section 2.2 below on direct vs. indirect theories of memory.

1.2.2 The status of a priori arguments

The chief arguments for discontinuism are indeed largely conceptual or a priori in character. Those of a naturalistic methodological bent — this includes many philosophers, in addition to most psychologists — might therefore doubt that the threat from discontinuism is particularly serious, suggesting that we base our understanding of the relationship between remembering the past and imagining the future on a demonstrably productive empirical research program, rather than ungrounded a priori intuitions. Indeed, since discontinuists do not deny the existence of important empirical symmetries between PMTT and FMTT, arguing, in many cases, for metaphysical and epistemological asymmetries that might obtain even if there were no empirical asymmetries whatsoever,[4] naturalists might suggest that the asymmetries between PMTT and FMTT with which discontinuists are concerned are not ones with which we should be concerned even in principle.

Discontinuism should, however, be of interest even to those who are sceptical of conceptual arguments. First, while conceptual arguments go beyond empirical MTT research, they ought to be consistent with the main findings of that research. Second, as the editors point out in their introduction to this volume, MTT research has proceeded at a breakneck pace that has left little time to develop the concepts needed to come to grips with its accumulating findings, and one potential benefit of engaging with philosophical resistance to the MTT framework is increased conceptual clarity: when we say that remembering the past and imagining the future are both forms of mental time travel, alike aside from their respective temporal orientations, what, exactly, do (or should) we mean?

2 Against metaphysical discontinuism

This section considers arguments for metaphysical discontinuism, discussing the implications of apparent causal (section 2.1), intentional (section 2.2), and phenomenological (section 2.3) discontinuities between episodic memory and FMTT.[5]

[4]Debus (2014), for example, argues that her discontinuist view is compatible with essentially all empirical MTT research.

[5]One might also argue for a functional discontinuity, suggesting that the function of episodic memory is to provide information which can be recombined to permit the simulation of possible future events, the latter being the function of FMTT. If this view of the respective functions of episodic memory and FMTT were accurate, it would indeed support a difference in kind between them. But the argument depends on an equivocation, treating episodic memory as a *store* — a source of raw materials for the simulation of episodes — and FMTT as a *process* that draws on the materials provided by episodic memory to simulate events. Bearing in mind the distinction between episodic *memory* (understood as a store) and episodic *remembering* (understood as a process), it is equally plausible to suggest that the function of episodic memory is to provide information which can be recombined to permit the simulation of past events, either counterfactual past events, in

2.1 Causal discontinuities

In addition to the quantitative representational-level differences noted above, an argument for discontinuism might be based on more fundamental apparent intentional differences between episodic memory and FMTT. One such argument is developed by Debus (2014), who sees intentional and causal discontinuities as intertwined. Debus's case for discontinuism rests on the view that, in episodic memory, the subject is experientially aware of a particular past event, whereas, in FMTT, the subject is not (and could not be) experientially aware of a particular future event.

She summarizes the argument as follows.

1. When a subject [episodically remembers] a past event, the subject is experientially aware of a past particular event.

2. When a subject [episodically imagines] a future event, the subject could not possibly be experientially aware of a future particular event.

3. Being experientially aware of a particular event and being in an experiential state in which one could not possibly be aware of a particular event are two different kinds of mental state.

4. Therefore, [episodic memories of] past events and [episodic imaginations] of future events are ultimately mental occurrences of two different kinds [i.e., metaphysical discontinuism is true].

Granting premise 3, this argument succeeds in establishing metaphysical discontinuism unless either premise 1 or premise 2 is false. Premise 1 has two components. First, it says that, in episodic memory, the subject enjoys *experiential awareness* of past events. Call this the positive relation claim. Second, it says that, in episodic memory, the subject is aware of *particular* past events. Call this the positive object claim. Premise 2 likewise has two aspects. First, it says that, in FMTT, the subject does not enjoy *experiential awareness* of future events. Call this the negative relation claim. Second, it says that, in FMTT, the subject is not aware of *particular* future events. Call this the negative object claim. This section discusses the relation claims (the causal aspect of the argument); section 2.2 discusses the object claims (the intentional aspect).[6]

episodic counterfactual thinking, or actual past events, in episodic remembering. As far as functional considerations go, then, the fact that episodic memory provides raw materials for the simulation of possible future events does not support an asymmetry between episodic remembering and FMTT. There is, moreover, a case to be made against the view that there is a dedicated store of episodic information. Klein (2015), in particular, argues that what makes a form of cognition episodic is a feature at the level of phenomenology, rather than information or content, namely, the involvement of autonoetic consciousness. If there is no dedicated episodic store, the functional argument does not get going in the first place.

[6]The initial stages of the discussion employ the language of representations; for reasons

2.1.1 The negative relation claim

In support of the negative relation claim, Debus argues, first, that, in order for a subject to be experientially aware of an event at a given time, the event must be "actual" "from the subject's own spatio-temporal perspective at the time" and, second, that future events are not actual. At first glance, this seems plausible enough. If experiential awareness is the kind of awareness that one has one when perceives an event, and if what it is for an event to be actual from the subject's spatiotemporal perspective is for it to occur at the relevant time, then premise 2 is uncontroversially right: one cannot be experientially aware of future events simply because one is not contemporaneous with them — they have not yet occurred. But if this is how we understand experiential awareness, premise 1 is likewise uncontroversially wrong: one cannot be experientially aware of past events simply because one is not contemporaneous with them – they have already occurred. Thus what Debus needs is a notion of experiential awareness on which a subject may be experientially aware of an event with which he is not contemporaneous, but only when the event in question is located in the past, not when it is located in the future.

2.1.2 The positive relation claim

The discontinuist therefore must understand experiential awareness in a looser, less literal sense. And indeed, Debus ultimately suggests that what it is for one to be aware of a past event, in the relevant sense, is for one to have experienced the event when it occurred and for one's current thought of the event to have the right sort of *causal* connection to one's original experience of it.[7] If this is what "experiential awareness" amounts to, then, while the label is misleading, a subject may be aware of an event that does not occur at the time at which he represents it, as long as he experienced at the time at which it did occur. And if enjoying this sort of awareness is definitive of episodic memory, then it does mark a qualitative asymmetry between remembering past events and imagining future events, since, even if one will eventually experience an event, one's current representation of an event cannot be caused by one's future experience of it. The tenability of continuism thus rests on the tenability of the denial of the claim that remembering requires a causal connection between the subject's current representation of a

given below, Debus herself rejects this language.

[7]Debus claims that experiential awareness *supervenes* on these factors. As philosophers use the notion of supervenience, experiential awareness will supervene on previous experience and causal connection just in case there can be no difference at the level of experiential awareness without a corresponding difference at the level of previous experience and causal connection. Debus resists the claim that experiential awareness *reduces* to these factors. Whether the latter claim is actually stronger than the former — i.e., whether supervenience is sufficient for reduction — is a technical question with which we need not be concerned here.

past event and his previous experience of that event, a condition which is at the heart of the influential causal theory of memory (Martin and Deutscher, 1966).

Before discussing the causal condition, some background is in order. Philosophers generally take it for granted that we can meaningfully distinguish between genuine and merely apparent memory. As it is standardly drawn, this distinction presupposes two claims. The first is the semantic claim that genuine remembering is "factive", where factivity means, in the case of semantic memory, that the content of the apparent memory is true or correct and, in the case of episodic memory, that the apparently remembered event actually occurred. If a subject seems to remember a proposition but the proposition is false, the apparent memory is said to be merely apparent. If a subject seems to remember an event but the event did not actually occur, the apparent memory is likewise said to be merely apparent. Factivity will be discussed further in section 2.2. The second claim is that, since an apparent memory might be accurate for the wrong sort of reason (e.g., due to blind luck), something more than simple correspondence between the subject's current representation and the event must be required to distinguish genuine memory from merely apparent memory — i.e., factivity is not enough for genuine memory. Different philosophical theories propose different accounts of the nature of this "something more". The causal theory, in particular, says that the subject's current representation must be caused (in the right way)[8] by his experience of the event itself.

Though it has been influential in philosophy, and though something like it may occasionally be presupposed in discussions of false memory in psychology, there is considerable empirical evidence against the causal theory, and hence the theory in turn can provide little support for metaphysical discontinuism. When arguing from empirical memory research to the conclusion that the causal theory is incorrect, there are two different paths that we might take.[9] First, we might start from the research on mental time travel discussed above. In outline — see Michaelian (forthcoming) for a more detailed version of this argument — the strategy here would be to argue that the overlap that has been demonstrated between remembering the past, imagining the future, episodic counterfactual thought, and other forms of episodic cognition shows that, intuitions to the contrary notwithstanding, there is nothing special about episodic memory — episodic memory is

[8]The qualification that causation must occur "in the right way" is required to cope with deviant causal chains, a problem which need not concern us here; see Michaelian 2011 for background.

[9]Both paths assume that the causal link supposed to be characteristic of genuine remembering is an informational link, that is, that the causal connection between the retrieved memory representation and the original experience goes via stored information deriving from the experience. In principle, the causal theorist might deny this assumption, but it is not clear that a non-informational characterization of non-deviant causal chains can be given.

just one form of episodic imagination among others. Like any other form of episodic imagination, episodic remembering is a simulational process which draws on information originating in a variety of different sources to construct a representation of a target episode. Just as other forms of episodic imagination of necessity do not draw on any information originating in the episodes they target (since the subject has not experienced those episodes), episodic memory need not draw on any information originating in the episodes it targets. In many cases, information originating in the subject's experience of the target episode will in fact end up being incorporated into the subject's memory representation of the event. In such cases, memory may ensure an appropriate causal connection between the memory representation and the original experience. But even in such cases, information originating in experience of the target episode may represent only a small fraction of the total content of the memory representation; there may thus be equally strong or stronger connections to another episode or episodes, making the relevance of the causal connection to the target episode unclear. And in many other cases, *no* information originating in the subject's experience of the target episode ends up being incorporated into the subject's memory representation of the event. In such cases, memory involves no causal connection with the original experience whatsoever.

If this picture of episodic memory as one form of episodic imagination among others is right, the causal condition does not provide an adequate account of the "something else" distinguishing genuine from merely apparent remembering,[10] and discontinuists may not appeal to the causal theory to ground the claim that there is a qualitative metaphysical difference between episodic memory and FMTT. Discontinuists might in principle respond to this argument by drawing a distinction within episodic memory, arguing that only representations of events from the personal past which are produced by processes that do satisfy the causal condition qualify as genuine memories. This response, however, sits ill with the workings of the episodic construction system. There is overwhelming evidence that, while the system aims, within certain limits, at producing *accurate* representations of past episodes, it is simply not designed to ensure *causal connections* with past episodes. Remembering regularly fails to involve causal connections, and there is no difference, at the neurocognitive level, between (apparent)

[10]What sort of account is available, if this picture is right? One possibility worth exploring is that genuine remembering is distinguished from merely apparent remembering simply by the fact that the subject's episodic construction system takes an episode from the subject's actual personal past, as opposed to the personal future or the counterfactual personal past, as its target. See Michaelian (forthcoming) for discussion of this possibility. Another possibility is that genuine remembering is distinguished from merely apparent remembering by the subject's personal-level intentions. See Hoerl (2014) for discussion of this possibility. The remainder of the argument of this chapter does not depend on a particular alternative to the causal condition as a means of distinguishing between genuine and merely apparent memory.

memories that do involve a causal connection and (apparent) memories that do not. Given that we want our philosophical account of remembering to correspond to a process that unfolds at the neurocognitive level, there is simply no justification for drawing this sort of distinction within episodic memory.

Situating episodic memory as a form of PMTT further undermines the causal condition. In addition to producing representations of episodes from the actual personal past, the episodic construction system frequently aims at producing representations of counterfactual episodes from the personal past. And while it may be intuitively plausible to claim that episodic memory is distinguished from FMTT by the fact that it the former necessarily involves a causal connection with the target episode, it is considerably less plausible to claim that it is distinguished from episodic counterfactual thought in this manner. The boundary between episodic memory and episodic counterfactual thought is fluid, in the sense that we are able to shift freely back and forth between (attempting to) remember events as they actually occurred and (attempting to) imagine them occurring in more or less different ways. In fact, given that episodic remembering need not be fully accurate to count as remembering, so that it is compatible with a certain amount of modification of the remembered event, and given that much episodic counterfactual thinking consists not in imagining wholly novel events but rather in modifying aspects of past events, in some cases it may be indeterminate whether the subject is engaged in episodic remembering or episodic counterfactual thinking. Episodic counterfactual thought, of course, does not presuppose a causal connection between the subject's current representation and his previous experience of the event, simply because counterfactual events (by definition) did not occur or did not occur as represented. Given the fluidity of the boundary between episodic memory and episodic counterfactual thought, it becomes difficult to justify imposing the causal condition on remembering.

Now, this first path to the rejection of the causal theory risks being question-begging in the current context, since it requires us to grant that there is no qualitative distinction between episodic memory and other forms of episodic imagination, which is precisely the matter at issue in the debate between continuists and discontinuists. The second path from empirical memory research to the rejection of the causal theory avoids this difficulty — see Sutton and Michaelian (forthcoming) for a more detailed version of this argument. Both continuists and discontinuists ought to accept the basic lesson of constructive memory research: remembering is never a matter of simply retrieving a stored "copy" of an experience but is always to some extent reconstructive, where reconstruction implies, inter alia, the inclusion, in the "retrieved" memory representation, of information not originating in the subject's experience of the target episode. Call the inclusion of such information "supplementation". Is the occurrence of supplementation com-

patible with satisfaction of the causal condition? One might argue that the causal condition can be satisfied in cases of supplementation, as long as the retrieved representation does not include too much information originating in sources other than the subject's experience of the target episode: some degree of supplementation is compatible with genuine remembering; only in cases where supplementation results in the inclusion of too much information originating in such sources does remembering fail to occur (Michaelian, 2011). One obvious difficulty with this position is its vagueness: how much supplementation is too much? Setting this difficulty aside, though, the position attempts to draw a distinction where none is to be found. The process that is at work when the subject, in an attempt to remember a given past episode, constructs a representation that includes a large fraction of information originating in his experience of the episode is not distinct, at the neurocognitive level, from the process that is at work when he constructs a representation that includes only a small fraction of information originating in his experience of the episode; nor is it distinct from the process that is at work when he constructs a representation that includes no information whatsoever originating in his experience of the episode. If no information originating in his experience of the episode is preserved, the causal condition is not satisfied. Thus the constructive character of remembering rules out the claim that genuine, as opposed to merely apparent, remembering presupposes satisfaction of the causal condition. Again, the causal condition does not provide an adequate account of the "something else" distinguishing genuine from merely apparent memory, and discontinuists should not appeal to the causal theory to ground the claim that there is a qualitative metaphysical difference between episodic memory and FMTT.[11]

2.2 Intentional discontinuities

In short, while there is much more that might be said both for and against the causal theory, the ground on which the theory stands at this point is arguably too shaky for it to provide significant support for metaphysical discontinuism. Since it is at best unclear whether genuinely remembering a past event does require a causal connection with the remembered event, we should look elsewhere for qualitative descriptive differences between episodic memory and FMTT. Debus maintains both that episodic memory involves causal connection to events (the positive relation claim) and that it involves awareness of particular events (the positive object claim). She maintains, further, that FMTT does not involve causal connection to events (the negative relation claim) and that it does not involve awareness of particular events (the negative object claim). While the negative relation claim is unproblematic, we have seen that there is reason to doubt the positive relation

[11] If the causal theory is incorrect, how can the reliability of remembering be explained? This question is taken up in section 4.

claim. Nevertheless, Debus's argument for metaphysical discontinuism will go through if both object claims are correct. There are reasons to doubt each of these claims.

2.2.1 The negative object claim

If we set aside the negative relation claim, the negative object claim collapses into the claim that it is impossible for subjects to represent particular future events. By way of argument for this claim, Debus begins by sketching a two-step account of what is involved in imagining future events. On her account, imagining a future event (e.g., having lunch with a friend next week) involves, first, imagining a general type of event (having lunch with a friend) and, second, imaginatively projecting that event into the future (next week). She then argues that, given this account, imagining a future event never amounts to imagining a *particular* future event. There are two problems with this strategy. First, it does not show that we cannot imagine particular future events by employing some other procedure. Second, it does not convincingly show that we cannot imagine particular future events by employing the two-step procedure.

As far as the first problem goes, it is by no means clear that the procedure described by Debus is in fact employed in all cases of FMTT, and other ways of engaging in FMTT may enable us to imagine particular future events. For example, it would seem to be possible to imagine a future event by, first, remembering a particular past event and, second, imaginatively projecting that event into the future. I might, for example, imagine having lunch with a friend next week by remembering having lunch with him last week and then imagining that event occurring next week rather than last week. If this procedure is employed, then, assuming that the representation produced by remembering qualifies as a representation of a particular past event, it should arguably not cease to count as a representation of a particular event simply because it is projected into the future.

The discontinuist might suggest that the representation does indeed cease to count as a representation of a particular event when it is projected into the future on the ground that this sort of projection involves imaginative modification of an important feature of the event, namely, its temporal location. Taking the counterfactual dimension of PMTT into account, however, undermines this suggestion. Thinking of a particular past event is compatible with imaginatively modifying various feature of the event. I might, for example, remember having lunch with a friend last week but imagine eating slightly different food. In this case, I am presumably still thinking of a particular event. I might go further, remembering having lunch with a friend but imagining eating in a different restaurant, modifying the spatial location of the event. In this case, I may or may not count as being engaged in remembering, but I am presumably still thinking of a particular event. If

imaginatively modifying their spatial locations is compatible with thinking of particular events, presumably imaginatively modifying their temporal location is as well. Nor does it seem to matter whether I modify the temporal location of the event by imagining it occurring at a different time in the past or by imagining it occurring in the future. I might remember having lunch with a friend last week but imagine the lunch occurring somewhat later — closer to the present — than it did. If I then shift from episodic counterfactual thought to FMTT and project the event forward into the future — say, tomorrow, rather than yesterday — nothing would seem to change, as far as the particularity of the event that I am imagining is concerned.

Conceding this point, the discontinuist might argue that the possibility of imagining a particular future event by remembering a particular past event and projecting it into the future is compatible with the spirit of discontinuism, since one's imagination of a particular future event in such cases is parasitic on one's memory of a particular past event, which might be enough to secure a metaphysical asymmetry between episodic memory and FMTT. The second problem for Debus's strategy, however, undermines this move. The problem is that, even in cases in which imagining future events employs the original two-step procedure — thinking of a general type of event and projecting it into the future — it is unclear whether the agent necessarily fails to imagine a particular future event. As Debus admits, whether one can, by means of this procedure, imagine a particular event may depend on how we individuate events. If events are individuated in a thin manner, simply by reference to spatiotemporal location, the procedure clearly enables subjects to think of particular future events. There is no barrier to imagining a general type of event occurring at a given spatiotemporal location.

Debus's argument thus presupposes a thicker way of individuating events. If events are individuated in a slightly thicker manner, say, by reference to the particular agents they involve, in addition to spatiotemporal location, the procedure still enables subjects to think of particular future events. There is no barrier to imagining a general type of event occurring at a given spatiotemporal location and involving given people. We might suspect that, if events are individuated in a maximally thick manner, as fully determinate, concrete entities, the procedure will not enable subjects to think of particular future events, since imagination of a general type of event will necessarily fail to capture the full richness of the real, concrete event. At this point, however, the question of event individuation starts to look like a red herring. Regardless of how events are individuated, one need not, in order to think of a particular past event, recall it in every detail. In many cases, one's memory for past events may in fact be quite schematic. Analogously, one need not, in order to think of a particular future event, imagine it in every detail. Of course, if events are individuated in a maximally thick manner, detailed representations of future events will very often be somewhat inac-

curate, simply because there are limits to our powers of anticipation. But the same thing goes for detailed representations of past events, which are likewise often inaccurate — there are limits to our powers of recollection, just as there are limits to our powers of anticipation. At least within certain limits, inaccuracy need prevent us neither from thinking of particular past event nor from thinking of particular future events.

The suspicion that the two-step procedure may suffice to allow us to imagine particular future events is reinforced by the observation that, in many cases, the very same procedure is employed in episodic memory: in many cases, remembering a past event amounts precisely to imagining a general type of event and imaginatively projecting it into the past. Consider memory for frequently repeated events. In many such cases, one retains little information from any particular repetition of the event, but one does retain a general representation of the event type. Remembering a particular repetition of the event may amount to forming a representation of the event type and imaginatively projecting it back to the relevant past time. If remembering can put us in contact with particular events when it employs this procedure, then so, presumably, can imagining. Of course, a discontinuist might suggest that, in cases where remembering employs this procedure, it does not count as genuinely episodic remembering, i.e., that it does not in fact involve representing a particular past event. But this suggestion is difficult to reconcile with the reconstructive character of episodic memory, which routinely draws both on information originating in experience of particular events and on generic knowledge, including scripts for common event types. In fact, as several of the contributions to this volume make clear (Irish; Klein and Steindam; Szpunar et al.), both PMTT and FMTT may in some cases depend exclusively on semantic knowledge. Alternatively, the discontinuist might grant that remembering sometimes employs the two-step procedure but suggest that it counts as genuinely episodic remembering only if there is a shift from projecting a general type of event to a past time to thinking of a particular event, strictly understood. But any plausible explanation of what such a shift might consist in is likely to invoke the causal criterion, and, as we have seen, there are independent problems for that criterion.

In light of these two problems for Debus's strategy, the discontinuist might look for another means of motivating the negative object claim. Rather than motivating it by appealing to the nature of *representations* of future events, the discontinuist might attempt to motivate it by appealing to the nature of future events themselves, arguing that we cannot represent particular future events simply because there is no such thing as a future event. This alternative strategy attempts to base a metaphysical asymmetry between episodic memory and FMTT on a metaphysical asymmetry between the past and the future, invoking an indeterminist metaphysics according to which, while there is a matter of fact about what happened in the past, there is no determinate matter of fact about what will happen in the fu-

ture. While such an approach would provide no guarantee that we are able to think of particular past events, it would guarantee that we are unable to think of particular future events and so would constitute a step towards metaphysical discontinuism.

This alternative strategy is not particularly promising means of motivating the negative object claim. First, the metaphysical openness of the future is a controversial matter (Debus this volume); the strategy is thus risky, as it depends on the outcome of an ongoing debate in the metaphysics of time. Second, the strategy seems to miss the point. What matters, in the debate between continuists and discontinuists, is not the *metaphysical* openness of the future but rather its *cognitive* openness. Discontinuism becomes no less plausible if we suppose that metaphysical determinism is true: even if the future is metaphysically closed, the discontinuist can maintain that it is cognitively open and that we therefore cannot enter into cognitive contact with particular future events. By the same token, the continuist can maintain that, even if the past is metaphysically closed, it is cognitively open (since our knowledge of the past is limited) and that we therefore cannot enter into cognitive contact with particular past events. Invoking a metaphysical asymmetry between the past and the future is thus unlikely to provide adequate support for a metaphysical asymmetry between past- and future-oriented MTT.

2.2.2 The positive object claim

While the considerations raised in the previous section cast doubt on the negative object claim, according to which we are unable to imagine particular future events, they do not establish that we are in fact able to imagine particular future events. If it should turn out that we are unable to imagine particular future events, the discontinuist might fall back on the positive object claim, according to which are able to remember particular past events, to motivate metaphysical discontinuism. Thus our assessment of the case for metaphysical discontinuism requires an examination of the case for the positive object claim. This section argues that, while the discussion so far has taken for granted that we are able to remember particular past events, there is a real possibility that the positive object claim is false, i.e., that we are unable to imagine particular past events.

In order to make a case for the claim that episodic memory puts the subject in contact with particular past events, Debus appeals to her relational account of memory (Debus, 2008). In addition to the positive relation claim (discussed above), the relational account includes the claim that, when a subject remembers, the remembered event itself is (part of) the content of his memory. The relational account thus amounts to a form of direct realism about the objects of memory (hence Debus cites Reid 1764 as a precursor

of her view).[12] In view of the serious problems facing direct realism, the relational account can provide little support for the positive object claim.

The distinction between direct and indirect (or representative) realism has been most extensively discussed in the philosophy of perception. The direct realist about perception claims that, in perceiving, what the subject is related to in the first instance is an external object. The indirect realist, in contrast, claims that what the subject is related to in the first instance is an internal representation; the subject counts as seeing the external object when his internal representation is related to it in the right way. Because it posits internal representations that function as intermediaries between the perceiving subject and the perceived object, the story told by indirect realism about the nature of perception is more complicated than that told by direct realism. But there is a powerful motivation for the introduction of this additional complexity, in the form of the argument from illusion, which motivates the introduction of internal representations by pointing out that perception is not always veridical.

In cases of hallucination, for example, the subject "sees" something that is not in fact there. Veridical perception seems to have something important in common with hallucination; in principle, after all, hallucination might be indistinguishable, from the subject's perspective, from veridical perception. Direct realism, since it views perception in terms of a direct relation between subject and object, has difficulty accounting for this fact. Indirect realism, in contrast, can accommodate it without difficulty, since it can point to the internal representation as the factor common to veridical perception and hallucination. The upshot is that direct realists are forced to adopt a "disjunctive" view of perception: veridical perception and hallucination are not two kinds of perception; only veridical perception is perception; hallucination is another kind of state entirely.

The direct realist about memory claims that, in remembering, what the subject is related to in the first instance is a past event. The indirect realist, in contrast, claims that what the subject is related to in the first instance is an internal representation; the subject counts as remembering a past event when his internal representation is related to it in the right way. Just as indirect realism about perception can be motivated by an argument from perceptual illusion, indirect realism about memory can be motivated by an argument from memory illusions — distorted or inaccurate memories. In the case of "memory hallucinations", i.e., wholly false memories, the subject "remembers" an event that did not in fact occur. Accurate memory seems to have something important in common with memory hallucination; in principle, memory hallucination might be indistinguishable, from the sub-

[12] As noted above, this means that it is misleading to describe Debus's view using the language of memory *representations*, as we have been doing so far; but nothing in the foregoing argument depends on this.

ject's perspective, from accurate memory. And just as direct realism about perception has difficulty accommodating perceptual hallucination, direct realism about memory has difficulty accommodating memory hallucination. Indirect realism about memory, in contrast, can accommodate it without difficulty, since it can point to the memory representation as the factor common to accurate memory and memory hallucination. The upshot is that, just as direct realists about perception are forced to adopt a disjunctive view of perception, direct realists about memory are forced to adopt a disjunctive view of memory: accurate memory and memory hallucination are not two kinds of memory; only accurate memory is memory; memory hallucination is another sort of state entirely.

Though it has its defenders, disjunctivism is implausible not only from an intuitive but also from a naturalistic perspective. From a naturalistic perspective, we should not draw distinctions between mental states or processes where none are to be found at the neurocognitive level; this is precisely what disjunctivism requires us to do. One way to see this is to note that, just as direct realism implies disjunctivism, it implies factivity: as Debus puts it, "a subject could not possibly have an [episodic memory] of an event that did not take place" (2008, 410). As noted above, it is a commonplace in philosophy that memory is factive, but factivity is empirically untenable. At the neurocognitive level, there need be no qualitative difference between the process at work when the episodic construction system produces a wholly inaccurate representation of a past event and the process at work when the system produces an accurate representation of a past event. It is tempting to suppose that there must be at least some quantitative difference, such as greater fluency in the case of accurate memory than in the case of memory hallucination (Perrin this volume), but there need not be even a quantitative difference. A counterfactual event that has been frequently imagined, for example, might be remembered just as fluently as an experienced event (Loftus, 1996). Factivity — and disjunctivism — should thus be rejected.[13]

Since direct realism about perception entails disjunctivism, it encounters an analogous objection. But direct realism about memory also faces a problem which is not shared by direct realism about perception. The objects of perception are present (i.e., they are contemporaneous with the perceiving subject); hence, setting aside disjunctivism, it is relatively easy to understand how they might form part of the content of the subject's perceptual states. The objects of episodic memory, in contrast, are in the past; hence it is much more difficult to understand how they might form part of the content of the subject's memories. The memory exists now; the event does

[13]As far as the distinction between genuine and merely apparent memory is concerned, the rejection of factivity implies that the distinction rests entirely on the second sort of difference between genuine and merely apparent memory that rival theories of remembering attempt to capture. Thus it might rest, for example, entirely on the aims of the episodic construction system or the subject's personal-level intentions.

not; how, then, might the latter be part of the content of the former? One possibility here is to invoke mnemic causation (Russell, 1921), the idea that past events can be direct causes of our present memories of them. Were the notion of mnemic causation legitimate, it might provide the sort of direct connection between present representations and past events needed to allow the latter to constitute part of the former. The problem, of course, is that the notion of causation at a temporal distance is borderline unintelligible (Bernecker, 2008), and hence relying on it to save direct realism would be a desperate move.

While this brief discussion cannot provide a decisive case against the disjunctivism entailed by direct realism, we have at least good prima facie reason to reject the positive object claim, if it is understood along the direct realist lines suggested by Debus. The alternative is to understand it along indirect realist lines, as saying that remembering puts subjects in contact with particular past events by means of internal representations that are appropriately related to those events. As we saw above, however, if we understand the negative object claim in the same way, there is little reason to think that imagination does not put subjects in contact with particular future events. In sum, neither Debus's object claims nor her relation claims go through, and her argument does not establish metaphysical discontinuism.

2.3 Phenomenological discontinuities

If the argument so far is on track, neither causal nor intentional considerations establish metaphysical discontinuism. The possibility remains that there are qualitative phenomenological discontinuities between episodic memory and FMTT. Perrin (this volume) claims to identify two such discontinuities, one concerning the role of autonoesis, the form of consciousness characteristic of MTT, and one concerning the mechanisms respectively responsible for the production of autonoesis in episodic memory and FMTT. This section discusses these alleged phenomenological discontinuities in turn.

2.3.1 The role of autonoesis

Continuists maintain that autonoesis plays essentially the same role in FMTT as it does in episodic memory — conferring a sense of "mineness" or ownership on an episodic representation, a sense that the represented event belongs to the subject's own personal past or personal future (as opposed to the life of some other subject or possibly of no one in particular). Perrin argues that this is a mistake. While he grants that FMTT, like episodic memory, draws on content originating in the subject's own experiences, he argues that the fact that FMTT draws on experiential content does not guarantee the involvement of autonoetic consciousness in the process. He points out, in particular, that a subject might, while drawing on the very same experiential

content, switch back and forth between imagining his own future experience and imagining a future experience belonging to another subject. Only in the former case will the subject enjoy autonoetic consciousness; the latter case will involve a different sort of phenomenology. Thus Perrin seeks to drive a wedge between autonoetic consciousness and mere episodic subjectivity: the latter, unlike the former, "does not imply the identity of the self who is episodically thinking with the self whose experience is being thought about". If this is right, then autonoetic consciousness is not a necessary component of FMTT, whereas it is a necessary component of episodic memory.[14]

Given that autonoesis is distinct from chronesthesia — with autonoesis referring specifically to consciousness of the self in subjective time and chronesthesia referring to consciousness of subjective time more broadly (Tulving, 2002) — the continuist should admit that autonoesis is not intrinsic to episodic cognition. The involvement of experiential content in FMTT does not imply the involvement of autonoesis, and for precisely the reason that Perrin identifies: FMTT enables the subject to imagine not only his own future experiences but also those of other subjects. Situating episodic memory as one form of PMTT, however, it is clear that the same thing goes for PMTT. The involvement of experiential content in PMTT does not imply the involvement of autonoesis, for precisely the same reason: PMTT enables the subject to imagine not only his own past experiences (i.e., to remember) but also to imagine those of other subjects.

At this point, the discontinuist might retreat from broad to narrow discontinuism, restricting the scope of the claim about the role of autonoesis to episodic memory in particular, as opposed to PMTT as a whole, to argue that episodic memory necessarily involves autonoesis, while FMTT and other forms of PMTT do not. But there are two problems with this move. First, even disregarding the question of the status of animal "episodic-like" memory (on which see, e.g., Martin-Ordas this volume, Corballis this volume, Thom and Clayton this volume), it is not in fact clear whether episodic memory necessarily involves autonoesis. There are, for example, cases of subjects who have intact episodic memory, in the sense that they are able to construct detailed representations of events from the personal past, despite suffering from impaired autonoesis, in the sense that they do not feel as if the events in question belong to them (e.g., Klein and Nichols 2012).

This does not provide an open and shut case against the move; for example, the discontinuist might maintain that autonoesis is necessarily involved in ordinary, non-pathological cases of episodic memory, but the second problem comes into play at this point. Even if we grant that episodic memory does necessarily involve autonoesis, Perrin's argument asks us to compare apples and oranges — or rather to compare a single apple selected from

[14]Klein has made a related point; see Klein forthcoming, Klein and Steindam this volume.

basket of varied apples to a whole basket of varied oranges. There are multiple forms of FMTT and multiple forms of PMTT. Thus it is illegitimate to select one form of PMTT, namely, episodic memory, point out that it has a certain phenomenological feature, note that some forms of FMTT lack that phenomenological feature, and argue on that basis for an asymmetry between episodic memory and FMTT. Since there are multiple forms of FMTT and PMTT, what the discontinuist needs to show is that the particular form of FMTT that corresponds to the particular form of PMTT in question — episodic memory — lacks or can lack the phenomenology characteristic of that form of PMTT. And there is no reason to think that it does or can: when I imagine an event of my personal future, the event seems to belong to me, just as remembered events do. There may be subtle phenomenological differences between MTT into the personal future and episodic memory, stemming from the greater uncertainty of the future. But the fact that I cannot be sure whether an anticipated event will actually occur as represented does not imply that I cannot be sure that I am the person at issue in my representation of the event, just as the fact that I cannot be sure whether a remembered event actually occurred as represented does not imply that I cannot be sure that I am the person at issue in my representation of the event.[15] Moreover, episodic memory is not the only form of PMTT to involve autonoetic consciousness, since the latter is also involved in episodic counterfactual thought; and any phenomenological difference between MTT into the personal future is unlikely to be greater than the phenomenological difference between episodic memory and episodic counterfactual thought.

2.3.2 Mechanisms of autonoesis

The second apparent discontinuity identified by Perrin concerns not the role but rather the production of autonoesis, specifically, the mechanism that brings autonoetic consciousness into play in a given episodic constructive process. He argues, first, that, in episodic memory, autonoesis is the result of an automatic, subpersonal-level monitoring process (roughly, a type 1 process), suggesting that it derives from the detection of certain features of the first-order episodic simulation process, perhaps its fluency. Due to its automatic, subpersonal-level character, autonoesis is impenetrable to belief in cases of episodic memory: even if one reflectively concludes that the represented event does not in fact belong to one's personal past, one will continue to feel as if it does. He argues, second, that, in FMTT, autonoesis (when it occurs) is the result of a deliberate, personal-level process (roughly, a type 2 process); for example, one might anticipate the occurrence of a certain event and infer, on the basis of semantic knowledge, that one will

[15]Perrin might grant this point, arguing that the difference concerns the source of the certainty: FMTT is, and episodic memory is not, immune to error through misidentification. This argument is discussed in section 3.

participate in the event in a given role. Due to its deliberate, personal-level character, autonoesis *is* penetrable to beliefs in cases of FMTT: if one concludes that the represented event does not in fact belong to one's personal future, one will cease to feel as if it does,

The first part of this argument is plausible. Given that other forms of PMTT can be just as fluent as episodic memory, the monitoring process is likely to rely on multiple factors, rather than fluency alone, but the production of autonoesis in episodic remembering might indeed result from a subpersonal-level monitoring process. The second part of the argument is much less plausible. On the one hand, if it were indeed possible for a personal-level process to generate autonoesis in the case of FMTT, the same process could in principle be responsible for generating autonoesis in the case of PMTT, including episodic memory. It is unclear, however, how a personal-level process might, in either case, result in the sort of primitive phenomenal state that autonoesis is generally taken to be. On the other hand, it is by no means clear that a subpersonal-level monitoring process could not be responsible for generating autonoesis in the case of FMTT. Whatever factors such a monitoring process might draw on in the case of episodic memory will also be available in the case of FMTT. Fluency alone will not discriminate perfectly between representations of one's own future experiences and those of other subjects, but neither does it discriminate perfectly between representations of one's own past experiences and those of other subjects.[16] Moreover, an account that posits two distinct mechanisms, one responsible for generating autonoesis in the case of PMTT, the other responsible for generating autonoesis in the case of FMTT, is unparsimonious; if an account positing a single mechanism is available, it is to be preferred. On the whole, then, it is possible that a subpersonal-level mechanism generates autonoesis in the case of episodic memory, while a personal-level mechanism generates autonoesis in the case of FMTT, but the more likely possibility is that a subpersonal-level mechanism is at work in both cases.

In short, neither of the apparent phenomenological discontinuities between episodic memory and FMTT identified by Perrin appears to be bona fide. Overall, the case for metaphysical discontinuism is weak.

3 Against epistemological discontinuism

If metaphysical discontinuism is incorrect, epistemological discontinuism is unlikely to be correct. Nevertheless, it is worth looking separately, if more briefly, at arguments for epistemological discontinuism. Even if metaphysical discontinuism is correct, epistemological discontinuism might turn out

[16]See Michaelian 2012, Michaelian forthcoming for suggestions as to other potential factors that might be involved in metacognitive monitoring of episodic construction.

to be incorrect, since metaphysical discontinuities do not necessarily entail epistemological discontinuities. A discontinuist unconvinced by the considerations raised against metaphysical discontinuism in the foregoing might thus still be persuaded to reject epistemological discontinuism.

This section therefore discusses possible epistemological discontinuities between episodic memory and FMTT, looking at the epistemic openness of the future and the past (section 3.1), the directness or indirectness of our knowledge of past and future events (section 3.2), and immunity to error through misidentification — or the lack thereof — in episodic memory and FMTT (section 3.3).

3.1 The epistemic openness of the future and the past

A natural first thought to have about potential epistemological discontinuities between episodic memory and FMTT is that, if FMTT can provide us with any knowledge of the future at all, that knowledge is bound to be dramatically less secure than the knowledge of the past with which we are provided by episodic memory. As noted in section 2, the future is cognitively or epistemically open, whether or not it is metaphysically open: even if there is a determinate matter of fact about what will happen, we cannot be *certain* about what will happen. We might go further: even if there is a determinate matter of fact about what will happen, we cannot be certain that *anything at all* will happen — no matter how detailed our projections, it is entirely consistent with our epistemic situation vis-à-vis the future that the entire world will blink out of existence five minutes from now. In contrast, one might maintain, even if we cannot be fully certain that the details of our memories of past events are accurate, we can at least be certain that the world did not blink into existence five minutes ago. The contrast between our inability to rule out the unreality of the future and our ability to rule out the unreality of the past might thus seem to underwrite a qualitative asymmetry between our knowledge of future events and our knowledge of past events.

While this is a natural thought to have, it is mistaken. As Russell (1921) notoriously pointed out, we are not in fact able to rule out the unreality of the past. No matter how detailed our memories, no matter how subjectively convincing they are, it is entirely consistent with our epistemic situation vis-à-vis the past that the entire world blinked into existence five minutes ago, complete with our apparent memories of past events. This is, of course, just a particular instance of a general point about scepticism in epistemology: if we raise the standards for knowledge sufficiently high, requiring not just reliability but certainty, we deprive ourselves of virtually all knowledge, including knowledge of the very existence of the external world — if certainty is a prerequisite for knowledge, we have knowledge neither of past events, nor of future events, nor even of present events. Non-sceptics —

and that includes virtually all of us — should thus set the standards for knowledge well below the level of certainty. And if knowledge presupposes only some lower level of reliability, we have no reason (so far) to think that there is a qualitative asymmetry between our knowledge of past events and our knowledge of future events. Our beliefs about future events may well be somewhat less reliable than our beliefs about past events, but this gives us a merely quantitative epistemological difference between episodic memory and FMTT.

3.2 Direct and indirect knowledge

If a qualitative epistemological difference between episodic memory and FMTT cannot be located at the level of reliability, perhaps it can be located elsewhere, either in the directness or indirectness of our knowledge of past and future events, or in the extent of our indirect knowledge.

3.2.1 Indirect knowledge

As far as indirect or inferential knowledge goes, our beliefs about the future appear to be on a par with our beliefs about the past. Future events are not given but inferred, and our knowledge of them thus depends on the reliability of the inferences that we perform, on the basis of the evidence available to us in the present, about what is going to happen. But the very same thing goes for our knowledge of past events. As Kneale (1971, 11) puts it,

> We may ... be misled by a truism, that what is past has happened and that what is future has not, into the belief that the past is necessarily more accessible for knowledge than the future. After all, it is *there* and we should be able to find out about it. But this does not follow. Past events have happened, but whether they can be known depends upon whether they are connected with present events by laws which allow of backward inference, and whether present observation, memory, and records provide us with sufficient information to apply such laws.

In general, past events are no more given than are future events, and our knowledge of them depends on the reliability of the inferences that we perform, on the basis of the evidence available to us in the present, about what must have happened. This point holds both for systematic, scientific inquiries into the past and the future and for episodic simulation of past and future events.

Why, if there is no qualitative difference between our indirect knowledge of past events and our indirect knowledge of future events, do we ordinarily take ourselves to know much more about the past than we do about the

future? Kneale suggests that this habit is explained by two broad psychological tendencies. On the one hand, when we think about the past, we tend to focus on what we *do* know. Her examples are primarily historical, but this goes for episodic memory as well: there are whole periods of one's life about which one remembers little or nothing, but one is normally only dimly aware of these, preferring to focus on well-remembered events ("meeting my wife", "the birth of my child", ...), as it is the latter that tend to be especially significant. On the other hand, when we think about the future, we tend to focus on what we *do not* know. Again, Kneale's examples are primarily historical, but this goes for FMTT as well: there are many future events which can anticipate in great detail, but one normally focuses on especially significant events ("my child leaving home", "my retirement", ...), and these are often difficult to envision in any detail.

This explanation might itself be taken to provide an alternative route to an epistemological discontinuity between episodic memory and FMTT. Our knowledge of the past, the discontinuist might argue, includes a great deal of specific information about events, whereas our knowledge of the future consists mostly of information about generic, repeated features of events — this is why we know more about the personally important events of the past than we do about the personally important events of the future. If the claim is merely that we know more about particular past events than we do about particular future events, it is probably right but does not establish a qualitative difference between episodic memory and FMTT.[17] If the claim is, more strongly, that, while we can have knowledge of particular past events, we cannot have knowledge of particular future events, it is vulnerable to the considerations raised in section 2.2 against Debus's positive and negative object claims. Thus the suggested route does not look promising.

3.2.2 Direct knowledge

With respect to indirect or inferential knowledge, Kneale argues for continuity between our knowledge of the future and our knowledge of the past. With respect to direct or non-inferential knowledge, however, she grants that there is a fundamental discontinuity between them, in that only memory provides us with direct knowledge of events. Indeed, she treats "the capacity to have non-inferential self-certifying beliefs about particular matters of fact in the past" as definitive of memory and argues that "we have no such self-certifying beliefs about the future" (1971, 1-2). In this, she is in agreement with a number of other philosophers. Swinburne, for example, argues that "men have memory but not foreknowledge. Their non-inferential knowledge is of the past but not of the future" (1966, 167).

[17] Even if it is right, there may be a chicken-and-egg problem here: to what extent do we remember events because they are important to us, and to what extent are events important to us because we remember them?

While this view has been popular among philosophers, it is not immediately clear what it might mean to say that memory is a direct source of knowledge. Kneale makes clear that her point is not about reliability: she acknowledges that memory is imperfectly reliable and suggests that, even if anticipation were just as reliable as memory, it would not count as a direct source of knowledge of the future. Ultimately, her point boils down to the claim that it is part of the very concept of memory that "the memory event [i.e., the occurrence of remembering] should have as a part-cause the occurrence of the event recollected" (Kneale, 1971, 2). Given this assumption, as well as the innocuous assumption that causes necessarily precede their effects, it follows that we are necessarily unable to "remember the future", i.e., that we are necessarily unable to spontaneously anticipate future events, in parallel to the way we spontaneously remember past events. The apparent epistemological difference between our knowledge of the past and our knowledge of the future thus collapses into the apparent metaphysical difference which provides the basis for the causal theory of memory. And, as we saw in section 2.1, while the causal theory is intuitively plausible, there is good reason to reject it.

Kneale's discussion of "direct" memory knowledge is nevertheless instructive, in that it provides an important clue to the source of the appeal of the thought that there must be a fundamental epistemological asymmetry between episodic memory and FMTT. Consider her discussion of alleged cases of "pre-cognition" (1971, 2):

> The pre-cognitions investigated by the Society for Psychical Research do not correspond to memories, because the investigator at least, and generally the percipient too, regards them as subject to confirmation, and no serious scientist or historian would think of such "precognitions" as providing premises for further inferences. ... We have no faculty for recognizing truths about the future analogous to memory.

The point is not that precognition does not exist. The point — and this is independent of the causal theory — is rather that, even if it were to exist, we would be reluctant to regard it as a source of knowledge in the same way that we regard memory as a source of knowledge, because it would always be subject to further confirmation. Similarly, we can observe that, whereas we regularly treat memory, for example in legal contexts, as a more or less decisive source of evidence about what happened, we do not have any equivalent practice of treating FMTT as a source of evidence about what will happen. More generally, when someone claims to remember the occurrence of an event, we often accept his claim at face value, but when someone anticipates the occurrence of an event, we normally want additional confirming evidence.

There is thus an asymmetry between our epistemic practices with respect to episodic memory and our epistemic practices with respect to FMTT. Whether this difference at the level of epistemic *practice* reflects an underlying difference of epistemic *status* is another question. There is good reason to think that it does not. Paradoxically, the practical difference seems to arise because confirming or disconfirming evidence is often available, at least in principle, at the time at which a subject remembers an event, whereas confirming evidence cannot be available at the time at which a subject anticipates an event. The fact that a memory can be confirmed at the time at which it occurs reassures us; we are therefore less inclined to ask for confirmation. The fact that a prediction cannot be confirmed at the time at which it occurs disquiets us; we are therefore more inclined to ask for confirmation. But as far as the epistemic status of the representations respectively produced by episodic memory and FMTT are concerned, there may yet be no qualitative difference between them.

Within a reliabilist framework (e.g., Goldman 2012) — analogous points can be made within most other epistemological frameworks — the epistemic status of a representation is determined simply by the reliability of the process that produced it. If episodic memory and FMTT are both sufficiently reliable for knowledge, they both produce knowledge, even if we have available processes for *confirming* that we have knowledge in the case of episodic memory (e.g., seeking corroborating physical evidence) and no such processs available for confirming that we have knowledge in the case of FMTT. The point is that one can know without being in a position to know that one knows. When, seeming to remember an event, I ask myself whether I know that the event occurred, I often have available some means of answering the question. In contrast, when, anticipating an event, I ask myself whether I know that the event will occur, I normally have available no means of answering the question. But the fact that I cannot know whether I know does not imply that I do not know. Our epistemic practices with respect to epistemic memory and FMTT thus suggest at most a difference between our knowledge of our knowledge of the past and our knowledge of our knowledge of the future, but this does not entail a difference between our knowledge of the past itself and our knowledge of the future itself.

3.3 Immunity to error through misidentification

Perrin (this volume) argues for a final, more subtle epistemological discontinuity between episodic memory and FMTT: on his view, FMTT displays immunity to error through misidentification, and episodic memory does not. Some background is required in order to understand this claim. First, we need the notion of *error through misidentification*. A thought involves an error through misidentification when it is incorrect due to the fact that it attributes an incorrect identity to something or someone. For example, sup-

pose that I hear the phone ring and think "my wife is calling" when, in reality, it is my colleagues who is calling. In this case, my thought correctly describes an event — someone really is calling — but it is nonetheless incorrect overall because I am mistaken about the identity of one of the people involved in the event. Second, we need the notion of *immunity* to error through misidentification. A thought manifests immunity to error through misidentification when it is impossible for it to be incorrect due to the fact that it attributes an incorrect identity to something or someone. My thought that my wife is calling is, obviously, not immune to error through misidentification. Even thoughts about oneself need not be immune to error through misidentification. For example, if I mistake my reflection in a mirror for that of a colleague and think "my colleague is wearing a sportcoat" when in fact I am the one wearing a sportcoat, I commit an error through misidentification. But certain thoughts about oneself are arguably immune to error through misidentification. For example, if I introspect and think "I feel hungry right now", I arguably cannot be mistaken about the identity of the person who feels hunger.

The first part of Perrin's claim, then, is that, when a subject episodically remembers, his thoughts are *not* immune to error through misidentification: it is possible for them to be mistaken due to the fact that they attribute an incorrect identity to something or someone. Errors through misidentification in episodic memory may result from attributions of incorrect identity to oneself,[18] but they may also result from attributions of incorrect identity to other subjects. For example, I might remember talking to a specific colleague but misidentify him, perhaps because I mistook him at the time of the conversation for another colleague. The second part of Perrin's claim is that, when a subject imagines a future episode, his thoughts *are* immune to error through misidentification: it is impossible for them to be mistaken due to the fact that they attribute an incorrect identity to something or someone. The thought here is that, when, for example, I imagine talking to a colleague tomorrow, I cannot attribute an incorrect identity to the colleague in question. One cannot be mistaken about the identities of the entities involved in imagined events, simply because the identities of the entities one imagines are wholly determined by one's intentions. In contrast, one can be mistaken about the identities of the entities involved in remembered events, because the identities of the entities one remembers are not wholly determined by one's intentions. This is due to the fact that episodic memory is "world-involving" (Recanati, 2007): the content of one's memory is partly determined by one's interactions with the world.

Faced with this argument, the epistemological continuist has three options: first, he might maintain that *neither* episodic memory nor FMTT is

[18] As in hypothetical cases of "quasi-memory", in which the memories of one subject are "transplanted" into another subject (Shoemaker, 1970).

immune to error through misidentification; second, he might maintain that *both* episodic memory and FMTT are immune to error through misidentification; third, he might maintain that *some* instances of episodic memory and *some* instances of FMTT are immune to error through misidentification, while others are not. Given the characterizations of FMTT and episodic memory developed in previous sections, the last option seems likely to be correct.

Might some instances of FMTT fail to be immune to error through misidentification? We have already seen why it is plausible to hold that error through misidentification is possible in episodic memory: one's past interactions with the entities involved in the events that one remembers affect the content of one's memories, with the consequence that one's own judgement about the identities of the entities involved in the events that one remembers may be mistaken. The same line of thought, however, suggests that error through misidentification is possible in FMTT. As noted in section 2, one may imagine a future event by remembering a past event and imaginatively projecting that event into the future. In such cases, one clearly imagines events involving entities with which one has previously interacting. Even when one imagines more novel events, one often imagines events involving entities with which one has previously interacted, and if previous interactions suffice to fix identity in cases of episodic memory, they presumably do so as well in cases of FMTT. Consider again a case in which I imagine talking to a specific colleague tomorrow. While it is natural to suppose that my intentions determine the identity of the colleague in question (that I effectively stipulate his identity), this need not be the case. Episodic memory is world-involving, if it is, because the subject has causally interacted with the entities involved in his memories; those past causal interactions can arguably partly determine the content of his memory representations. But the same mechanism allows FMTT to be world-involving: if the subject imagines a conversation with a person with whom he has interacted in the past, the past causal interaction arguably determines the identity of the person of whom he is thinking. Thus when I imagine a conversation with a specific colleague, I might be mistaken about the identity of the colleague in question, perhaps because I mistook him during our past interactions for another colleague. If this is right, then in at least some cases FMTT is not immune to error through misidentification.[19]

[19]The discontinuist might object that the the world-involvingness of FMTT is parasitic on that of episodic memory, since it depends on traces laid down during past causal interactions — that is, on episodic memories. But this objection depends on ignoring the distinction (noted above in connection with alleged functional disanalogies between episodic memory and FMTT) between memory as a store and remembering as a process. As we have seen, the same stored content is employed both in episodic remembering and in FMTT; there is no convincing reason to assign any priority to the former process, rather than the latter.

This argument only shows that *some* instances of FMTT are not immune to error through misidentification; if other instances of FMTT are immune to error through misidentification and no instances of episodic memory are immune to error through misidentification, the discontinuist can maintain that there is an important epistemological asymmetry between episodic memory and FMTT. But might some instances of episodic memory be immune to error through misidentification? Perrin's argument tacitly assumes that episodic memory always satisfies the factivity constraint: genuine memory is always memory of events that actually occurred. But we saw in section 2 that memory of events that did not actually occur has just as much claim to be genuine memory as memory of events that actually occurred, i.e., that episodic memory does not always satisfy the factivity constraint. In cases where the factivity constraint is not satisfied, the subject need not have interacted with the entities involved in the event he remembers. Thus there are no causal interactions available to fix the identities of those entities, and their identities are presumably fixed by the subject's intentions: just as, when I imagine a wholly novel future event, my intentions determine the identities of the things and persons involved, when I remember a wholly novel past event, my intentions determine the identities of the things and persons involved.

Thus it appears that both parts of Perrin's claim are to be rejected. It is not the case that, when a subject episodically remembers, his thoughts are never immune to error through misidentification; in some cases, they may not be, but in other cases they may be. And it is not the case that, when a subject imagines a future episode, his thoughts are always immune to error through misidentification; in some cases, they may be, but in other cases they may not be. In short, episodic memory and FMTT appear to be on a par, as far as immunity to error through misidentification is concerned.

4 Towards a positive case for epistemological continuism

This chapter has largely been devoted to responding to the available arguments for metaphysical and epistemological continuism. Supposing that the response has been successful in undermining those arguments, should we move to metaphysical and epistemological continuism? The focus of positive continuist arguments so far has been on metaphysical continuism, with psychologists and philosophers appealing to mental time travel research to support a picture of remembering the past and imagining the future as qualitatively similar constructive processes. What remains to be provided, then, is a positive case for epistemological continuism, the view that our knowledge of future events is qualitatively similar to our knowledge of past events.

Such a case may be less difficult to develop than one might initially suspect. Against the background of metaphysical continuism, Michaelian (forthcoming) develops a reliabilist account of our knowledge of past events, arguing that episodic memory provides reliable access to the personal past despite its simulational character. The account involves two core ideas: first, that mental simulation of past events is governed by reliable heuristics; second, that the threat of unreliability introduced by the simulational character of remembering is further reduced by reliable metacognitive monitoring. It may be possible to extend this account to our knowledge of future events in a straightforward manner. Once we give up the idea that remembering necessarily involves a causal connection to the relevant past event (i.e., the causal theory) and the idea that remembering puts us in direct contact with the relevant past event (i.e., direct realism), it becomes clear that memory knowledge, in line with other forms of indirect, inferential knowledge, is underwritten entirely by the reliability of the process of remembering. FMTT, of course, involves neither a causal connection to nor direct contact with the relevant future event, which suggests that here, too, knowledge is underwritten entirely by the reliability of the process of imagining. Thus our knowledge of the personal future need pose no special epistemological challenge: we acquire knowledge of future events, just as we acquire knowledge of past events, through reliable simulation and reliable metacognitive monitoring.

Such a reliabilist account of our knowledge of the personal future dovetails nicely with recent responses to the charge that imagination is "epistemically irrelevant" (Kind forthcoming) (cf. Balcerak Jackson forthcoming). As Kind points out, philosophers such as Sartre (1948) and Wittgenstein (1980) have argued that the imagination is fundamentally epistemically irrelevant: while it can play an incidental role in the generation of knowledge, for example by leading us to formulate hypotheses which can then be verified by other means, it cannot itself generate new knowledge. Kind identifies three assumptions which have been used to ground this conclusion: first, imagination is typically under our voluntary control; second, it is not world-sensitive, in the sense that "its content is determined by the imaginer, not by the world"; third, imagination is uninformative, in the sense that "an act of imagining can provide us with no new information".

Kind grants the first two assumptions but argues that they do not imply that imagination cannot aim at the truth. The epistemological continuist might also challenge the first two assumptions more directly. As far as the first assumption is concerned, much FMTT is spontaneous or involuntary (Berntsen and Jacobsen, 2008). As far as the second assumption is concerned, this amounts to a version of the claim that FMTT is not world-involving, and we have already seen that this claim can be challenged. But Kind focuses her resistance to the charge of epistemic irrelevance on the third assumption, arguing that imagination can indeed provide the imag-

iner with new information. The thought behind the third assumption is that a representation produced by imagination does not contain anything that the imaginer did not put into it, which suggests that it cannot provide him with new information. Kind points out that this is simply mistaken. A computer program, for example, contains nothing that the programmer did not put into it, and yet it can of course provide new information when it is run. The same thing goes for imagination: in episodic imagination, one simulates an event, but the simulation can provide one with new information. Consider, e.g., the case of mindreading or theory of mind: when one simulates the mental states of another subject, the simulation contains nothing that one did not put into it but can nevertheless, assuming that the simulation process is reliable, provide one with new information about the target's mental states. Similarly, when one simulates one's own possible future experience, the simulation will contain nothing that one did not put into it but can nevertheless, assuming that the simulation process is reliable, provide one with new information about those experiences.

This is only a sketch of a positive account, but it is enough to suggest that the reliabilist explanation of how episodic memory provides us with knowledge of past events can be extended to explain how FMTT provides us with knowledge of future events.[20] The future of continuism about our knowledge of past and future events thus looks bright.

References

Balcerak Jackson, M. (Forthcoming). Justification by imagination. In Dorsch, F. and Macpherson, F., editors, *Perceptual Memory and Perceptual Imagination*. Oxford University Press.

Bernecker, S. (2008). *The Metaphysics of Memory*. Springer.

Berntsen, D. and Jacobsen, A. S. (2008). Involuntary (spontaneous) mental time travel into the past and future. *Consciousness and Cognition*, 17(4):1093–1104.

Corballis, M. C. (This volume). The future of memory, mental time travel, and mind wandering.

[20] An alternative strategy for securing the epistemological continuity of episodic memory and FMTT is to grant the epistemic irrelevance of FMTT but to argue that episodic memory is also epistemically irrelevant. Hopkins, for example, argues that "[i]f in memory we represent the past to ourselves as being a certain way, and if representing it so is an action of ours, it seems we can only get the past right if we in some way already know how it was. Remembering cannot be the source of knowledge of how things were, since it presupposes such knowledge" (2014, 321); this amounts to the suggestion that episodic memory cannot provide us with new information. Those who acknowledge the generative, simulational character of remembering will not be tempted to avail themselves of this argument.

D'Argembeau, A. and Van der Linden, M. (2004). Phenomenal characteristics associated with projecting oneself back into the past and forward into the future: Influence of valence and temporal distance. *Consciousness and Cognition*, 13(4):844–858.

De Brigard, F. (2014). Is memory for remembering? Recollection as a form of episodic hypothetical thinking. *Synthese*, 191(2):155–185.

De Brigard, F., Addis, D., Ford, J., Schacter, D., and Giovanello, K. (2013). Remembering what could have happened: Neural correlates of episodic counterfactual thinking. *Neuropsychologia*, 51(12):2401–2414.

De Brigard, F. and Gessell, B. S. (This volume). Time is not of the essence: Understanding the neural correlates of mental time travel.

Debus, D. (2008). Experiencing the past: A relational account of recollective memory. *Dialectica*, 62(4):405–432.

Debus, D. (2014). "Mental time travel": Remembering the past, imagining the future, and the particularity of events. *Review of Philosophy and Psychology*, 5(3):333–350.

Debus, D. (This volume). Temporal perspectives in imagination: On the nature and value of imagining the future.

Goldman, A. I. (2012). *Reliabilism and Contemporary Epistemology: Essays*. Oxford University Press, Oxford.

Hoerl, C. (2014). Remembering events and remembering looks. *Review of Philosophy and Psychology*, 5(3):351–372.

Hopkins, R. (2014). Episodic memory as representing the past to oneself. *Review of Philosophy and Psychology*, 5(3):313–331.

Hume, D. (1739). *A Treatise of Human Nature*. John Noon, London.

Irish, M. (This volume). Semantic memory as the essential scaffold for future oriented mental time travel.

Kind, A. (Forthcoming). How imagination gives rise to knowledge. In Dorsch, F. and Macpherson, F., editors, *Perceptual Memory and Perceptual Imagination*. Oxford University Press.

Klein, S. B. (2013). The complex act of projecting oneself into the future. *Wiley Interdisciplinary Reviews: Cognitive Science*, 4(1):63–79.

Klein, S. B. (2015). What memory is. *Wiley Interdisciplinary Reviews: Cognitive Science*, 6(1):1–38.

Klein, S. B. (Forthcoming). Autonoetic awareness: Reconsidering the role of episodic memory in future-oriented mental time travel. *Quarterly Journal of Experimental Psychology*.

Klein, S. B. and Nichols, S. (2012). Memory and the sense of personal identity. *Mind*, 121(483):677–702.

Klein, S. B. and Steindam, C. (This volume). The role of subjective temporality in future-oriented mental time travel.

Kneale, M. (1971). Our knowledge of the past and of the future. *Proceedings of the Aristotelian Society*, 72:1–12.

Loftus, E. (1996). *Eyewitness Testimony*. Harvard University Press, Cambridge, MA, second edition.

Martin, C. B. and Deutscher, M. (1966). Remembering. *Philosophical Review*, 75(2):161–196.

Martin-Ordas, G. (This volume). With the future in mind: towards a comprehensive understanding of the evolution of future-oriented cognition.

Michaelian, K. (2011). Generative memory. *Philosophical Psychology*, 24(3):323–342.

Michaelian, K. (2012). (Social) metacognition and (self-)trust. *Review of Philosophy and Psychology*, 3(4):481–514.

Michaelian, K. (Forthcoming). *Mental Time Travel: Episodic Memory and Our Knowledge of the Personal Past*. MIT Press.

Perrin, D. (This volume). Asymmetries in subjective time.

Pezzulo, G. (This volume). The mechanisms and benefits of a future-oriented brain.

Recanati, F. (2007). *Perspectival Thought: A Plea for (Moderate) Relativism*. Oxford University Press, Oxford.

Reid, T. (1764). *An Inquiry into the Human Mind*. Printed for Alexander Ewing, Dublin.

Russell, B. (1921). *The Analysis of Mind*. George Allen & Unwin, London.

Sartre, J.-P. (1948). *The Psychology of Imagination*. Philosophical Library.

Schacter, D. L. and Addis, D. R. (2007). The cognitive neuroscience of constructive memory: Remembering the past and imagining the future. *Philosophical Transactions of the Royal Society of London B: Biological Sciences*, 362(1481):773–786.

Schacter, D. L., Addis, D. R., Hassabis, D., Martin, V. C., Spreng, R. N., and Szpunar, K. K. (2012). The future of memory: Remembering, imagining, and the brain. *Neuron*, 76(4):677–694.

Shoemaker, S. (1970). Persons and their pasts. *American Philosophical Quarterly*, 7(4):269–285.

Suddendorf, T. and Corballis, M. C. (2007). The evolution of foresight: What is mental time travel, and is it unique to humans? *Behavioral and Brain Sciences*, 30:299–313.

Sutton, J. and Michaelian, K. (Forthcoming). Memory. In Zalta, E. N., editor, *Stanford Encyclopedia of Philosophy*.

Swinburne, R. G. (1966). Knowledge of past and future. *Analysis*, 26(5):166–172.

Szpunar, K. K. (2010). Episodic future thought: An emerging concept. *Perspectives on Psychological Science*, 5(2):142–162.

Szpunar, K. K., Spreng, R. N., and Schacter, D. L. (This volume). Toward a taxonomy of future thinking.

Szpunar, K. K., Watson, J. M., and McDermott, K. B. (2007). Neural substrates of envisioning the future. *Proceedings of the National Academy of Sciences*, 104(2):642–647.

Thom, J. M. and Clayton, N. S. (This volume). Evolutionary perspectives on prospective cognition.

Tulving, E. (2002). Chronesthesia: Conscious awareness of subjective time. In Knight, R. T., editor, *Principles of Frontal Lobe Function*, pages 311–325. Oxford University Press, Oxford.

Wittgenstein, L. (1980). *Remarks on the Philosophy of Psychology*. University of Chicago Press.